

---

# BMapEst: Estimation of Brain Tissue Probability Maps using a Differentiable MRI Simulator

---

Utkarsh Gupta<sup>1</sup> Emmanouil Nikolakakis<sup>2</sup> Moritz Zaiss<sup>3,4</sup> Razvan Marinescu<sup>1</sup>

## Abstract

Reconstructing digital brain phantoms in the form of voxel-based, multi-channeled tissue probability maps for individual subjects is essential for capturing brain anatomical variability, understanding neurological diseases, as well as for testing image processing methods. We demonstrate the first framework that estimates brain tissue probability maps (Grey Matter - GM, White Matter - WM, and Cerebrospinal fluid - CSF) with the help of a Physics-based differentiable MRI simulator that models the magnetization signal at each voxel in the volume. Given an observed  $T_1/T_2$ -weighted MRI scan, the corresponding clinical MRI sequence, and the MRI differentiable simulator, we estimate the simulator’s input probability maps by back-propagating the L2 loss between the simulator’s output and the  $T_1/T_2$ -weighted scan. This approach has the significant advantage of not relying on any training data and instead uses the strong inductive bias of the MRI simulator. We tested the model on 20 scans from the BrainWeb database and demonstrated a highly accurate reconstruction of GM, WM, and CSF. Our source code is available online: <https://github.com/BioMedAI-UCSC/BMapEst>.

## 1. Introduction

Quantitative  $T_1$  ( $qT_1$ ) and  $T_2$  ( $qT_2$ ) MRI plays a pivotal role in understanding various neurological conditions, including multiple sclerosis [1], brain tumors [2], and neurodegenerative diseases such as Alzheimer’s [3], by providing valuable

insights into tissue characteristics and pathological changes. However, even more information than  $T_1/T_2$ -weighted MRI can be obtained by reconstructing the tissue probability maps of the digital brain phantoms. Such multi-channeled tissue probability maps typically include the percentage of a particular tissue (GM, WM, CSF, blood vessels, glia cells, etc.) that exists at each 3D voxel and thus give significantly more information than standard  $T_1/T_2$ -weighted MRIs. In addition, probability maps are critical as priors in accurate self-supervised semantic segmentation or classification of a 3D volume reconstruction. They also help regularize and eliminate artifacts such as partial volume effects, motion artifacts, or bias field artifacts, which typically require external atlas-based priors [4], [5]. BrainWeb phantoms [6], [7] are popular examples that store probability maps for 11 tissues at each voxel, offering more precise anatomical detail than  $qT_1/qT_2$  maps, customizable tissue properties, and the ability to simulate realistic imaging scenarios.

A related problem to tissue probability map estimation is tissue segmentation, where a single label is assigned for each voxel. Current state-of-the-art methods for tissue segmentation use supervised deep-learning [8], [9] and clustering methods [10]–[13]. However, tissue segmentation is a different and easier problem than probability map estimation since models don’t need to estimate partial volume effects at all voxels. Despite their utility, we are unaware of any methods that fully infer tissue probability maps from  $T_1/T_2$ -weighted scans. Additionally, the problem is particularly challenging due to its ill-posedness at two different levels: 1) multiple probability map combinations can produce the same  $qT_1/qT_2$ /PD maps, and 2) multiple  $qT_1/qT_2$ /PD maps can produce the same  $T_1/T_2$ -weighted image.

MRI imaging simulators [14]–[26] offer a solution for inferring probability maps of digital brain phantoms. They act as a forward model to generate k-space measurements and subsequently reconstruct  $T_1/T_2$ -weighted scans. This represents a paradigm shift in medical imaging simulation and optimization. By leveraging simulators’ inherent physics-based inductive biases, it becomes feasible to fine-tune the parameters of clinical sequences, leading to enhanced image quality. A recent innovation in the field is the building of MR-zero [27], a *differentiable* MRI simulator implement-

---

<sup>1</sup>University of California, Santa Cruz, Computer Science and Engineering Department <sup>2</sup>University of California, Santa Cruz, Electrical and Computer Engineering Department <sup>3</sup>Department of Artificial Intelligence in Biomedical Engineering, Friedrich-Alexander-Universität Erlangen <sup>4</sup>Department of Neuroradiology, Universitätsklinik Erlangen, Erlangen, Germany. Correspondence to: Utkarsh Gupta <utgupta@ucsc.edu>.

Published at the 2<sup>nd</sup> Differentiable Almost Everything Workshop at the 41<sup>st</sup> International Conference on Machine Learning, Vienna, Austria. July 2024. Copyright 2024 by the author(s).

ing the Phase Distribution Graph [28] (Appendix A), thus alleviating the need for slow derivative-free optimizations.

In this paper, we introduce a novel framework leveraging the capabilities of MR-zero to estimate the tissue probability maps of digital brain phantoms (Fig. 2) representing the CSF, GM, and WM. Unlike supervised learning, our framework does not need a-priori training pairs of inputs and outputs and can further estimate tissue probability maps for any arbitrary set of MRI sequences (e.g.,  $T_1$ -only,  $T_1+T_2$ ,  $T_2+T_2^*+GRE$ , etc.) at any arbitrary echo times. Our contributions are:

- We demonstrate the first method of estimating brain tissue probability maps using a differentiable MRI simulator that conducts forward inference to generate a  $T_1/T_2$ -weighted image by backpropagating a loss function to the brain tissue probability maps. Our approach is versatile, applicable to many different MRI sequences, and does not require learnable parameters.
- We overcome the ill-posedness of probability maps estimation by using the inductive bias of the simulator and multiple  $T_1/T_2$  contrasts.
- We validate our approach on BrainWeb’s 20 subjects with the popular Fast Low Angle Shot (FLASH) sequence variants and obtain state-of-the-art results compared to supervised deep learning and clustering methods.

## 2. Methodology

Our framework optimizes brain tissue probability maps given one or more observed  $T_1/T_2$ -weighted MRI scans ( $G$ ) and the corresponding MRI clinical sequence ( $S$ ) used to generate each scan (Fig. 1). The optimization pipeline begins by transforming the brain tissue probability maps into  $qT_1$ ,  $qT_2$ , and Proton Density (PD) maps with the help of known relaxation time values for a 1.5T scanner [29]. These maps serve as input for the MR-zero (Appendix B) MRI differentiable simulator that produces K-space measurements as its output of the forward pass. By applying the Inverse Fourier Transform ( $I$ ), the simulated  $T_1/T_2$ -weighted scan ( $G'$ ) is reconstructed from K-space. Therefore, the end-to-end forward pass is given as follows:

$$G' = I(MRZero(f(GM, WM, CSF), S)) \quad (1)$$

where  $f$  is a fixed linear transformation mapping tissue probability maps to  $qT_1/qT_2/PD$  maps using known values from the literature. The pixel-wise L2 loss  $\|G - G'\|_2^2$  is calculated between the observed  $T_1/T_2$ -weighted scan  $G'$  and the ground truth scan  $G$  and backpropagated to the input probability maps (GM/WM/CSF).

## 3. Experiments

### 3.1. Dataset

We demonstrate our method on the BrainWeb dataset, which comprises brain tissue probability maps from 20 subjects. We downsampled the 3D brain volumes to 64x64 (higher resolutions were not possible on our machine due to RAM requirements exceeding 64GB) and selected the middle slice, as it contains all three types of tissues, resulting in CSF, GM, and WM probability maps. Due to the simulator’s compute-intensive nature, we only select a single slice, but our method is equally generalizable to other slices of the volume as well.

### 3.2. Ill-posed problem

The estimation of tissue probability maps presents a challenge due to the inherent ambiguity stemming from the ill-posed nature of the problem, wherein multiple potential solutions (tissue probabilities) exist for a single observed  $T_1/T_2$ -weighted scan. This can be seen in Fig. 2, where the optimization of all tissue probability maps of a subject with only a single observed contrast obtained from the  $T_1$  inversion recovery sequence give a blurry estimation of tissue maps. Furthermore, the sensitivity of tissue delineation is influenced by the specific MRI clinical sequence utilized. For instance, the Fluid Attenuated Inversion Recovery (FLAIR) sequence effectively diminishes the signal from CSF, making CSF map estimation from this sequence highly ill-posed. Unlike supervised deep learning methods, the flexibility of our method allows us to augment the optimization process by incorporating additional output contrasts obtained from capturing images at varying echo times and by including a diverse array of sequences such as  $T_2$ -weighted imaging,  $T_2^*$  imaging through echo time alterations, FLAIR, Double Inversion Recovery (DIR), and Diffusion Weighted Imaging (DWI) sequences.

### 3.3. Setup

Our optimization process spans 501 epochs with a learning rate of 0.01, and takes 5-6min for a single contrast and under 3 hours for 24 contrasts (6 sequences with 4 echo times each). We concurrently optimize all probability maps (CSF, GM, and WM), leveraging the output contrasts from up to six different MRI clinical sequences. Loss computation is conducted across reconstructed image spaces, representing the output  $T_1/T_2$ -weighted scans (each with 4 contrasts per sequence). This is the baseline configuration settings for comparing the results. We performed leave-one-out cross-validation. Our approach adds no extra learnable parameters beyond the probability maps themselves.

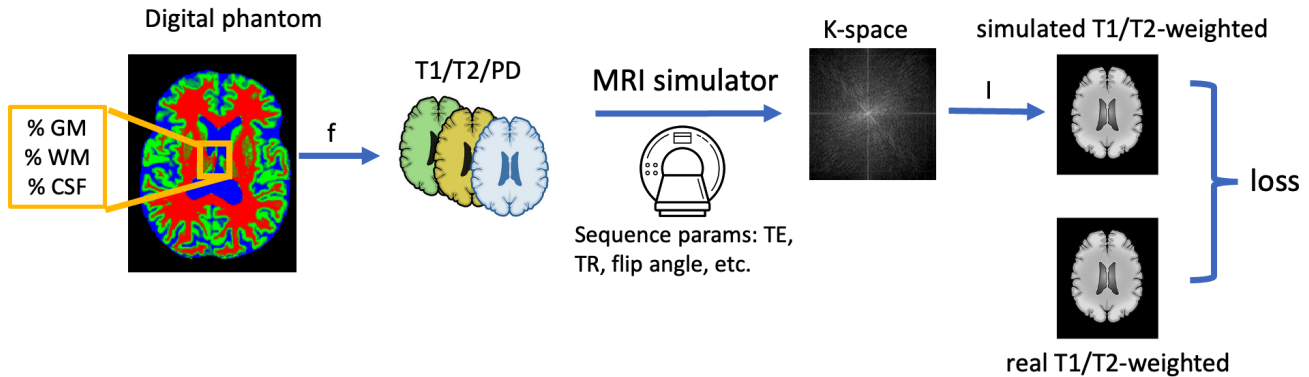


Figure 1: Optimization pipeline: The CSF, GM, and WM probability maps are converted into  $qT_1$ ,  $qT_2$ , and PD maps using known  $T_1$  and  $T_2$  relaxation times. These maps are input into the MR-zero simulator to generate K-space measurements, and the image is reconstructed via Inverse Fourier Transform  $I$ . The L2 loss is computed against a real MRI and backpropagated to the probability maps.

## 4. Results

As illustrated in Table 1 and in Fig. 2, our method achieves very good reconstructions when at least 4 contrasts (echoes) from a single sequence are observed. Only when one single contrast is used does our method perform poorly due to the strongly ill-posed problem. Best reconstructions are achieved using 5 sequences together:  $T_1$ ,  $T_2$ ,  $T_2^*$ , double inversion recovery and FLAIR.

In addition, we also compared against a supervised U-Net baseline [30]. We trained the U-Net by creating pixel-wise normal distributions of probability maps from the middle slices of ten (5, 18, 38, 42, 44, 46, 48, 50, 52, 54) BrainWeb subjects. The input to the U-Net is a 24-channelled MRI contrasts tensor, and the output is a 3-channelled CSF, GM, and WM probability maps. Since training a U-Net is difficult in this low-data regime, we augmented the dataset by building pixel-wise normal distributions and isotropically sampling 500 images from the distribution. We tested the model on the remaining subjects and the same brain slices of training data subjects on which our method’s metrics are reported. Table 1 shows that our method has better PSNR and SSIM than the U-Net across all tissues. Also, our method has a better DICE score than the U-Net on CSF and a similar DICE score on GM and WM.

In Table 1, we also tested the BCEFCM [31], a fuzzy C-means clustering method. BCEFCM performed sub-optimally for the DICE score and poorly on PSNR and SSIM metrics, which is evident from Appendix Figure 4, as clustering-based methods cannot estimate the exact probability values. Moreover, we have taken Table 2 from [10] which shows the DICE metrics obtained for CSF, GM, and WM of the BrainWeb dataset for 4 other clustering methods: LNLFCM [11], DCT-LNLFCM [12], KWFLICM [13], and Double Estimation [10] based fuzzy C-means. Though the slices used for the metric would differ, they are from the

same BrainWeb dataset on which we have reported our results. Our method outperforms all the other techniques for GM and CSF and performs equally well for WM estimation.

For visual results, refer to Figure 4 in the Appendix for inferred probability map results for our baseline, U-Net, and BCEFCM methods. Also, Appendix C and D show detailed experiments and other visual results respectively.

## 5. Conclusion

We have developed BMapEst, the first framework that estimates digital brain phantoms using an MRI differentiable simulator. In future work, we plan to estimate probability maps on large brain MRI datasets such as OASIS or ADNI. Furthermore, in ill-posed settings, our framework can also be extended using variational inference to capture a *distribution* over the tissue probability maps that generate the same  $T_1/T_2$ -weighted MRI scan. This will allow us to capture the variability in tissue characteristics across different individuals having similar  $T_1/T_2$ -weighted scans. Another direction we plan to pursue is to optimize the MRI clinical sequence parameters towards a downstream goal, such as better image contrast, better segmentation of structures, or better disease detection. By leveraging the differentiable MRI simulators, we can efficiently and accurately model the impact of various sequence parameters on the resulting images, enabling us to fine-tune these parameters to enhance diagnostic utility and achieve optimal imaging sequence.

One of the limitations of differentiable simulators is their high computational requirements during optimization, in particular GPU memory. To address this, we plan to investigate the use of advanced optimization techniques and hardware acceleration, such as leveraging GPUs and CPU RAM in tandem to balance memory consumption, reduce computational load, and improve processing speed. To speed-up

Table 1: Dice score, PSNR, and SSIM comparisons for a combination of MRI clinical sequences. Except for the first row, all the experiments use 4 contrasts per sequence. † represents the baseline configuration. Our method is compared against two baselines: **U-Net** and **BCEFCM**. **Note: BCEFCM** can only take a single contrast as input.

Contrasts	CSF			GM			WM		
	Dice↑	PSNR ↑	SSIM ↑	Dice↑	PSNR ↑	SSIM ↑	Dice↑	PSNR ↑	SSIM ↑
$T_1$ single contrast	0.38±0.07	17.84±1.18	0.58±0.07	0.67±0.05	14.22±0.97	0.57±0.08	0.62±0.05	12.83±1.1	0.38±0.07
	<b>0.31±0.05</b>	<b>18.29±0.86</b>	<b>0.63±0.05</b>	<b>0.66±0.04</b>	<b>14.94±0.74</b>	<b>0.64±0.05</b>	<b>0.65±0.04</b>	<b>13.5±0.83</b>	<b>0.51±0.06</b>
	<b>0.04±0.01</b>	<b>11.47±0.4</b>	<b>0.13±0.02</b>	<b>0.40±0.02</b>	<b>7.53±0.27</b>	<b>0.16±0.02</b>	<b>0.38±0.04</b>	<b>7.68±0.27</b>	<b>0.18±0.02</b>
$T_1$ four contrasts	0.55±0.04	27.57±2.11	0.89±0.03	0.7±0.04	16.16±0.97	0.74±0.05	0.76±0.02	18.15±0.98	0.78±0.04
	<b>0.43±0.04</b>	<b>23.57±0.85</b>	<b>0.86±0.02</b>	<b>0.72±0.03</b>	<b>16.52±0.53</b>	<b>0.76±0.02</b>	<b>0.76±0.02</b>	<b>17.38±0.48</b>	<b>0.75±0.03</b>
$T_1+T_2$	0.55±0.04	30.9±3.93	0.97±0.02	0.76±0.03	23.5±1.98	0.93±0.03	0.79±0.01	25.63±2.43	0.93±0.03
	<b>0.43±0.04</b>	<b>22.63±0.74</b>	<b>0.85±0.02</b>	<b>0.76±0.02</b>	<b>19.19±0.34</b>	<b>0.87±0.01</b>	<b>0.79±0.02</b>	<b>19.71±0.34</b>	<b>0.83±0.02</b>
$T_1+T_2+T_2^*$	<b>0.56±0.03</b>	33.43±4.72	0.97±0.02	0.76±0.03	26.38±2.35	0.96±0.03	0.79±0.01	25.16±2.34	0.91±0.03
	<b>0.47±0.04</b>	<b>24.08±0.8</b>	<b>0.88±0.02</b>	<b>0.76±0.02</b>	<b>20.18±0.36</b>	<b>0.9±0.01</b>	<b>0.79±0.02</b>	<b>19.91±0.35</b>	<b>0.83±0.02</b>
$T_1+T_2+T_2^*+DIR$	<b>0.56±0.03</b>	33.68±4.51	0.97±0.02	0.76±0.03	26.33±2.38	0.96±0.03	0.79±0.01	25.16±2.45	0.91±0.04
	<b>0.45±0.04</b>	<b>24.33±0.91</b>	<b>0.89±0.02</b>	<b>0.76±0.02</b>	<b>20.48±0.37</b>	<b>0.91±0.01</b>	<b>0.79±0.02</b>	<b>20.49±0.36</b>	<b>0.84±0.02</b>
$T_1+T_2+T_2^*+DIR+FLAIR$	0.55±0.04	<b>34.64±5.29</b>	<b>0.98±0.02</b>	0.76±0.02	<b>27.71±2.48</b>	<b>0.97±0.02</b>	<b>0.8±0.01</b>	<b>27.0±2.69</b>	<b>0.93±0.03</b>
	<b>0.45±0.04</b>	<b>24.08±0.78</b>	<b>0.89±0.02</b>	<b>0.77±0.02</b>	<b>20.56±0.38</b>	<b>0.9±0.01</b>	<b>0.79±0.02</b>	<b>20.34±0.39</b>	<b>0.84±0.18</b>
$T_1+T_2+T_2^*+DIR+FLAIR+DWI$ †	0.55±0.04	34.45±4.87	<b>0.98±0.02</b>	0.76±0.02	26.8±2.54	0.96±0.02	0.79±0.01	25.99±2.69	0.92±0.03
	<b>0.45±0.04</b>	<b>24.23±0.85</b>	<b>0.89±0.07</b>	<b>0.76±0.02</b>	<b>20.38±0.37</b>	<b>0.9±0.01</b>	<b>0.79±0.02</b>	<b>20.42±0.39</b>	<b>0.85±0.02</b>




Figure 2: From the left are CSF, GM, and WM. The first row shows the raw tissue probability maps for subject 42 in the Brain-Web dataset. The second row shows the estimated probability maps using our baseline configuration. The third row shows the ill-posed estimation using the single contrast from the  $T_1$  inversion recovery sequence.

Table 2: Dice score, PSNR, and SSIM comparisons with other clustering-based methods. † represents the baseline configuration.

Method	CSF	GM	WM
Our †	<b>0.55±0.04</b>	<b>0.76±0.02</b>	0.79±0.01
Double Estimation	0.41±0.31	0.60±0.37	0.80±0.23
LNLFCM	0.41±0.32	0.58±0.38	0.79±0.23
DCT-LNLFCM	0.41±0.31	0.60±0.36	<b>0.80±0.22</b>
KWFLICM	0.40±0.27	0.59±0.37	0.80±0.23

the convergence, one can also run a pre-trained supervised deep-learning method to get an initial probability map estimate, which can then be further refined using our method.

## References

- [1] C. Granziera, J. Wuerfel, F. Barkhof, *et al.*, “Quantitative magnetic resonance imaging towards clinical application in multiple sclerosis,” *Brain*, vol. 144, no. 5, pp. 1296–1311, 2021.
- [2] I. Blystad, J. M. Warntjes, Ö. Smedby, P. Lundberg, E.-M. Larsson, and A. Tisell, “Quantitative mri for analysis of peritumoral edema in malignant gliomas,” *PLoS One*, vol. 12, no. 5, e0177135, 2017.
- [3] S. Leandrou, S. Petroudi, P. A. Kyriacou, C. C. Reyes-Aldasoro, and C. S. Pattichis, “Quantitative mri brain studies in mild cognitive impairment and alzheimer’s disease: A methodological review,” *IEEE reviews in biomedical engineering*, vol. 11, pp. 97–111, 2018.
- [4] G. Schramm and C. N. Ladefoged, “Metal artifact correction strategies in mri-based attenuation correction in pet/mri,” *BJR—Open*, vol. 1, no. 1, p. 2019033, 2019.
- [5] Y. Wang, M. Zhan, A. Roebroek, P. De Weerd, S. Kashyap, and M. J. Roberts, “Inconsistencies in atlas-based volumetric measures of the human nucleus basalis of meynert: A need for high-resolution alternatives,” *NeuroImage*, vol. 259, p. 119421, 2022.
- [6] C. A. Cocosco, “Brainweb: Online interface to a 3d mri simulated brain database.” (No Title), 1997.
- [7] B. Aubert-Broche, A. C. Evans, and L. Collins, “A new improved version of the realistic digital brain phantom,” *NeuroImage*, vol. 32, no. 1, pp. 138–145, 2006.
- [8] K. V. Kumari and S. S. Barpanda, “Residual unet with dual attention—an ensemble residual unet with dual attention for multi-modal and multi-class brain mri segmentation,” *International Journal of Imaging Systems and Technology*, vol. 33, no. 2, pp. 644–658, 2023.
- [9] X. Zhang, Y. Liu, Q. Zhang, and F. Yuan, “Multi-modality reconstruction attention and difference enhancement network for brain mri image segmentation,” *IEEE Access*, vol. 10, pp. 31058–31069, 2022.
- [10] M. Tavakoli-Zaniani, Z. Sedighi-Maman, and M. H. F. Zarandi, “Segmentation of white matter, grey matter and cerebrospinal fluid from brain mr images using a modified fcm based on double estimation,” *Biomedical Signal Processing and Control*, vol. 68, p. 102615, 2021.
- [11] J. Wang, J. Kong, Y. Lu, M. Qi, and B. Zhang, “A modified fcm algorithm for mri brain image segmentation using both local and non-local spatial constraints,” *Computerized medical imaging and graphics*, vol. 32, no. 8, pp. 685–698, 2008.
- [12] C. Singh and A. Bala, “A dct-based local and non-local fuzzy c-means algorithm for segmentation of brain magnetic resonance images,” *Applied Soft Computing*, vol. 68, pp. 447–457, 2018.
- [13] M. Gong, Y. Liang, J. Shi, W. Ma, and J. Ma, “Fuzzy c-means clustering with local information and kernel metric for image segmentation,” *IEEE transactions on image processing*, vol. 22, no. 2, pp. 573–584, 2012.
- [14] J. Bittoun, J. Taquin, and M. Sauzade, “A computer algorithm for the simulation of any nuclear magnetic resonance (nmr) imaging method,” *Magnetic Resonance Imaging*, vol. 2, no. 2, pp. 113–120, 1984.
- [15] R. M. Summers, L. Axel, and S. Israel, “A computer simulation of nuclear magnetic resonance imaging,” *Magnetic resonance in medicine*, vol. 3, no. 3, pp. 363–376, 1986.
- [16] P. Shkarin and R. G. Spencer, “Direct simulation of spin echoes by summation of isochromats,” *Concepts in Magnetic Resonance*, vol. 8, no. 4, pp. 253–268, 1996.
- [17] P. Shkarin and R. G. Spencer, “Time domain simulation of fourier imaging by summation of isochromats,” *International journal of imaging systems and technology*, vol. 8, no. 5, pp. 419–426, 1997.
- [18] R.-S. Kwan, A. C. Evans, and G. B. Pike, “Mri simulation-based evaluation of image-processing and classification methods,” *IEEE transactions on medical imaging*, vol. 18, no. 11, pp. 1085–1097, 1999.
- [19] D. A. Yoder, Y. Zhao, C. B. Paschal, and J. M. Fitzpatrick, “Mri simulator with object-specific field map calculations,” *Magnetic resonance imaging*, vol. 22, no. 3, pp. 315–328, 2004.
- [20] H. Benoit-Cattin, G. Collewet, B. Belaroussi, H. Saint-Jalmes, and C. Odet, “The simri project: A versatile and interactive mri simulator,” *Journal of Magnetic Resonance*, vol. 173, no. 1, pp. 97–115, 2005.
- [21] T. H. Jochimsen, A. Schäfer, R. Bammer, and M. E. Moseley, “Efficient simulation of magnetic resonance imaging with bloch–torrey equations using intra-voxel magnetization gradients,” *Journal of Magnetic Resonance*, vol. 180, no. 1, pp. 29–38, 2006.
- [22] T. Stöcker, K. Vahedipour, D. Pflugfelder, and N. J. Shah, “High-performance computing mri simulations,” *Magnetic resonance in medicine*, vol. 64, no. 1, pp. 186–193, 2010.
- [23] K. G. Baum, G. Menezes, and M. Helguera, “Simulation of high-resolution magnetic resonance images on the ibm blue gene/l supercomputer using simri,” *Journal of Biomedical Imaging*, vol. 2011, pp. 1–8, 2011.
- [24] C. G. Xanthis, I. E. Venetis, A. Chalkias, and A. H. Aletras, “Mrisimul: A gpu-based parallel approach to mri simulations,” *IEEE Transactions on Medical Imaging*, vol. 33, no. 3, pp. 607–617, 2013.
- [25] C. G. Xanthis, I. E. Venetis, and A. H. Aletras, “High performance mri simulations of motion on multi-gpu systems,” *Journal of Cardiovascular Magnetic Resonance*, vol. 16, no. 1, pp. 1–15, 2014.
- [26] Z. Cao, S. Oh, C. T. Sica, *et al.*, “Bloch-based mri system simulator considering realistic electromagnetic fields for calculation of signal, noise, and specific absorption rate,” *Magnetic resonance in medicine*, vol. 72, no. 1, pp. 237–247, 2014.
- [27] A. Loktyushin, K. Herz, N. Dang, *et al.*, “Mrzero-automated discovery of mri sequences using supervised learning,” *Magnetic Resonance in Medicine*, vol. 86, no. 2, pp. 709–724, 2021.
- [28] J. Endres, S. Weinmüller, H. N. Dang, and M. Zaiss, “Phase distribution graphs for fast, differentiable, and spatially encoded bloch simulations of arbitrary mri sequences,” *Magnetic Resonance in Medicine*, 2023. DOI: 10.1002/mrm.30055.
- [29] *Relaxation times for brain tissues*, <https://mri-q.com/why-is-t1--t2.html>.
- [30] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, Springer, 2015, pp. 234–241.

- [31] C. Feng, D. Zhao, and M. Huang, "Image segmentation using cuda accelerated non-local means denoising and bias correction embedded fuzzy c-means (bcefcmm)," *Signal processing*, vol. 122, pp. 164–189, 2016.
- [32] M. Weigel, "Extended phase graphs: Dephasing, rf pulses, and echoes-pure and simple," *Journal of Magnetic Resonance Imaging*, vol. 41, no. 2, pp. 266–295, 2015.
- [33] B. A. Jung and M. Weigel, "Spin echo magnetic resonance imaging," *Journal of Magnetic Resonance Imaging*, vol. 37, no. 4, pp. 805–817, 2013.
- [34] M. Markl and J. Leupold, "Gradient echo imaging," *Journal of Magnetic Resonance Imaging*, vol. 35, no. 6, pp. 1274–1289, 2012.
- [35] J. Hennig, A. Nauerth, and H. Friedburg, "Rare imaging: A fast imaging method for clinical mr," *Magnetic resonance in medicine*, vol. 3, no. 6, pp. 823–833, 1986.
- [36] A. Haase, J. Frahm, D. Matthaei, W. Hanicke, and K.-D. Merboldt, "Flash imaging. rapid nmr imaging using low flip-angle pulses," *Journal of Magnetic Resonance (1969)*, vol. 67, no. 2, pp. 258–266, 1986.

## A. Extended Phase Graph

Extended Phase Graphs (EPG) [32] are a powerful way to simulate the magnetization response of various MR sequences. A key aspect of EPG involves the manipulation of magnetization vectors through changes in a Fourier basis. Initially, the longitudinal  $M_z(z)$  and transverse  $M_{xy}$  and  $M_{xy}^*$  magnetization components are projected on a Fourier basis to obtain the Fourier components  $F_n^+$ ,  $F_n^-$ , and  $Z_n$ :

$$\begin{aligned} F_n^+ &= \int_0^1 M_{xy}(z) e^{-2\pi i n z} dz \\ F_n^- &= \int_0^1 M_{xy}^*(z) e^{-2\pi i n z} dz \\ Z_n &= \int_0^1 M_z(z) e^{-2\pi i n z} dz \\ Q &= \begin{bmatrix} F_0^+ & F_1^+ & F_2^+ & \dots & F_N^+ \\ F_0^{+*} & F_1^- & F_2^- & \dots & F_N^- \\ Z_0 & Z_1 & Z_2 & \dots & Z_N \end{bmatrix} \end{aligned}$$

The  $F$  states are then combined across  $n$  harmonics to form the EPG basis  $Q$  (shown above), which describes the magnetization state in one isochromat. Using the above EPG formalization, we can model the spin precession as a matrix operation:

$$\begin{bmatrix} F_n^+ \\ F_n^- \\ Z_n \end{bmatrix}' = \begin{bmatrix} e^{i\theta} & 0 & 0 \\ 0 & e^{-i\theta} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} F_n^+ \\ F_n^- \\ Z_n \end{bmatrix} \quad (2)$$

where  $\theta$  is the accumulated phase. Similarly, an RF pulse that flips the magnetization longitudinally by  $\alpha$  and transversally by  $\phi$  as:

$$\begin{bmatrix} F_n^+ \\ F_n^- \\ Z_n \end{bmatrix}' = \begin{bmatrix} \cos^2(\alpha/2) & e^{2i\phi} \sin^2(\alpha/2) & -ie^{i\phi} \sin(\alpha) \\ e^{-2i\phi} \sin^2(\alpha/2) & \cos^2(\alpha/2) & ie^{-i\phi} \sin(\alpha) \\ -\frac{i}{2e^{-i\phi} \sin(\alpha)} & \frac{i}{2e^{i\phi} \sin(\alpha)} & \cos(\alpha) \end{bmatrix} \begin{bmatrix} F_n^+ \\ F_n^- \\ Z_n \end{bmatrix}$$

Also, the evolution of relaxation for Transverse states  $F_n'(t)$  and Longitudinal states  $Z_n(t)$  and  $Z_0(t)$  follows the equations:

$$\begin{aligned} F_n'(t) &= F_n e^{-t/T_2} \\ Z_n'(t) &= Z_n e^{-t/T_1} \\ Z_0(t) &= M_0(1 - e^{-t/T_1}) + Z_0 e^{-t/T_1} \end{aligned}$$

Equivalent linear operations can also be used to model gradients. Therefore, an MRI clinical sequence of events (RF pulses, gradients, precession, relaxation, etc.) is simulated as a multiplication of  $Q$  (representing the voxel's state) through several matrices corresponding to precession, RF rotations, relaxation operators, and gradients. Until recently, EPG could only describe echo amplitudes, but it was extended by Phase Distribution Graphs (PDG) [28] to describe full echo shapes, including gradient echoes that provide spatial encoding.

## B. MRI Differentiable Simulator

MR-zero [27] is an MRI simulation tool that uses PDG to simulate the magnetization changes based on the MRI sequence events and input  $T_1$ ,  $T_2$ , and Proton Density (PD) maps. MR-zero can simulate arbitrary MRI sequences, including clinical MRI sequences, such as Spin-Echo [33], Gradient Recalled Echo [34], Rapid Acquisition with Relaxation Enhancement [35], and Fast Low Angle Shot [36], etc. The simulation process in MR-zero comprises two passes: The first pass swiftly simulates numerous states to provide an initial signal estimate. This step identifies crucial states contributing to the signal and discards irrelevant ones. Secondly, the main pass utilizes the precise signal equation, considering all voxels but minimizing the number of simulated states based on information gleaned from the pre-pass. Doing so generates a meaningful signal




Figure 3: The first row shows the raw tissue probability maps for subject 42 in the BrainWeb dataset. The second row shows the checkerboard pattern enclosed in a red-marked region while optimizing 19 linear coefficients instead of using probability maps directly.

suitable for reconstructing  $T_1/T_2$ -weighted scans as output. As opposed to previous simulators such as [14]–[20], MR-zero is implemented in PyTorch and is differentiable by design. This allows us to incorporate MR-zero as part of the forward inference pipeline through which backpropagation can be performed, allowing one to efficiently optimize the input  $T_1/T_2$ /PD maps or the MRI sequence given a  $T_1/T_2$ -weighted image.

## C. Results

### C.1. Direct pixel optimization

We optimized probability values directly and got the PSNRs and SSIMs for CSF, GM, and WM, respectively (Table 3). The direct pixel optimization outperforms all the other methods. One of those methods includes optimizing either 1 coefficient (*scalar optimization*) or 19 coefficients (*linear coefficients*) per pixel. The PSNR and SSIM values for coefficient optimization failed to surpass those of the baseline configuration and are off by a margin of 5.09 and 0.02 points, respectively, on average for all the maps. Additionally, examination of the results depicted in Fig 3 revealed a checkerboard pattern present in the optimized probability maps generated using this technique. Also, we tried optimizing a scalar coefficient per probability map, with all the contrasts of six sequences, which performed poorly compared to all the methods by being 14.39 and 0.47 points lower on average, respectively.

### C.2. All 3 maps vs. single map

In Table 4, we report DICE, PSNR, and SSIM metrics for single map optimization (e.g., only optimize CSF while GM and WM are assumed to be known). For all tissues, single map optimization has outperformed the optimization of all three maps concurrently. Single map optimization is significantly easier because the two other maps are assumed to be known and fixed,



Table 3: Dice score, PSNR, and SSIM comparisons of different types of optimization procedures. † represents the baseline configuration.

Method	CSF			GM			WM		
	Dice↑	PSNR ↑	SSIM ↑	Dice↑	PSNR ↑	SSIM ↑	Dice↑	PSNR ↑	SSIM ↑
Direct pixel estimation †	<b>0.55±0.04</b>	<b>34.45±4.87</b>	<b>0.98±0.02</b>	<b>0.76±0.02</b>	<b>26.8±2.54</b>	<b>0.96±0.02</b>	<b>0.79±0.01</b>	<b>25.99±2.69</b>	<b>0.92±0.03</b>
Linear coefficients	0.54±0.05	28.26±4.26	0.96±0.03	0.75±0.03	22.55±3.06	0.94±0.04	<b>0.79±0.01</b>	21.19±2.94	0.9±0.05
Scalar coefficient	0.32±0.09	16.46±1.1	0.42±0.09	0.66±0.05	13.69±1.01	0.51±0.1	0.63±0.04	13.92±1.09	0.53±0.07

Table 4: Dice score, PSNR, and SSIM comparisons on optimizing either all maps concurrently or only a single probability map given the other two fixed. † represents the baseline configuration.

Maps	CSF			GM			WM		
	Dice↑	PSNR ↑	SSIM ↑	Dice↑	PSNR ↑	SSIM ↑	Dice↑	PSNR ↑	SSIM ↑
All 3 †	0.55±0.04	34.45±4.87	0.98±0.02	0.76±0.02	26.8±2.54	0.96±0.02	0.79±0.01	25.99±2.69	0.92±0.03
CSF	<b>0.56±0.03</b>	<b>40.34±10.73</b>	<b>0.99±0.01</b>	-	-	-	-	-	-
GM	-	-	-	<b>0.78±0.02</b>	<b>47.23±10.43</b>	<b>0.99±0.01</b>	-	-	-
WM	-	-	-	-	-	-	<b>0.8±0.01</b>	<b>43.33±10.65</b>	<b>0.99±0.01</b>

Table 5: Dice score, PSNR, and SSIM comparisons on tissue probability maps while computing loss in Image-space and K-space domain. † represents the baseline configuration.

Maps	CSF			GM			WM		
	Dice↑	PSNR ↑	SSIM ↑	Dice↑	PSNR ↑	SSIM ↑	Dice↑	PSNR ↑	SSIM ↑
Image-space †	<b>0.55±0.04</b>	<b>34.45±4.87</b>	<b>0.98±0.02</b>	<b>0.76±0.02</b>	<b>26.8±2.54</b>	<b>0.96±0.02</b>	<b>0.79±0.01</b>	<b>25.99±2.69</b>	<b>0.92±0.03</b>
K-space	0.48±0.11	23.16±8.15	0.71±0.27	0.72±0.07	18.21±5.02	0.75±0.21	0.74±0.06	18.11±4.4	0.72±0.17




Figure 4: From the left are CSF, GM, and WM. The first row shows the ground truth probability maps for subject 20 of BrainWeb. The second row shows the estimated map using our baseline method. The third and fourth row shows the inferred probability maps using U-Net and BCEFCM.

thus acting as a strong prior for estimating the third map.

### C.3. Image-space vs. K-space loss

In Table 5, we report comparisons between an image-space loss and a k-space loss, i.e. computing the loss directly at k-space values, without applying the inverse Fourier transform. The PSNR and SSIM values of estimated probability maps while computing the loss in K-space domain are  $(0.48 \pm 0.11, 23.16 \pm 8.15, 0.71 \pm 0.27)$ ,  $(0.72 \pm 0.07, 18.21 \pm 5.02, 0.75 \pm 0.21)$ , and  $(0.74 \pm 0.06, 18.11 \pm 4.4, 0.72 \pm 0.17)$  for CSF, GM, and WM respectively. These values, on average, are off by a huge margin of 9.25 and 0.23 points for PSNR and SSIM, respectively, from the baseline configuration metrics of the experiment done in Image-space. We found these results surprising since the pipeline with the k-space loss has one step less than the image-space pipeline. We plan to explore this more in our future work.

### D. More probability maps estimation results

This section of the Appendix contains the brain probability maps optimization for all other subjects of the BrainWeb dataset.




Figure 5: From the left are CSF, GM, and WM. The first and third rows show the raw tissue probability maps for subjects 04 and 05, respectively, of the BrainWeb dataset. The second and fourth rows show the optimized probability maps using the baseline configuration for subjects 04 and 05, respectively.




Figure 6: From the left are CSF, GM, and WM. The first and third rows show the raw tissue probability maps for subjects 06 and 18, respectively, of the BrainWeb dataset. The second and fourth rows show the optimized probability maps using the baseline configuration for subjects 06 and 18, respectively.




Figure 7: From the left are CSF, GM, and WM. The first and third rows show the raw tissue probability maps for subjects 20 and 38, respectively, of the BrainWeb dataset. The second and fourth rows show the optimized probability maps using the baseline configuration for subjects 20 and 38, respectively.




Figure 8: From the left are CSF, GM, and WM. The first and third rows show the raw tissue probability maps for subjects 41 and 43, respectively, of the BrainWeb dataset. The second and fourth rows show the optimized probability maps using the baseline configuration for subjects 41 and 43, respectively.




Figure 9: From the left are CSF, GM, and WM. The first and third rows show the raw tissue probability maps for subjects 44 and 45, respectively, of the BrainWeb dataset. The second and fourth rows show the optimized probability maps using the baseline configuration for subjects 44 and 45, respectively.






Figure 10: From the left are CSF, GM, and WM. The first and third rows show the raw tissue probability maps for subjects 46 and 47, respectively, of the BrainWeb dataset. The second and fourth rows show the optimized probability maps using the baseline configuration for subjects 46 and 47, respectively.




Figure 11: From the left are CSF, GM, and WM. The first and third rows show the raw tissue probability maps for subjects 48 and 49, respectively, of the BrainWeb dataset. The second and fourth rows show the optimized probability maps using the baseline configuration for subjects 48 and 49, respectively.




Figure 12: From the left are CSF, GM, and WM. The first and third rows show the raw tissue probability maps for subjects 50 and 51, respectively, of the BrainWeb dataset. The second and fourth rows show the optimized probability maps using the baseline configuration for subjects 50 and 51, respectively.




Figure 13: From the left are CSF, GM, and WM. The first and third rows show the raw tissue probability maps for subjects 52 and 53, respectively, of the BrainWeb dataset. The second and fourth rows show the optimized probability maps using the baseline configuration for subjects 52 and 53, respectively.




Figure 14: From the left are CSF, GM, and WM. The first row shows the raw tissue probability maps for subject 54 of the BrainWeb dataset. The second row shows the optimized probability maps using the baseline configuration for subject 54.