



Published in final edited form as:

Med Image Anal. 2023 April ; 85: 102742. doi:10.1016/j.media.2023.102742.

Deep learning prediction of diffusion MRI data with microstructure-sensitive loss functions

Geng Chen,

Yoonmi Hong,

Khoi Minh Huynh,

Pew-Thian Yap*

Department of Radiology, University of North Carolina, Chapel Hill, NC, USA

Biomedical Research Imaging Center (BRIC), University of North Carolina, Chapel Hill, NC, USA

Abstract

Deep learning prediction of diffusion MRI (DMRI) data relies on the utilization of effective loss functions. Existing losses typically measure the signal-wise differences between the predicted and target DMRI data without considering the quality of derived diffusion scalars that are eventually utilized for quantification of tissue microstructure. Here, we propose two novel loss functions, called microstructural loss and spherical variance loss, to explicitly consider the quality of both the predicted DMRI data and derived diffusion scalars. We apply these loss functions to the prediction of multi-shell data and enhancement of angular resolution. Evaluation based on infant and adult DMRI data indicates that both microstructural loss and spherical variance loss improve the quality of derived diffusion scalars.

Keywords

Diffusion MRI; Deep learning; Loss functions; Data prediction

1. Introduction

Prediction of diffusion MRI (DMRI) data is central to problems such as reconstruction (Hong et al., 2019a; Mani et al., 2020), phase correction (Aggarwal et al., 2019), denoising (Fadnavis et al., 2020), super-resolution (Tanno et al., 2017), harmonization (Mirzaalian et al., 2016, 2018; Tax et al., 2019; Ning et al., 2020; Moyer et al., 2020), longitudinal data imputation (Hong et al., 2019b), and microstructural analysis (Ye et al., 2019; Raffelt et al., 2012). Predicting unacquired data from acquired data is often regarded as a regression problem, which can be tackled using methods such as sparse representation (St-Jean et al.,

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

*Correspondence to: Bioinformatics Building, 130 Mason Farm Road, Chapel Hill, NC, 27599, USA., ptyap@med.unc.edu (P.-T. Yap).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

2016) and kernel regression (Chen et al., 2019a). However, the problem can be ill-posed and challenging to solve when the amount of data that needs to be predicted increases. Learning-based algorithms, particularly those based on deep learning (Schultz, 2012; Ye, 2017; Koppers et al., 2017; Ye et al., 2019; Gibbons et al., 2019; Zhang et al., 2020), can be more effective for ill-posed regression problems since they can leverage information learned from training data (Hong et al., 2019b).

A key factor in determining success in prediction is the loss function, which defines the optimization target (LeCun et al., 2015). Common loss functions, including ℓ_1 , ℓ_2 , and Huber losses (Huber, 1992; Kim et al., 2018), measure point-wise differences between the prediction and the target and can yield good peak signal-to-noise ratio (PSNR) performance (Johnson et al., 2016; Dosovitskiy and Brox, 2016; Ledig et al., 2017). DMRI analysis often focuses on derived diffusion scalars rather than raw DMRI data; therefore, prediction accuracy of diffusion scalars needs to be taken into consideration when designing loss functions.

In this work, we propose two novel loss functions, called the microstructural loss (ML) and the spherical variance loss (SVL), to explicitly consider the quality of derived diffusion scalars during DMRI data prediction. The ML includes an ℓ_1 loss for DMRI data and an ℓ_1 loss for a diffusion scalar of interest, estimated by a microstructure estimator sub-network. The SVL is based on the spherical variance of the diffusion signal profile — a signal statistic of each shell that is related to tissue diffusion properties (Minh Huynh et al., 2019; Huynh et al., 2019). The SVL is a composite of an ℓ_1 loss for the DMRI data and an ℓ_1 loss for the spherical variance.

To demonstrate the effectiveness of the two loss functions, we apply them to two DMRI data prediction tasks: Multi-shell (MS) data prediction and angular resolution enhancement. The first task aims to predict MS data by estimating additional shells from data with fewer shells. MS data allows the use of advanced diffusion models, e.g., neurite orientation dispersion and density imaging (NODDI) (Zhang et al., 2012) and diffusion kurtosis imaging (DKI) (Jensen et al., 2005), for probing tissue microstructure. However, collecting MS data requires longer acquisition times. Few methods were proposed for predicting MS data. A popular method described in Koppers et al. (2016) employs a multi-layer perceptron (MLP) to learn data prediction with an ℓ_1 loss function.

Angular resolution enhancement can be achieved by predicting high angular resolution (HAR) DMRI data from its low-resolution counterpart. HAR DMRI data allows the use of sophisticated diffusion models for better estimation of local fiber geometric structures. Angular resolution enhancement allows increasing the angular resolution post-acquisition without prolonging the acquisition time (Yin et al., 2018; Chen et al., 2019b; Lyon et al., 2022).

The effectiveness of the ML and SVL will be demonstrated with data from the Baby Connectome Project (BCP) (Howell et al., 2019) and Human Connectome Project (HCP) (Van Essen et al., 2013). The experimental results indicate that these losses improve data prediction accuracy, both qualitatively and quantitatively. In what follows, we will first give

a detailed description of our losses in Section 2 and report experiment results in Section 3. We will then provide additional discussions in Section 4 and finally conclude this work in Section 5.

2. Methods

In this section, we will first briefly introduce our data prediction network. We will then explain the details of our losses, i.e., the network-based ML and the statistic-based SVL. Finally, we will summarize and compare these two loss functions.

2.1. Data predictor

The data predictor (left in Fig. 1) estimates a non-linear function $f_{\hat{p}}(\cdot)$ for predicting \hat{T} , an estimation of the ground truth target data T , from input data S . To be consistent, we employ the same data predictor for the ML, the SVL, and all baseline losses.

In practice, the data predictor can be designed using a variety of network architectures, e.g., MLPs, convolutional neural networks (LeCun et al., 1998), and graph convolutional neural networks (Bronstein et al., 2017). In this work, we implement the data predictor utilizing an MLP for three main reasons: (i) MLP is widely employed in DMRI data prediction (Golkov et al., 2016; Koppers et al., 2016); (ii) MLP is straightforward to implement; and (iii) Most importantly, the main focus of our work is to investigate the effectiveness of novel loss functions; therefore we opt to adopt a simple and well-studied network for the data predictor. A comprehensive investigation of different data predictor architectures is beyond the scope of this paper.

MLP is a feedforward artificial neural network consisting of at least three layers. The first/last layer is an input/output layer. Therefore, the number of neural nodes in the input/output layer is determined by the dimension of the input/output data. The intermediate layers are hidden layers that encode features of different levels. Following (Golkov et al., 2016), we utilize three hidden layers, each with 150 neural nodes, matching approximately the number of our input measurements. The feature vector l_i of layer i is computed using

$$l_i = a(W_i l_{i-1} + b_i), \quad (1)$$

where $a(\cdot)$ is a non-linear activation function, ReLU. W_i and b_i are learned weights and biases. The linear representations of the features of a layer are non-linearly transformed with a ReLU activation function to give the features of the next layer. MLP is trained with a backpropagation algorithm (Rumelhart et al., 1988), which propagates errors, computed with respect to a loss function, through the network in a back-to-front manner so that trainable network parameters can be updated to eventually minimize the loss.

We predict data missing in the source and then utilize a concatenation branch to combine the predicted data and the source data for a complete output (Fig. 1). This strategy is useful for angular resolution enhancement, where the low angular resolution data is a subset of its high angular resolution counterpart. However, this strategy needs to be modified for MS data prediction (Section 3.2), since the source and target data are acquired using two noncollinear

sets of different gradients. In this case, we remove the concatenation branch and directly feed the outcome of the data predictor to the microstructure estimator.

2.2. Microstructure Loss (ML)

Microstructure loss is based on a microstructure estimator network (right in Fig. 1) that estimates a microstructure scalar of interest based on the output of the data predictor. The microstructure loss consists of an ℓ_1 term for DMRI data and an ℓ_1 term for the microstructure scalar, i.e.,

$$L_{\text{ML}} = \frac{1}{N(T)} \|\hat{T} - T\|_1 + \lambda_{\text{ML}} \|f_{\text{E}}(\hat{T}) - f_{\text{E}}(T)\|_1, \quad (2)$$

where \hat{T} is the prediction of target data T , $N(T)$ is the number of elements in T , λ_{ML} is a parameter balancing the contributions of the two ℓ_1 terms, and $f_{\text{E}}(\cdot)$ is a non-linear function realized via the microstructure estimator.

We first train the microstructure estimator using T and the microstructure scalar estimated from T . We then fix the microstructure estimator and train the data predictor using paired source and target data.

Similar to the data predictor, we design the microstructure estimator f_{E} using MLP with three hidden layers, each with 150 neural nodes. The number of neural nodes in the input layer is determined by the dimension of T .

2.3. Spherical Variance Loss (SVL)

Unlike the network-based ML, the SVL is a mathematical loss based on signal statistics (Fig. 2). The spherical variance of the signal is related to tissue microstructural properties (Novikov et al., 2018; Xu et al., 2018; Minh Huynh et al., 2019). Unlike point-wise losses, such as ℓ_1 , ℓ_2 , and Huber losses (Johnson et al., 2016; Dosovitskiy and Brox, 2016; Ledig et al., 2017), the SVL promotes q -space sharpness and avoids bias in derived diffusion scalars, such as fractional anisotropy, due to q -space smoothing.

We denote a target diffusion signal as $T_b(g)$, where b and $g \in \mathbb{S}^2$ are the b -value and gradient direction, respectively. The spherical variance V_b is computed as the second-order central moment of signals acquired for the same b -shell, i.e.,

$$V_b = \frac{1}{4\pi} \int_{\mathbb{S}^2} (T_b(g) - M_b)^2 dg, \quad (3)$$

where M_b is the corresponding first order raw moment defined as

$$M_b = \frac{1}{4\pi} \int_{\mathbb{S}^2} T_b(g) dg. \quad (4)$$

We combine the ℓ_1 terms for DMRI data and spherical variances to form the SVL:

$$L_{\text{SVL}} = \frac{1}{N(T)} \|\hat{T} - T\|_1 + \frac{\lambda_{\text{SVL}}}{N(V)} \|\hat{V} - V\|_1, \quad (5)$$

where \hat{V} is a vector containing the spherical variances for different shells of the predicted DMRI data. Correspondingly, V is the vector for the target DMRI data. $N(V)$ is the number of elements in V , which is determined by the number of shells. λ_{SVL} is a parameter trading off the contributions of the two ℓ_1 terms.

The SVL constrains the prediction to match not only the DMRI data, but also higher-level signal statistics, such as the spherical variance. In our implementation of the SVL, we utilize the same data predictor as in the ML for fair comparison.

2.4. Method summary

Two loss functions are introduced for improving the quality of diffusion scalars derived from the predicted DMRI data. They are however different in three aspects:

1. **Motivation:** The ML is a network-based loss function with a microstructure estimator sub-network to infer microstructural properties based on the predicted DMRI data. In contrast, the SVL is formulated based on the fact that tissue properties are related to the spherical variance.
2. **Training:** The SVL is trained in a standard end-to-end manner. Training the ML involves first training the microstructure estimator and then training the data predictor with the microstructure estimator fixed. Therefore, the SVL is more training friendly than ML.
3. **Implementation Complexity:** The ML integrates the microstructure estimator in its framework. In contrast, the SVL directly computes the loss value using a closed-form equation, which is easier to implement than the ML.

3. Experiments

3.1. Datasets

3.1.1. Dataset – Multi-shell data prediction—We demonstrate the effectiveness of the ML and SVL in MS data prediction with the BCP data (Howell et al., 2019). All data were acquired using a Siemens 3T Magnetom Prisma MRI scanner with two imaging protocols: (i) **2-Shell:** $b = 1500, 3000 \text{ s/mm}^2$ with a total of 74 non-collinear gradient directions and (ii) **6-Shell:** $b = 500, 1000, 1500, 2000, 2500, 3000 \text{ s/mm}^2$ with a total of 144 non-collinear gradient directions. The two protocols shared some common imaging parameters, including 140×105 imaging matrix, $1.5 \times 1.5 \times 1.5 \text{ mm}^3$ resolution, $\text{TE} = 88 \text{ ms}$, $\text{TR} = 2365 \text{ ms}$, and 32-channel receiver coil. Note that the 2-shell and 6-shell data were acquired in the same imaging session. EDDY (Andersson and Sotiropoulos, 2016) and TOPUP (Andersson et al., 2003) were performed on the 2-shell and 6-shell data combined so that they are aligned.

Here we aim to predict the 6-shell data from its 2-shell counterpart. Our dataset consisted of 10 randomly selected individuals 1 year to 2 years of age. We extracted voxels to construct a large-scale training dataset with more than one hundred thousand samples.

3.1.2. Dataset – Angular resolution enhancement—We utilized the HCP data (Van Essen et al., 2013) for angular resolution enhancement. The data was acquired using a customized Siemens 3T Connectome Skyra MR scanner with the following imaging protocol: $b = 1000, 2000, 3000$ s/mm², 90 gradient directions per shell, 145×174 imaging matrix, $1.25 \times 1.25 \times 1.25$ mm³ resolution, TE = 89 ms, TR = 5500 ms, and 32-channel receiver coil (Van Essen et al., 2013).

We downsampled the data angularly with a factor of 5 by uniformly selecting 18 gradient directions for each shell to form the low angular resolution (LAR) data with a total of 54 gradient directions. The data predictor was trained to estimate the data of the remaining 216 gradient directions, which were combined with the LAR data for a complete output of 270 gradient directions.

Our dataset consisted of 16 randomly selected subjects. We utilized 6 subjects for training and the other 10 subjects for testing. We constructed the training dataset using a sufficiently large number of voxel samples.

3.2. Implementation details

For training, we randomly extracted 20,000 voxels from the brain region of every training subject to construct a training dataset containing in total $6 \times 20,000 = 120,000$ samples. The two tasks were trained in a supervised manner. The proposed methods were implemented using TensorFlow 1.2 (Abadi et al., 2016). We trained the networks using the ADAM optimizer (Kingma and Ba, 2014) with an initial learning rate of 0.0001 and an exponential decay rate of 0.95. The batch size was set to 1024. Over-fitting was avoided by early stopping. The networks were trained using an NVIDIA GeForce GTX 1080 Ti GPU with 11 GB RAM. For ML, we utilized generalized fractional anisotropy (GFA) (Tuch, 2004), a widely-adopted model-free diffusion scalar, as the prediction target of the microstructure estimator.

3.3. Parameter settings

We set the two tuning parameters in the ML and SVL as $\lambda_{ML} = 0.5$ and $\lambda_{SVL} = 1.0$. Parameter selection was based on grid search with [0.01, 0.1, 0.5, 1.0, 1.5, 2.0, 10.0]. More details are described in Section 3.6.2.

3.4. Baseline losses

We compared our method with various baseline losses described below.

1. **ℓ_1 loss:** Due to its robustness to outliers, the ℓ_1 loss is one of the most popular loss functions used in regression tasks. We define the ℓ_1 loss as

$$L_1 = \frac{1}{N(T)} \|\hat{T} - T\|_1. \quad (6)$$

2. **ℓ_2 loss:** The ℓ_2 loss is another popular loss function and is defined as

$$L_2 = \frac{1}{N(T)} \|\hat{T} - T\|_2^2. \quad (7)$$

Compared with the ℓ_1 loss, the ℓ_2 loss is less robust to outliers, but eases training (Zhao et al., 2016).

3. **Huber loss:** The Huber loss can be viewed as a combination of ℓ_1 loss and ℓ_2 loss. Following Huber (1992) and Kim et al. (2018), we define the Huber loss as

$$L_{\text{Huber}} = \begin{cases} \frac{1}{2} \|\hat{T} - T\|_2^2, & \text{for } \|\hat{T} - T\|_1 < \delta, \\ \delta \|\hat{T} - T\|_1 - \frac{1}{2} \delta^2, & \text{otherwise} \end{cases} \quad (8)$$

where δ is a parameter balancing the ℓ_1 loss and the ℓ_2 loss. We set $\delta = 0.5$ following Kim et al. (2018).

3.5. Evaluation methods

1. **Peak signal-to-noise ratio (PSNR):** Quantitative evaluation was performed using PSNR:

$$\text{PSNR} = 20 \log_{10} \frac{\text{MAX}}{\text{RMSE}}, \quad (9)$$

where MAX is the maximum value of the data and RMSE is the root mean square error between the prediction and the ground truth.

2. **Structural similarity index (SSIM):** Quantitative evaluation was also performed using SSIM (Wang et al., 2004):

$$\text{SSIM}(\hat{T}, T) = \frac{(2\mu(\hat{T})\mu(T) + C_1)(2\sigma(\hat{T}, T) + C_2)}{(\mu^2(\hat{T}) + \mu^2(T) + C_1)(\sigma^2(\hat{T}) + \sigma^2(T) + C_2)}, \quad (10)$$

where $\mu(\hat{T})$ and $\sigma(\hat{T})$ are the mean and variance of the prediction, $\sigma(\hat{T}, T)$ is the covariance of the prediction and the ground truth, C_1 and C_2 are two positive constants that ensure that the denominator is non-zero.

3. **Microstructure scalars:** We considered various microstructure models for evaluation. The corresponding microstructure scalars included GFA given by q -ball imaging (Tuch, 2004), mean kurtosis (MK) given by DKI (Jensen et

al., 2005), as well as orientation dispersion (OD), neurite density (ND), and cerebrospinal fluid volume fraction (CSFVF) given by NODDI (Zhang et al., 2012). GFA and DKI scalars were computed using Dipy (Garyfallidis et al., 2014). NODDI scalars were computed using AMICO (Daducci et al., 2015).

4. **Absolute difference (AD) map:** The AD map reflects voxel-wise differences between the prediction and the ground truth.
5. **Mean absolute difference (MAD):** MAD was computed as the mean of the AD values in the brain region.
6. **Fiber orientation distribution functions (ODFs):** We generate fiber ODFs using the multi-shell multi-tissue model described in Yap et al. (2016).

3.6. Results – BCP multi-shell data prediction

3.6.1. Quantitative evaluation—Results shown in Figs. 3 and 4 indicate that both the ML and SVL significantly improve the average PSNR and SSIM for GFA, MK, and NODDI scalars (OD, ICVF, and CSFVF combined). Compared with the SVL, the ML performs better in improving the estimation of GFA, but not MK and NODDI scalars.

3.6.2. Influence of λ_{ML} and λ_{SVL} —Fig. 5 indicates that, with the increase of λ_{ML} , the GFA PSNR reaches maximum and then slightly fluctuates, while the MK PSNR continues to increase. The NODDI PSNR peaks at $\lambda_{ML} = 0.5$. We chose $\lambda_{ML} = 0.5$ based on NODDI scalars since GFA and MK improvements are already significant. Fig. 6 shows that the NODDI PSNR peaks at $\lambda_{SVL} = 1.0$. Experiments were performed using the dataset for MS data prediction.

3.6.3. Qualitative evaluation—We randomly selected one subject for qualitative inspection. Fig. 7 indicates that the ML gives the GFA image that is closest to the ground truth. This is further confirmed by the AD maps shown in the figure. The SVL gives the second-best AD map with smaller AD values. The MAD values indicate that the ML and SVL yield better performance than the ℓ_1 , ℓ_2 , and Huber losses.

We further investigated the influence of different losses on fiber ODFs. The results, shown in Fig. 8, indicate that the SVL and ML give clean and coherent fiber ODFs that are closer to the ground truth than the ℓ_1 , ℓ_2 , and Huber losses.

3.7. Results – HCP angular resolution enhancement

3.7.1. Quantitative evaluation—The quantitative results, shown in Figs. 9 and 10, are consistent with our observations in Figs. 3 and 4, where the ML and SVL outperform baseline losses. Specifically, both the ML and SVL improve the GFA, MK, and NODDI results in comparison with the baseline losses. The ML gives the best results for GFA and the SVL gives the best results for MK and NODDI. Note that the optimal parameters determined in Section 3.6.2 were used.

3.7.2. Qualitative evaluation—Fig. 11 indicates that the ML yields the GFA image that is closest to the ground truth. The SVL is the second-best loss and reduces the AD values significantly.

Fig. 12 indicates that the fiber ODFs given by the SVL and ML are closer to the ground truth than the baseline losses. The improved fiber ODFs given by our losses can benefit subsequent processing, such as tractography (Aydogan and Shi, 2016) and connectomics (Li et al., 2016).

4. Discussion

A straightforward approach to take into account diffusion properties in the designing of loss functions is by directly computing the microstructure estimations from the predicted data. However, the fitting of microstructural models can be non-linear and the computational complexity can be too high for efficient training of deep neural networks, especially considering the fact the loss function needs to be repeatedly computed for numerous times. This limitation is overcome by the ML and SVL, which are computationally efficient.

In addition to the spherical variance, other spherical moments can be considered to design the loss. A typical example is the spherical mean, which is highly related to the tissue properties of brain white matter (Kaden et al., 2016). However, our evaluation shows that incorporating spherical mean into the loss affects the results only marginally. Currently, we did not consider other high-order spherical moments since (i) their connection with the diffusion properties of brain white matter is unclear and (ii) they are sensitive to noise (Flusser et al., 2009).

The advantages arising from our losses are further reflected by the derivatives of the predicted DMRI data. For example, by improving fiber ODFs via MS data prediction (Fig. 8), our losses improve the estimation of apparent fiber densities (Raffelt et al., 2012) (Fig. 13), orientations (Fig. 14), and (Fig. 14), and tractography (Fig. 15).

5. Conclusion

In this paper, we have proposed two novel loss functions for deep learning prediction of DMRI data. The proposed losses sufficiently consider the microstructural properties of brain white matter during data prediction, avoiding the drawbacks of existing losses that only consider signal prediction error. We demonstrated the effectiveness of our losses, the ML and the SVL, by applying them to the prediction of 6-shell DMRI data from their 2-shell counterparts and the angular resolution enhancement of q -space downsampled DMRI data. Experimental results demonstrate that the ML and SVL substantially outperform standard loss functions.

Acknowledgments

This work was supported in part by the United States National Institutes of Health (NIH) under grants MH125479, EB008374, and EB006733.

Data availability

The authors do not have permission to share data.

References

- Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, Devin M, Ghemawat S, Irving G, Isard M, et al., 2016. Tensorflow: A system for large-scale machine learning. In: 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16). pp. 265–283.
- Aggarwal HK, Mani MP, Jacob M, 2019. MoDL-MUSSELS: model-based deep learning for multishot sensitivity-encoded diffusion MRI. *IEEE Trans. Med. Imaging* 39 (4), 1268–1277. [PubMed: 31603819]
- Andersson JL, Skare S, Ashburner J, 2003. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *NeuroImage* 20 (2), 870–888. [PubMed: 14568458]
- Andersson JL, Sotiropoulos SN, 2016. An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *NeuroImage* 125, 1063–1078. [PubMed: 26481672]
- Aydogan DB, Shi Y, 2016. Probabilistic tractography for topographically organized connectomes. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 201–209.
- Bronstein MM, Bruna J, LeCun Y, Szlam A, Vandergheynst P, 2017. Geometric deep learning: going beyond euclidean data. *IEEE Signal Process. Mag.* 34 (4), 18–42.
- Chen G, Dong B, Zhang Y, Lin W, Shen D, Yap P-T, 2019a. Denoising of infant diffusion MRI data via graph framelet matching in x - q space. *IEEE Trans. Med. Imaging*.
- Chen G, Dong B, Zhang Y, Lin W, Shen D, Yap P-T, 2019b. XQ-SR: Joint x - q space super-resolution with application to infant diffusion MRI. *Med. Image Anal.* 57, 44–55. [PubMed: 31279215]
- Daducci A, Canales-Rodríguez EJ, Zhang H, Dyrby TB, Alexander DC, Thiran J-P, 2015. Accelerated microstructure imaging via convex optimization (AMICO) from diffusion MRI data. *NeuroImage* 105, 32–44. [PubMed: 25462697]
- Dosovitskiy A, Brox T, 2016. Generating images with perceptual similarity metrics based on deep networks. In: Advances in Neural Information Processing Systems. pp. 658–666.
- Fadnavis S, Batson J, Garyfallidis E, 2020. Patch2Self: Denoising diffusion MRI with self-supervised learning. *Adv. Neural Inf. Process. Syst.* 33, 16293–16303.
- Flusser J, Zitova B, Suk T, 2009. Moments and Moment Invariants in Pattern Recognition. John Wiley & Sons.
- Garyfallidis E, Brett M, Amirbekian B, Rokem A, Van Der Walt S, Descoteaux M, Nimmo-Smith I, 2014. Dipy, a library for the analysis of diffusion MRI data. *Front. Neuroinform.* 8, 8. [PubMed: 24600385]
- Gibbons EK, Hodgson KK, Chaudhari AS, Richards LG, Majersik JJ, Adluru G, DiBella EV, 2019. Simultaneous NODDI and GFA parameter map generation from subsampled q -space imaging using deep learning. *Magn. Reson. Med.* 81 (4), 2399–2411. [PubMed: 30426558]
- Golkov V, Dosovitskiy A, Sperl JI, Menzel MI, Czisch M, Sämann P, Brox T, Cremers D, 2016. q -space deep learning: twelve-fold shorter and model-free diffusion MRI scans. *IEEE Trans. Med. Imaging* 35 (5), 1344–1351. [PubMed: 27071165]
- Hong Y, Chen G, Yap P-T, Shen D, 2019a. Multifold acceleration of diffusion MRI via deep learning reconstruction from slice-undersampled data. In: International Conference on Information Processing in Medical Imaging. Springer, pp. 530–541.
- Hong Y, Kim J, Chen G, Lin W, Yap P-T, Shen D, 2019b. Longitudinal prediction of infant diffusion MRI data via graph convolutional adversarial networks. *IEEE Trans. Med. Imaging* 38 (12), 2717–2725. [PubMed: 30990424]
- Howell BR, Styner MA, Gao W, Yap P-T, Wang L, Baluyot K, Yacoub E, Chen G, Potts T, Salzwedel A, Li G, Gilmore JH, Piven J, Smith JK, Shen D, Ugurbil K, Zhu H, Lin W, Ellison JT, 2019.

- The UNC/UMN Baby Connectome Project (BCP): An overview of the study design and protocol development. *NeuroImage* 185, 891–905. [PubMed: 29578031]
- Huber PJ, 1992. Robust estimation of a location parameter. In: *Breakthroughs in Statistics*. Springer, pp. 492–518.
- Huynh KM, Xu T, Wu Y, Chen G, Thung K-H, Wu H, Lin W, Shen D, Yap P-T, Consortium UBCP, et al., 2019. Probing brain micro-architecture by orientation distribution invariant identification of diffusion compartments. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 547–555.
- Jensen JH, Helpert JA, Ramani A, Lu H, Kaczynski K, 2005. Diffusional kurtosis imaging: the quantification of non-gaussian water diffusion by means of magnetic resonance imaging. *Magn. Reson. Med.* 53 (6), 1432–1440. [PubMed: 15906300]
- Johnson J, Alahi A, Fei-Fei L, 2016. Perceptual losses for real-time style transfer and super-resolution. In: *European Conference on Computer Vision*. Springer, pp. 694–711.
- Kaden E, Kruggel F, Alexander DC, 2016. Quantitative mapping of the per-axon diffusion coefficients in brain white matter. *Magn. Reson. Med.* 75 (4), 1752–1763. [PubMed: 25974332]
- Kim J, Hong Y, Chen G, Lin W, Yap P-T, Shen D, 2018. Graph-based deep learning for prediction of longitudinal infant diffusion MRI data. In: *MICCAI Workshop on Computational Diffusion MRI*. Springer, pp. 133–141.
- Kingma DP, Ba J, 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Koppers S, Friedrichs M, Merhof D, 2017. Reconstruction of diffusion anisotropies using 3D deep convolutional neural networks in diffusion imaging. In: *Modeling, Analysis, and Visualization of Anisotropy* Springer, pp. 393–404.
- Koppers S, Haarbuerger C, Merhof D, 2016. Diffusion MRI signal augmentation: from single shell to multi shell with deep learning. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 61–70.
- LeCun Y, Bengio Y, Hinton G, 2015. Deep learning. *Nature* 521 (7553), 436. [PubMed: 26017442]
- LeCun Y, Bottou L, Bengio Y, Haffner P, et al. , 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86 (11), 2278–2324.
- Ledig C, Theis L, Huszár F, Caballero J, Cunningham A, Acosta A, Aitken AP, Tejani A, Totz J, Wang Z, et al. , 2017. Photo-realistic single image super-resolution using a generative adversarial network. In: *Computer Vision and Pattern Recognition, 2017. CVPR 2017. IEEE Computer Society Conference on*, Vol. 2. p.4.
- Li J, Shi Y, Toga AW, 2016. Mapping brain anatomical connectivity using diffusion magnetic resonance imaging: structural connectivity of the human brain. *IEEE Signal Process. Mag.* 33 (3), 36–51. [PubMed: 27212872]
- Lyon M, Armitage P, Álvarez MA, 2022. Angular super-resolution in diffusion MRI with a 3D recurrent convolutional autoencoder. arXiv preprint arXiv:2203.15598.
- Mani MP, Aggarwal HK, Ghosh S, Jacob M, 2020. Model-based deep learning for reconstruction of joint k-q under-sampled high resolution diffusion MRI. In: *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. IEEE, pp. 913–916.
- Minh Huynh K, Chen G, Wu Y, Shen D, Yap P-T, 2019. Multi-site harmonization of diffusion MRI data via method of moments. *IEEE Trans. Med. Imaging*.
- Mirzaalian H, Ning L, Savadjiev P, Pasternak O, Bouix S, Michailovich O, Grant G, Marx CE, Morey RA, Flashman LA, et al. , 2016. Inter-site and inter-scanner diffusion MRI data harmonization. *NeuroImage* 135, 311–323. [PubMed: 27138209]
- Mirzaalian H, Ning L, Savadjiev P, Pasternak O, Bouix S, Michailovich O, Karmacharya S, Grant G, Marx CE, Morey RA, et al. , 2018. Multi-site harmonization of diffusion MRI data in a registration framework. *Brain Imaging Behav.* 12 (1), 284–295. [PubMed: 28176263]
- Moyer D, Ver Steeg G, Tax CM, Thompson PM, 2020. Scanner invariant representations for diffusion MRI harmonization. *Magn. Reson. Med.* 84 (4), 2174–2189. [PubMed: 32250475]
- Ning L, Bonet-Carne E, Grussu F, Seppehrband F, Kaden E, Veraart J, Blumberg SB, Khoo CS, Palombo M, Kokkinos I, et al. , 2020. Cross-scanner and cross-protocol multi-shell diffusion MRI data harmonization: Algorithms and results. *Neuroimage* 221, 117128. [PubMed: 32673745]

- Novikov DS, Veraart J, Jelescu IO, Fieremans E, 2018. Rotationally-invariant mapping of scalar and orientational metrics of neuronal microstructure with diffusion MRI. *NeuroImage* 174, 518–538. [PubMed: 29544816]
- Raffelt D, Tournier J-D, Rose S, Ridgway GR, Henderson R, Crozier S, Salvado O, Connelly A, 2012. Apparent fibre density: a novel measure for the analysis of diffusion-weighted magnetic resonance images. *Neuroimage* 59 (4), 3976–3994. [PubMed: 22036682]
- Rumelhart DE, Hinton GE, Williams RJ, et al. , 1988. Learning representations by back-propagating errors. *Cogn. Model.* 5 (3), 1.
- Schultz T, 2012. Learning a reliable estimate of the number of fiber directions in diffusion MRI. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 493–500.
- St-Jean S, Coupé P, Descoteaux M, 2016. Non Local Spatial and Angular Matching: Enabling higher spatial resolution diffusion MRI datasets through adaptive denoising. *Med. Image Anal.* 32, 115–130. [PubMed: 27082655]
- Tanno R, Worrall DE, Ghosh A, Kaden E, Sotiropoulos SN, Criminisi A, Alexander DC, 2017. Bayesian image quality transfer with CNNs: exploring uncertainty in dMRI super-resolution. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 611–619.
- Tax CM, Grussu F, Kaden E, Ning L, Rudrapatna U, Evans CJ, St-Jean S, Leemans A, Koppers S, Merhof D, et al. , 2019. Cross-scanner and cross-protocol diffusion MRI data harmonisation: A benchmark database and evaluation of algorithms. *NeuroImage* 195, 285–299. [PubMed: 30716459]
- Tuch DS, 2004. Q-ball imaging. *Magn. Reson. Med.* 52 (6), 1358–1372. [PubMed: 15562495]
- Van Essen DC, Smith SM, Barch DM, Behrens TE, Yacoub E, Ugurbil K, Consortium W-MH, et al. , 2013. The WU-Minn human connectome project: an overview. *NeuroImage* 80, 62–79. [PubMed: 23684880]
- Wang Z, Bovik AC, Sheikh HR, Simoncelli EP, et al. , 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* 13 (4), 600–612. [PubMed: 15376593]
- Xu T, Chen G, Wu H, Lin W, Shen D, Yap P-T, 2018. Disentangling the effects of anisotropy and orientation dispersion using diffusion spherical mean spectrum imaging. In: *ISMRM*.
- Yap P-T, Chen Y, An H, Yang Y, Gilmore JH, Lin W, Shen D, 2011. SPHERE: SPHERical harmonic elastic REgistration of HARDI data. *NeuroImage* 55 (2), 545–556. [PubMed: 21147231]
- Yap PT, Zhang Y, Shen D, 2016. Multi-tissue decomposition of diffusion MRI signals via ℓ_1 sparse-group estimation. *IEEE Trans. Image Process.* 25 (9), 4340–4353. [PubMed: 27392357]
- Ye C, 2017. Estimation of tissue microstructure using a deep network inspired by a sparse reconstruction framework. In: *International Conference on Information Processing in Medical Imaging*. Springer, pp. 466–477.
- Ye C, Li X, Chen J, 2019. A deep network for tissue microstructure estimation using modified LSTM units. *Med. Image Anal.* 55, 49–64. [PubMed: 31022640]
- Yin S, You X, Yang X, Peng Q, Zhu Z, Jing X-Y, 2018. A joint space-angle regularization approach for single 4D diffusion image super-resolution. *Magn. Reson. Med.* 80 (5), 2173–2187. [PubMed: 29672917]
- Zhang F, Breger A, Cho KIK, Ning L, Westin C-F, O'Donnell LJ, Pasternak O, 2020. Deep learning based segmentation of brain tissue from diffusion MRI. *BioRxiv*.
- Zhang H, Schneider T, Wheeler-Kingshott CA, Alexander DC, 2012. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. *Neuroimage* 61 (4), 1000–1016. [PubMed: 22484410]
- Zhao H, Gallo O, Frosio I, Kautz J, 2016. Loss functions for image restoration with neural networks. *IEEE Trans. Comput. Imaging* 3 (1), 47–57.

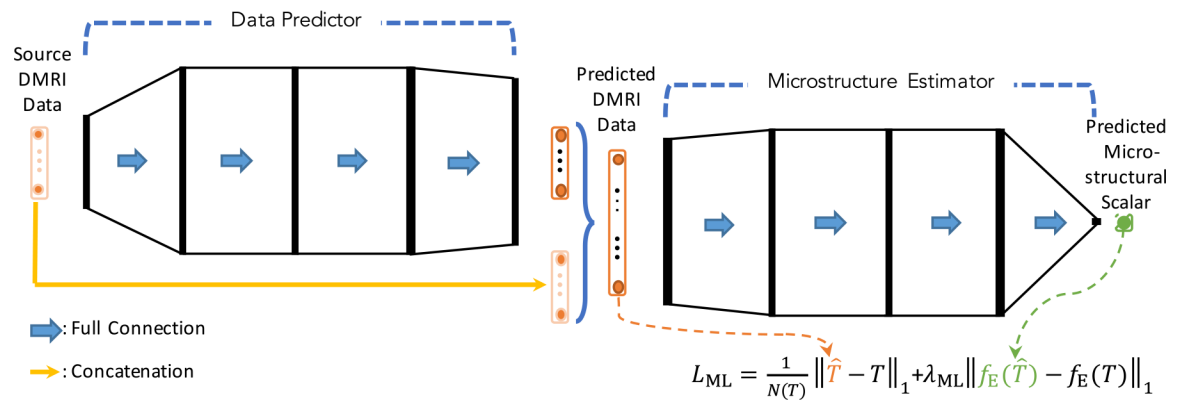


Fig. 1. Network Overview. The network consists of a data predictor and a microstructure estimator. The data predictor, shown on the left, is designed using an MLP. The concatenation branch, marked in yellow, reduces trainable parameters and preserves the source data. The microstructure estimator, shown on the right, is a key component of the ML and is realized using an MLP. f_E , T , and \hat{T} denote the microstructure estimator, target data, and predicted data, respectively.

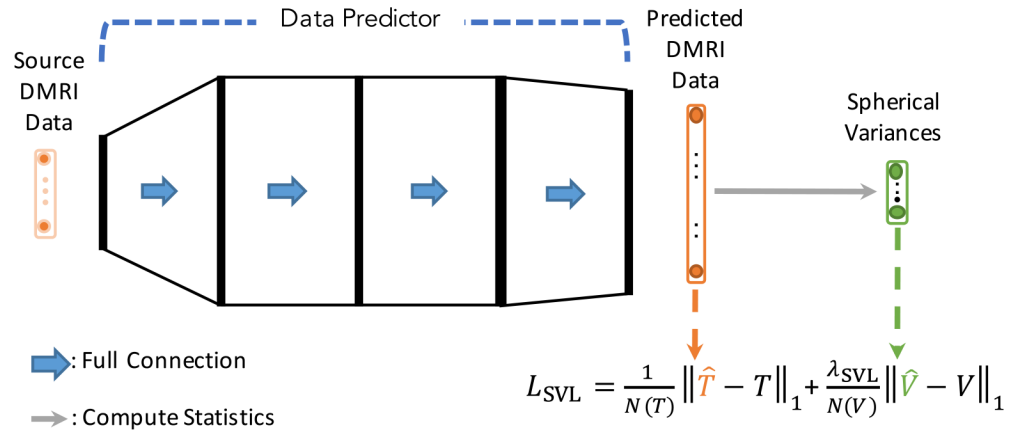


Fig. 2. Spherical variance loss. The spherical variance of the predicted DMRI data is employed to preserve q -space signal sharpness. T and \hat{T} denote the target data (ground truth) and predicted data, respectively. V and \hat{V} are the corresponding spherical variances.

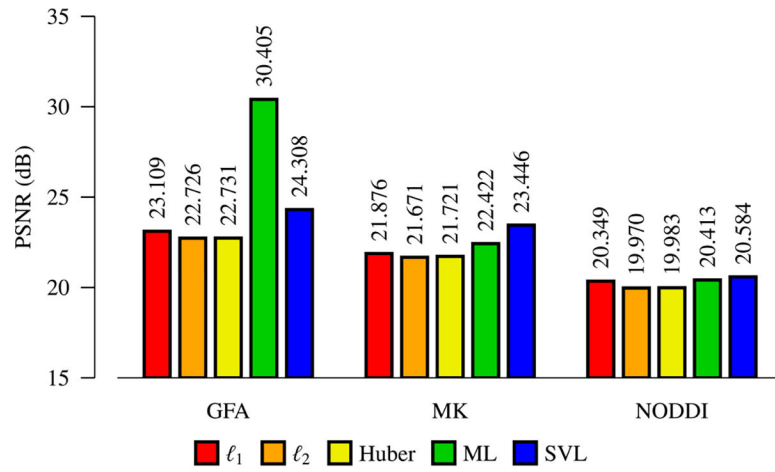


Fig. 3. Quantitative evaluation via PSNR — MS data prediction. PSNR comparison of different losses in terms of GFA, MK, and NODDI.

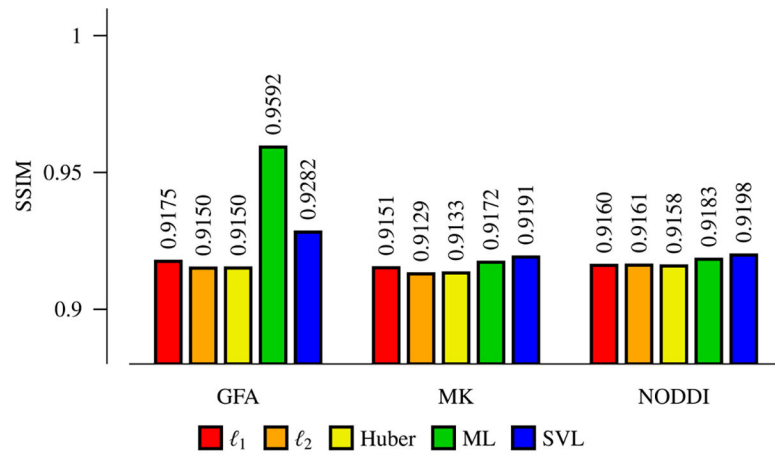


Fig. 4. Quantitative evaluation via SSIM — MS data prediction. SSIM comparison of different losses in terms of GFA, MK, and NODDI.




Fig. 5. Influence of λ_{ML} . PSNR changes with respect to λ_{ML} for GFA, MK, and NODDI. PSNR peaks at $\lambda_{ML} = 0.5$ for NODDI scalars.




Fig. 6. Influence of λ_{SVL} . PSNR changes with respect to λ_{SVL} for GFA, MK, and NODDI. PSNR peaks at $\lambda_{SVL} = 1.0$ for NODDI scalars.




Fig. 7. GFA Images — MS data prediction. Visual comparison of GFA images and associated AD maps. The numbers at top left corners are the mean AD values.




Fig. 8. Fiber ODFs — MS data prediction. The ML and SVL yield fiber ODFs that are closest to the ground truth.




Fig. 9. Quantitative evaluation via PSNR — Angular resolution enhancement. PSNR comparison of different losses in terms of GFA, MK, and NODDI.

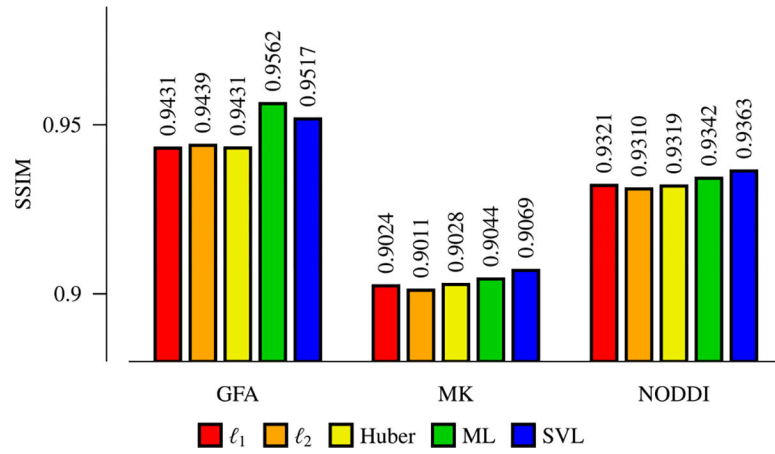


Fig. 10. Quantitative evaluation via SSIM — Angular resolution enhancement. SSIM comparison of different losses in terms of GFA, MK, and NODDI.

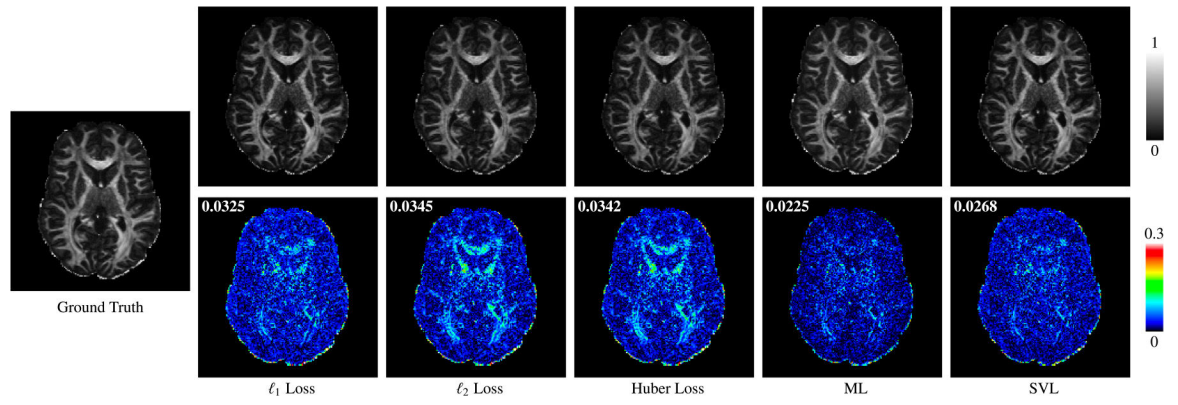


Fig. 11. GFA Images — Angular resolution enhancement. Visual comparison of GFA images and associated AD maps. The numbers at top left corners are the mean AD values.




Fig. 12. Fiber ODFs – Angular resolution enhancement. The ML and SVL yield fiber ODFs that are closest to the ground truth.

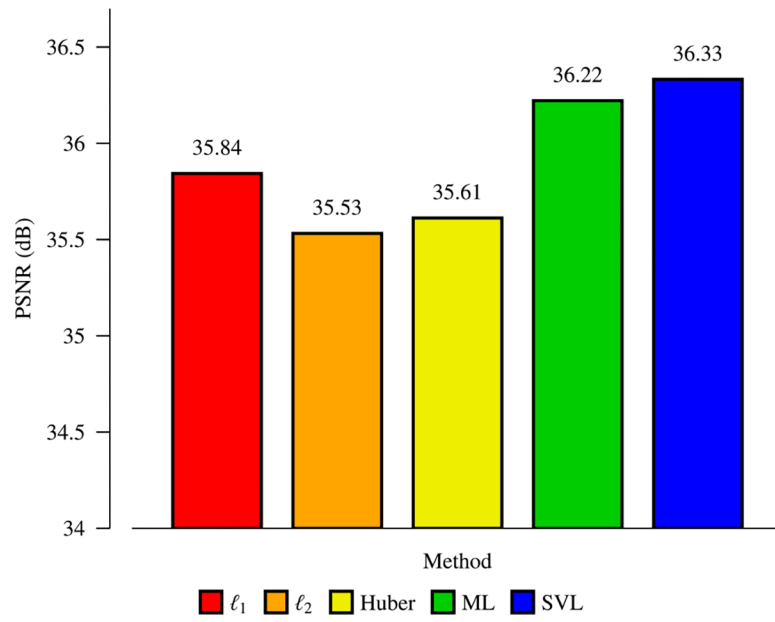


Fig. 13. Apparent fiber densities. PSNR comparison of losses in terms of apparent fiber densities derived from fiber ODFs.

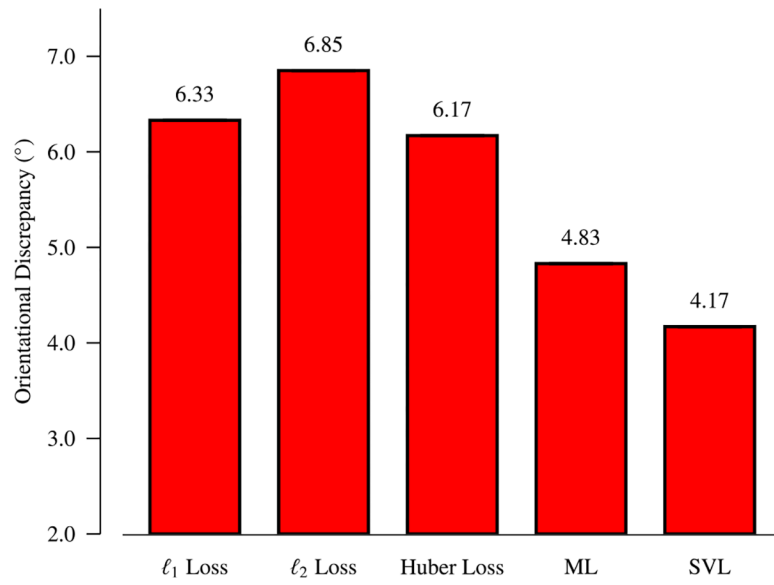


Fig. 14. Axonal orientations. Orientational discrepancy (Yap et al., 2011) of fiber ODFs.

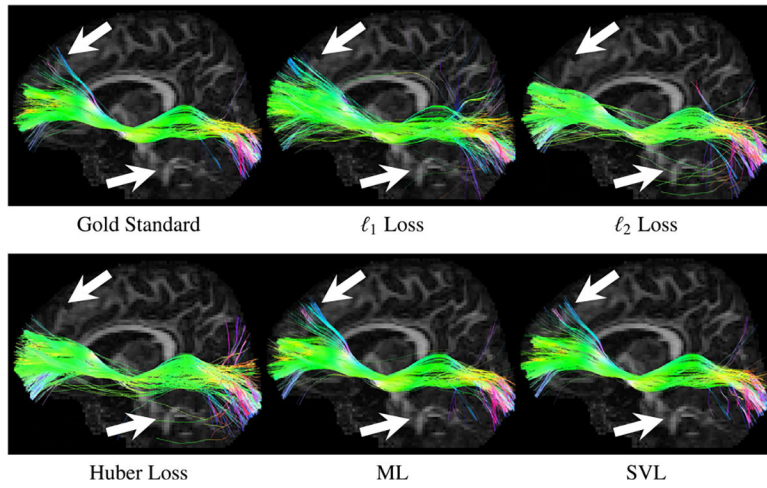
Inferior Fronto-Occipital Fasciculus (IFO)

Fig. 15. Tractography. Comparison of the inferior fronto-occipital fasciculus (IFO) given by the different losses.