
MRI Parameter Mapping via Gaussian Mixture VAE: Breaking the Assumption of Independent Pixels

Moucheng Xu ^{1,2,6*}

Yukun Zhou ^{1,6}

Tobias Goodwin-Allcock ^{1,6}

Kimia Firoozabadi ⁵

Joseph Jacob ^{1,2,3}

Daniel C. Alexander ^{1,4}

Paddy J. Slator ^{1,4,7,8}

¹UCL Hawkes Institute, UK ²UCL Satsuma Lab, UK ³UCL Respiratory, UK

⁴UCL Department of Computer Science, UK ⁵UCL Medical School, UK

⁶UCL Department of Medical Physics and Biomedical Engineering, UK

⁷Cardiff University, School of Computer Science and Informatics, UK

⁸Cardiff University Brain Research Imaging Center, UK

moucheng.xu.18@alumni.ucl.ac.uk

slatorp@cardiff.ac.uk

Abstract

We introduce and demonstrate a new paradigm for quantitative parameter mapping in MRI. Parameter mapping techniques, such as diffusion MRI and quantitative MRI, have the potential to robustly and repeatably measure biologically-relevant tissue maps that strongly relate to underlying microstructure. Quantitative maps are calculated by fitting a model to multiple images, e.g. with least-squares or deep learning. However, the overwhelming majority of model fitting techniques assume that each voxel is independent, ignoring any co-dependencies in the data. This makes model fitting sensitive to voxelwise measurement noise, hampering reliability and repeatability. We propose a self-supervised deep variational approach that breaks the assumption of independent pixels, leveraging redundancies in the data to effectively perform data-driven regularisation of quantitative maps. We demonstrate that our approach outperforms current model fitting techniques in dMRI simulations and real data. Especially with a Gaussian mixture prior, our model enables sharper quantitative maps, revealing finer anatomical details that are not presented in the baselines. Our approach can hence support the clinical adoption of parameter mapping methods such as dMRI and qMRI. Our code is available at <https://github.com/moucheng2017/MRI-GMM-VAE>.

1 Introduction

Multiple MRI techniques can produce quantitative maps of biophysical, chemical and physiological tissue properties. Such *quantitative parameter mapping* techniques include diffusion MRI (dMRI) and quantitative MRI (qMRI). dMRI and qMRI use an essentially identical approach; by acquiring multiple images then fitting a model to the images, intrinsic values of the relevant tissue properties in each voxel can be estimated. In dMRI the images have different diffusion weightings and directions, and model fitting enables estimation of microstructural parameters, such as diffusivity and kurtosis.

*Work completed at UCL, Moucheng has now moved to Medtronic.

In qMRI acquisition parameters such as echo time or inversion time are varied, and model fits produce maps of chemical tissue properties, such as T1 and T2.

In the vast majority of applications, such models are fit to the data using non-linear least squares techniques. Machine learning model fitting is emerging as a attractive alternative technique. Supervised learning has been demonstrated for a range of models Palombo et al. [2020], but the distribution of parameters in the training dataset introduces biases Gyori et al. [2022], Epstein et al. [2022]. Self-supervised learning has the potential to address this, but has only been demonstrated in a few models thus far, most prominently the intravoxel incoherent motion (IVIM) model Barbieri et al. [2020]. Hybrid approaches are also emerging that merge the benefits of supervised and self-supervised learning Epstein et al. [2022].

However, whilst these approaches offer improvements, the current generation of parameter mapping techniques fail to capitalize on the extensive inherent redundancies in the data. Specifically, the overwhelming majority of techniques fit models to each voxel separately, effectively assuming that each voxel is independent. This leads to high sensitivity to voxelwise measurement noise, which negatively affects the quality of derived parameter maps. Bayesian hierarchical modelling has been proposed as an approach that breaks these assumptions, but requires slow Markov chain Monte Carlo inference Orton et al. [2014]. Convolutional neural networks (CNNs) Vasylechko et al. [2021] have been demonstrated, but only learn spatial redundancies. Transformers have also been demonstrated Zheng et al. [2023] but rely on gold standard training data. One technique assumes a set of underlying tissue components to regularise quantitative maps Slator et al. [2021] in a data-driven and unsupervised way, but at the expense of voxelwise parameter estimates.

In this paper, we demonstrate a deep learning approach that breaks the paradigm of independent voxels. Analogously to recent approaches Slator et al. [2021], we seek a lower dimensional representation of the data to parameterise data redundancies. We therefore adapt ideas from the clustering literature Manduchi et al. [2022] to derive a deep variational autoencoder for quantitative parameter mapping. We show that our approach yields improved parameter estimates and maps in simulated and real data.

2 Methods

Problem Setting We assume a dMRI dataset and model throughout the paper. We assume an observed discrete series of, T , dMRI images, $\mathcal{S} = S_{(1)}, S_{(2)}, \dots, S_{(T)}$, where $S_{(i)}$ is a diffusion weighted image (DWI) with height, H , width, W , and depth, D . The DWIs are defined $S_{(i)} = \{s_{(i)}^{(h,w,d)}; h \in [1, 2, \dots, H], w \in [1, 2, \dots, W], d \in [1, 2, \dots, D]\}$, where $s_{(i)}^{(h,w,d)}$ is the signal at the voxel (h, w, d) for the i th DWI. We aim to estimate the tissue properties, i.e. model parameters, $X = \{\mathbf{x}^{(h,w,d)}; h \in [1, 2, \dots, H], w \in [1, 2, \dots, W], d \in [1, 2, \dots, D]\}$ where the tissue properties for a voxel located at (h, w, d) corresponds to the vector $\mathbf{x}^{(h,w,d)} = (x_1^{(h,w,d)}, \dots, x_M^{(h,w,d)})$ for the M tissue properties. The model parameters are estimated with respect to the signal model. Traditionally, each voxel would be estimated independently of all the other voxels, for example, the MLE at (h_1, w_1, d_1) is calculated by finding the parameters that maximise the probability of seeing the observed data; i.e. what maximises $p((s_{(1)}^{(h_1,w_1,d_1)}, \dots, s_{(T)}^{(h_1,w_1,d_1)}) | \mathbf{x}^{(h_1,w_1,d_1)})$. This is performed independently to the MLE at (h_2, w_2, d_2) .

Latent Variable Model In this work, we propose to jointly model the distribution of the voxels together as $p_\theta(\mathcal{S})$ using a latent variable (\mathbf{z}) model: $p_\theta(\mathcal{S} | \mathbf{z}) p_\theta(\mathbf{z}) d\mathbf{z}$. Considering dMRI-specific format, with independence assumption for expression clarity (we note that despite this assumption our proposed model does not treat voxels as independent due to the shared latent space): $p_\theta(\mathcal{S}) \approx \prod_{t=1}^T \prod_{d=1}^D \prod_{w=1}^W \prod_{h=1}^H p_\theta(s_{(t)}^{(h,w,d)} | \mathbf{z}) p_\theta(\mathbf{z})$. Our latent variable model enjoys two benefits. The first benefit is, by conditioning the data from each DWI from each voxel $s_{(i)}^{(h,w,d)}$ on \mathbf{z} , we absorb all the arbitrary dependencies among voxels into \mathbf{z} , a compact representation in a latent space. In the latent space which has lower dimension than the data space, voxels are clustered with their close voxels, therefore inter-voxels information must be captured. The second benefit is, the complicated underlying distribution of $p_\theta(\mathcal{S})$ can be learnt via learning a much simpler distribution $p_\theta(\mathbf{z})$. However, the marginal distribution of $p_\theta(\mathbf{z})$ is still intractable: $p_\theta(\mathbf{z} | \mathcal{S}) = p_\theta(\mathcal{S}, \mathbf{z}) / p_\theta(\mathcal{S})$ because the data density is unknown. To address this computational issue, we deploy a variational

approach to approximate the posterior of $p_{\theta}(\mathbf{z})$. We explore different implementations to obtain the posterior of $p(\mathbf{z})$.

Univariate Gaussian prior We start with a simple univariate Gaussian $\mathcal{N}(0, 1)$ for prior $q(\mathbf{z})$ and our first implementation is called VAE-UniG. We first use the encoder to parameterize \mathbf{z} from input observed signals: $\mu, \text{Log}(\sigma^2) = \theta_{Encoder}(\mathcal{S})$. Then, $p(\mathbf{z}|\mathcal{S}, \theta_{Encoder}) \approx \mathcal{N}(\mu, \sigma)$. We then use a decoder to map randomly drawn samples $\mathbf{z} \sim p(\mathbf{z})$ to physically-relevant dMRI model parameters (X) that are inputs to a closed-form dMRI model that reconstructs the MRI signal. We denote the closed-form physics decoding process as $\phi(\cdot)$. We emphasise that $\phi(\cdot)$ can be any dMRI (or qMRI) model. The complete decoding process is: $p(\mathcal{S}) = p(\phi(X|\mathbf{z}, \theta_{Decoder}))$. The loss function is: $\text{Log}(p_{\theta}(\mathcal{S})) \geq \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})}[\text{Log}(p(\phi(\mathbf{X}|\mathbf{z}, \theta)))] - KL(p(\mathbf{z})||q(\mathbf{z}))$. The likelihood of $\text{Log}(p(\phi(\mathbf{x}|\mathbf{z}, \theta)))$ is measured as a mean squared error loss between estimated signals and raw input signals.

Gaussian Mixture Prior Our second implementation is VAE-GMM with enhanced expressivity by considering a prior as a mixture of univariate Gaussians. We add an extra latent variable \mathbf{y} for controlling the index of the Gaussian component. The prior of \mathbf{y} is chosen as a Categorical distribution. The probabilistic model is: $p_{\theta}(\mathcal{S}) = \int_{\mathbf{z}} \int_{\mathbf{y}} p_{\theta}(\mathcal{S}|\mathbf{z})p(\mathbf{z}|\mathbf{y})p(\mathbf{y})d\mathbf{y}d\mathbf{z}$. In implementation, we build our Gaussian mixture VAE (VAE-GMM) following a hierarchical order. We first need to parametrize the mixing coefficients of each Gaussian using Gumbel Softmax Maddison et al. [2016]: $c = \text{Gumbel}(\theta_{Encoder_{1st}}(\mathcal{S}))$. Where c is a normalised vector indicating the weight for each Gaussian and the sum of the c is 1. We then concatenate c with input to parametrize the parameters of Gaussians: $\mu_k, \text{Log}(\sigma_k^2) = \theta_{Encoder_{2nd}}(\mathcal{S}, c)$. Also, $p(\mathbf{z}|\mathcal{S}, \theta_{Encoder}) \approx \mathcal{N}(\mu_k, \sigma_k)$. Where k means that the mean and the variance are for the K -th Gaussian. We apply the same decoding process as in the last section. And the loss function now becomes: $\text{Log}(p_{\theta}(\mathcal{S})) \geq \mathbb{E}_{\mathbf{z} \sim p(\mathbf{y}|\mathbf{z})p(\mathbf{z}|\mathbf{x})}[\text{Log}(p(\phi(\mathbf{X}|\mathbf{z}, \theta)))] - KL(p(\mathbf{z}, \mathbf{y})||q(\mathbf{z}, \mathbf{y}))$.

MRI physics models We test our approach on two dMRI models, the mean signal diffusion kurtosis imaging (MS-DKI)Henriques et al. [2019] model and ball-stick model Behrens et al. [2003]. The normalised signal for MS-DKI² is given by: $\phi(b) = \exp(-bD + b^2D^2K/6)$, where b is the b-value, D is the diffusivity and K the kurtosis. For ball-stick³ the normalised signal is: $\phi(b, \mathbf{g}) = f \exp(-bD_{||}(\mathbf{g}, \mathbf{n})) + (1 - f) \exp(-bD_{iso})$, where b is the b-value, \mathbf{g} the gradient direction, f is the stick volume fraction, $D_{||}$ is the parallel diffusivity of the stick, and D_{iso} is the ball isotropic diffusivity.

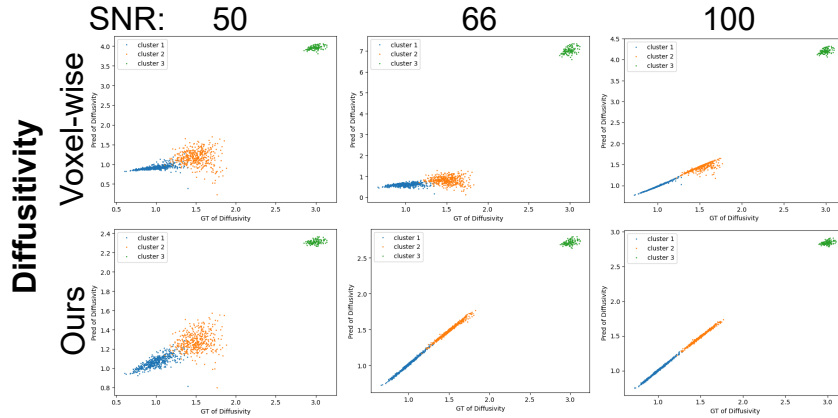


Figure 1: Diffusivity results on simulated model using MSDKI, comparisons between self-supervised baseline and our VAE-UniG. X axis: ground truth of simulated diffusivity. Y axis: prediction of diffusivity. SNR: signal-noise-ratio. Blue: cluster 1. Orange: cluster 2. Green: cluster 3. Ours vastly outperforms the baseline in recovery of the ground truth. See the kurtosis results in the appendix Fig.7.

²For MSDKI we implemented LSQ fitting with the dipy python package Garyfallidis et al. [2014].

³For ball-stick we used the dmipy package Fick et al. [2019].

3 Results

Simulated Data We use least squares fitting and voxel-wise self-supervised fitting (a 3-layer fully connected network Lim et al. [2022]) as baselines. We first test our method on a simulated MSDKI dataset with 3 clusters. We chose each cluster’s diffusivity and kurtosis to mimic white matter, grey matter, and CSF; the mean D, K values for each cluster were $\{1, 1.5\}$, $\{1.5, 1\}$, and $\{3, 0\}$ respectively, with diffusivity in units of $\mu\text{m}^2/\text{ms}$. We simulated 10,000 voxels, with relative weightings of each cluster $\{0.5, 0.4, 1\}$. The specific ground truth parameter value was simulated from a Gaussian with the relevant mean D and K for that cluster, and variance 0.1 for white matter and grey matter clusters, and 0.01 for the CSF cluster. As shown in the Fig. 1, across different settings of signal-noise ratio (SNR), our model’s predictions (row 2) are much better diagonally aligned, meaning that our latent variable model consistently outperform the voxel-wise self-supervised baseline (row 1).

Real Data We as well test our model on real data with ball-stick fit, with publicly-available dMRI data from HCP WU-Minn Consortium Van Essen et al. [2012]. We used preprocessed Glasser et al. [2013] data from a single subject from the 1200 Subjects Data Release (Release Date: Mar 01, 2017, available online at <http://humanconnectome.org>). As shown in the 1st row in Fig. 2, when predicting the diffusivity parameters, both of our models can drastically reduce the background noises in comparison with the baselines. Interestingly, our model VAE-GMM (4th column) can even discover new anatomical structures, as seen in highlighted areas in row3 in Fig. 2, showing potential clinical application promises. Evidently, our model VAE-GMM also successfully captured at least two mixing components of the latent variable, as shown in Fig. 3.

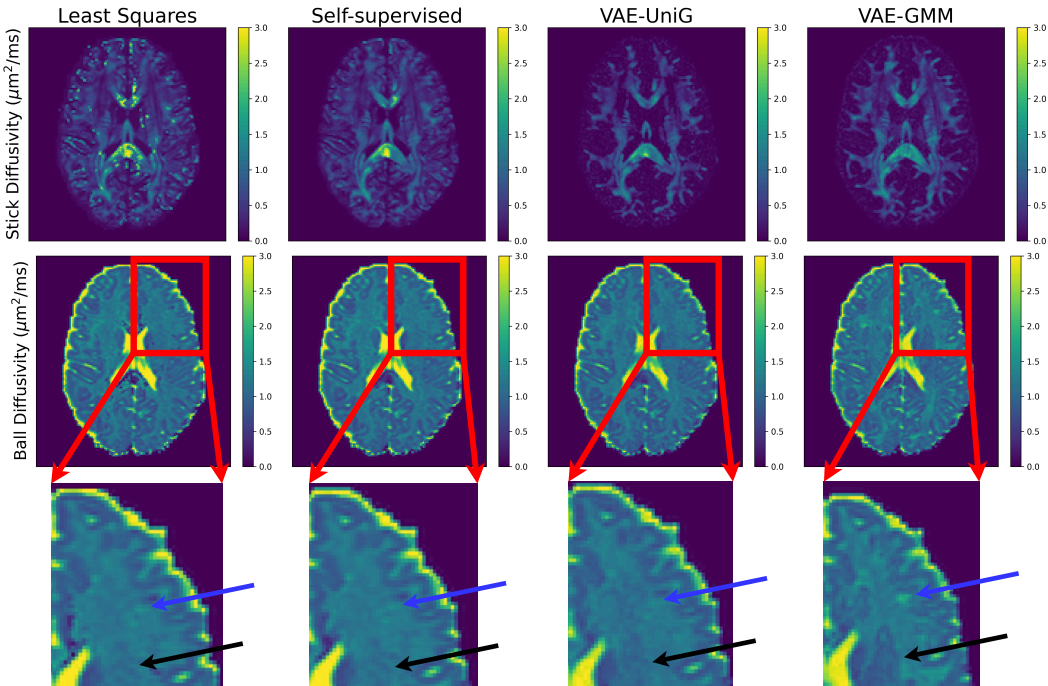


Figure 2: LSQ, self-supervised, VAE-UniG, VAE-GMM ball-stick fits to HCP dMRI subject. Both our VAE models drastically reduce noises (1st Row) for sharper white matter visualisation. More importantly, our VAE-GMM reveals finer anatomical structures with clear details which were not seen in the baselines (see the arrows locations in 3rd Row, 4th Col).




Figure 3: Visualisation of the learnt posterior distributions of the latent variable z training on HCP data. 1st row: VAE. 2nd row: GMM VAE.

4 Conclusion

We introduce a deep VAE model fitting method that exploits data redundancies to maximise information extraction in parameter mapping techniques like dMRI and qMRI. Our approach outperforms baseline methods in both simulated and real data using the ball-stick model, improving ground truth estimations and revealing new anatomical details. It can enhance existing acquisition sequences for better tissue maps or shorten scan times without loss of quality. However, limitations include variability in parameter maps and trade-offs between parameter accuracy. Future work will focus on generalising the method to handle arbitrary acquisition schemes.

5 Acknowledgements

MCX was supported by GSK (BIDS3000034123) and UCL Engineering Dean’s Prize. DCA is supported by UK EPSRC grants M020533, R006032, R014019, V034537, Wellcome Trust UNS113739, Wellcome Trust 221915/Z/20/Z. JJ is supported by Wellcome Trust Clinical Research Career Development Fellowship 227835/Z/23/Z. DCA, and JJ are supported by the NIHR UCLH Biomedical Research Centre, UK. This research was funded in whole or in part by the Wellcome Trust [209553/Z/17/Z].

6 Impact Statement

Medical physics plays a pivotal role in advancing physical sciences for the betterment of humanity. Among the many branches of this field, Magnetic Resonance Imaging (MRI) stands out as a critical area of study due to its wide-ranging applications in diagnostics, treatment planning, and medical research. MRI has revolutionized healthcare, enabling non-invasive imaging with detailed insights into the human body, making it indispensable for modern medicine. Our proposed method aims to further enhance MRI parameter mapping, improving both accuracy and efficiency. This has the potential to significantly impact the medical community by providing more reliable tools for diagnosis and research. Moreover, the method is versatile and can be adapted for use in a wide variety of parameter mapping applications beyond MRI, broadening its utility across multiple domains in medical physics and other scientific fields. By offering a solution that is both robust and adaptable, our

work contributes to the broader goal of leveraging machine learning for physical sciences, providing meaningful societal benefits through improved medical imaging techniques and related applications.

References

- Sebastiano Barbieri, Oliver J. Gurney-Champion, Remy Klaassen, and Harriet C. Thoeny. Deep learning how to fit an intravoxel incoherent motion model to diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 83(1):312–321, January 2020. ISSN 0740-3194, 1522-2594. doi: 10.1002/mrm.27910.
- T.E.J. Behrens, M.W. Woolrich, Mi Jenkinson, H. Johansen-Berg, R.G. Nunes, S. Clare, P.M. Matthews, J.M. Brady, and S.M. Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, 50(5):1077–1088, November 2003. ISSN 0740-3194. doi: 10.1002/mrm.10609.
- Sean C. Epstein, Timothy J. P. Bray, Margaret Hall-Craggs, and Hui Zhang. Choice of training label matters: How to best use deep learning for quantitative MRI parameter estimation. *ArXiv Preprint*, May 2022. doi: 10.48550/arxiv.2205.05587.
- Rutger H.J. Fick, Demian Wassermann, and Rachid Deriche. The Dmipy Toolbox: Diffusion MRI Multi-Compartment Modeling and Microstructure Recovery Made Easy. *Frontiers in Neuroinformatics*, 13(October):1–26, 2019. doi: 10.3389/fninf.2019.00064.
- Eleftherios Garyfallidis, Matthew Brett, Bagrat Amirbekian, Ariel Rokem, Stefan van der Walt, Maxime Descoteaux, Ian Nimmo-Smith, and Dipy Contributors. Dipy, a library for the analysis of diffusion MRI data. *Frontiers in Neuroinformatics*, 8, February 2014. ISSN 1662-5196. doi: 10.3389/fninf.2014.00008.
- Matthew F. Glasser, Stamatios N. Sotiropoulos, J. Anthony Wilson, Timothy S. Coalson, Bruce Fischl, Jesper L. Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R. Polimeni, David C. Van Essen, and Mark Jenkinson. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124, October 2013. ISSN 10538119. doi: 10.1016/j.neuroimage.2013.04.127.
- Noemi G. Gyori, Marco Palombo, Christopher A. Clark, Hui Zhang, and Daniel C. Alexander. Training data distribution significantly impacts the estimation of tissue microstructure with machine learning. *Magnetic Resonance in Medicine*, 87(2):932–947, February 2022. ISSN 0740-3194, 1522-2594. doi: 10.1002/mrm.29014.
- Rafael Neto Henriques, Sune N. Jespersen, and Noam Shemesh. Microscopic anisotropy misestimation in spherical-mean single diffusion encoding MRI. *Magnetic Resonance in Medicine*, 81(5): 3245–3261, 2019. doi: 10.1002/mrm.27606.
- Jason P. Lim, Stefano B. Blumberg, Neil Narayan, Sean C. Epstein, Daniel C. Alexander, Marco Palombo, and Paddy J. Sator. Fitting a Directional Microstructure Model to Diffusion-Relaxation MRI Data with Self-Supervised Machine Learning, October 2022.
- Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *ArXiv*, abs/1611.00712, 2016. URL <https://api.semanticscholar.org/CorpusID:14307651>.
- Laura Manduchi, Ričards Marcinkevičs, Michela C. Massi, Thomas Weikert, Alexander Sauter, Verena Gotta, Timothy Müller, Flavio Vasella, Marian C. Neidert, Marc Pfister, Bram Stieltjes, and Julia E. Vogt. A Deep Variational Approach to Clustering Survival Data. In *ICLR 2022*. arXiv, 2022.
- Matthew R. Orton, David J. Collins, Dow-Mu Mu Koh, and Martin O. Leach. Improved intravoxel incoherent motion analysis of diffusion weighted imaging by data driven Bayesian modeling. *Magnetic Resonance in Medicine*, 71(1):411–420, 2014. doi: 10.1002/mrm.24649.
- Marco Palombo, Andrada Ianus, Michele Guerreri, Daniel Nunes, Daniel C. Alexander, Noam Shemesh, and Hui Zhang. SANDI: A compartment-based model for non-invasive apparent soma and neurite imaging by diffusion MRI. *NeuroImage*, 215:116835–116835, July 2020. doi: 10.1016/j.neuroimage.2020.116835.

- Paddy J. Slator, Jana Hutter, Razvan V. Marinescu, Marco Palombo, Laurence H. Jackson, Alison Ho, Lucy C. Chappell, Mary Rutherford, Joseph V. Hajnal, and Daniel C. Alexander. Data-Driven multi-Contrast spectral microstructure imaging with InSpect: INtegrated SPECTral component estimation and mapping. *Medical Image Analysis*, 71:102045–102045, 2021. doi: 10.1016/j.media.2021.102045.
- D.C. Van Essen, K. Ugurbil, E. Auerbach, D. Barch, T.E.J. Behrens, R. Bucholz, A. Chang, L. Chen, M. Corbetta, S.W. Curtiss, S. Della Penna, D. Feinberg, M.F. Glasser, N. Harel, A.C. Heath, L. Larson-Prior, D. Marcus, G. Michalareas, S. Moeller, R. Oostenveld, S.E. Petersen, F. Prior, B.L. Schlaggar, S.M. Smith, A.Z. Snyder, J. Xu, and E. Yacoub. The Human Connectome Project: A data acquisition perspective. *NeuroImage*, 62(4):2222–2231, October 2012. ISSN 10538119. doi: 10.1016/j.neuroimage.2012.02.018.
- Serge Didenko Vasylechko, Simon K. Warfield, Onur Afacan, and Sila Kurugol. Self-supervised IVIM DWI parameter estimation with a physics based forward model. *Magnetic Resonance in Medicine*, 00:1–11, 2021. doi: 10.1002/mrm.28989.
- Tianshu Zheng, Guohui Yan, Haotian Li, Weihao Zheng, Wen Shi, Yi Zhang, Chuyang Ye, and Dan Wu. A microstructure estimation transformer inspired by sparse representation for diffusion mri. *Medical Image Analysis*, 86:102788, 5 2023. ISSN 1361-8415. doi: 10.1016/J.MEDIA.2023.102788.

A The architectures of the neural networks




Figure 4: Architectures of our model implementations. Row 1: with univariate prior. Row 2: with Gaussian mixture model prior. The encoder is three fully connected layers. The decoder is one fully connected layer.

B Results on MSDKI MRI model




Figure 5: Comparisons on MS-DKI fits on HCP dMRI subject. Our approach has less obvious improvements when the MRI model is relatively simple, but doesn't hallucinate spurious anatomical features.

C Results on the sensitivity of hyper parameters




Figure 6: Hyper-parameter sensitivities of our model with diffusivity results on simulated model using MSDKI, comparisons between self-supervised baseline and ours with Gaussian prior. X axis: ground truth of simulated diffusivity. Y axis: prediction of diffusivity. We observe that both latent dimension and kl loss strength have optimal values, but those values might be data dependent.

D Results on kurtosis estimation of MSDKI with VAE and simulated data




Figure 7: Kurtosis results on simulated model using MSDKI, comparisons between self-supervised baseline and our VAE-UniG. X axis: ground truth of simulated diffusivity. Y axis: prediction of diffusivity. SNR: signal-noise-ratio. Blue: cluster 1. Orange: cluster 2. Green: cluster 3. Ours vastly outperforms the baseline in recovery of the ground truth.