


Geometric Data Analysis Project Report

Mathis Reymond
mathis.reymond74@gmail.com
ENS Paris-Saclay
Saclay, France

Théotime de Charrin
theotime.de.charrin@ens.psl.eu
ENS-PSL
Paris, France



1 INTRODUCTION

In this report, we provide a review of the article titled "Modeling the Shape of the Brain Connectome via Deep Neural Networks" [5]. In this 2023 paper, the authors present a novel Riemannian method aimed at inferring brain connectomes using Diffusion Weighted Imaging (DWI) data acquired through Magnetic Resonance Imaging (MRI).

This method is part of a bigger project from the team [2], which works in finding a robust and multimodal framework to statistically analyze and quantify geometric variability of a population of brain connectomes.

Neural brain connectomes, often represented by fibers tracts, are a representation of whole brains' white matter. They have become a clinical standard for studying and diagnosing neuro-diseases on a non-invasive manner. They allow for both functional studies from autism to Parkinson's Disease [17][3], but also structural and anatomical studies (tractography), mainly pre-surgery neuro-oncology planning[18][10]. Indeed, an accurate and meaningful "connectome" allows physicians to assess the location of vital fiber tracts surrounding a tumor, white matter density of cognitive pathways...

Several imaging modalities are available to produce a brain's tractography, the most useful being DWI/DTI and its variants, all based on Magnetic Resonance Imaging. MRI relies on the principles of nuclear magnetic resonance to create high-resolution images of internal organs and tissues, by subjecting the body to a strong magnetic field. It exploits the behavior of hydrogen protons within




Figure 1: Different usecases for a single brain connectome. Adapted from [6]

water molecules, aligning them to the magnetic field, then perturbing this alignment to generate signals that are translated into detailed cross-sectional images.

Building on this core principle, DWI specifically measures the diffusion of water molecules in axonal fibers. As a reminder, the axon is the transmitting part of the neuron, a long "cylinder" surrounded by a myelin sheath (which appears white). These axons represent the information highway in our body. Thus, disentangling the $90 \cdot 10^9$ neurons of the brain into distinct tracts wiring one region to another seems quite vital to have a structural and functional understanding of our neural system.

DWI relies on the fact that water molecules are more likely to diffuse along the long axis of the fibers. Specifically, it allows to extract quantitative measures of the direction of the neuronal tracts.

However, one of MRI's (*a fortiori* DWI) main downsides is that its sequence acquisition and results are highly dependant on the device used, leaving a backdoor to mis-registration with other anatomical images. Furthermore, DWI suffers from several technical and financial limitations such as poor resolution, and inaccuracies generated

by a rather low Signal-to-Noise Ratio (SNR)[10]. All these factors combined together make it currently difficult to compare on a statistical level between hospitals, countries... There is a need for further refinement of DWI tractography methods, by reducing voxel size, reducing distortion, and developing better techniques for dealing with crossing fibers and multiple fiber directions.

Here, the authors address these challenges by proposing to apply a Riemannian geometry framework along with deep neural network's abilities to efficiently solve partial differential equations. Indeed, studying the brain's shape (mostly its white matter pathways) should be similar to studying any shape, and a Riemannian framework seems particularly relevant. Moreover, constructing a Riemannian manifold for the connectome should give access to a very ideal toolbox for later statistical analysis. There are a lot of tools available one can use to calculate an average brain connectome, which quantifies the geometric variability over a population [2]. To construct such a manifold, one must find a global Riemannian metric compatible with the DTI.

2 MATHEMATICAL BACKGROUND

Usually, classic methods for tractography rely on computing the integral curves over the DTI vector field, to find the most likely fiber tract for each voxel. However, such a technique is highly sensitive to noise, impossible to avoid in DWI measurements. For this reason, the authors have successfully combined deep learning models with the Riemann manifold.

A Riemann manifold is a real, differentiable mathematical manifold, namely a sort of domain where we can perform some differentiation and apply ML techniques. The shape of the Riemann manifold can infer the white matter pathways and, together with deep learning, it is possible to have a model to find the perfect curve that fits for shaping a connectome by computing its geodesics.

If the connectome lives in a Riemann manifold it means that we can have Riemannian metrics. A Riemannian metric can tell us how to find the right path for satisfying and creating a connectome; the way to find this perfect path being to solve a partial differential equation (PDE). The solution for this PDE is geodesic curves, a family of curves that satisfy the Riemannian metrics. Since this problem is hard to be solved, the authors reversed the question point of view.

What we need is to find a Riemannian metric on the manifold which minimizes a given energy functional form. When the minimization is achieved, it means we have found the right geodesic, namely the right path/connectome, which fits our mathematical domain. Unlike tractographic methods based on following the flow of principal eigenvectors of the diffusion tensor, these geodesic paths are well-defined even in regions where the tensor diffusion is isotropic.

This problem can be nowadays solved with deep learning models such as PINNs (physics-informed neural networks)[24], which can estimate the solution to PDE problems. In particular, the authors have employed the extension of PINN, namely the convolutional encoder-decoder neural network (CEDNN), which constructs multi-scale features from high-dimensional input.

The main objective of this neural network is to perform the optimization of a loss function between the ground truth data and the estimated metric. Hence, the CEDNN minimizes the following loss function :

$$\varepsilon(g) = \sum_{i=1}^m \|\nabla_{v_i}^g v_i - \sigma_i v_i\|_2 + \alpha \text{Reg}(g) \quad (1)$$

Here, the loss ensures that the DWI vector fields v_i are close to being geodesic vector fields under the current estimated metric g . The penalization term is said to be of no particular matter, but should ensure that the estimated metric is not too aberrant.

This representation in the form of a Riemannian metric for each individual facilitates the process of analyzing connectomes, viewing them as points in an infinite-dimensional manifold. After equipping this space with a natural metric structure, one can apply object-oriented statistical analysis to define an atlas as the Fréchet mean of the population of Riemannian metrics.

3 ALGORITHM

DWI data appears as a vector field where a diffusion orientation distribution function (dODF) is attached to each voxel, providing the directions of water diffusion. These vector fields are obtained through any state of the art tractography method, namely DSI-Studio [22] or MICA [4] pipelines. They can be interpreted as fiber-tangent directions at each voxel. Here, the authors are not interested in computing this distribution function, but rather using it directly to compute the projection onto the Riemannian space. Thus, and for the sake of "proof of concept", here the authors only use the most basic analysis available : eigen decomposition of the diffusion tensor. The interesting part is the analysis of the vector field they propose, and the way they identify the Riemannian metric whose geodesics most accurately match the DWI data.

Their architecture is able to take as input several vector fields from a single individual (unlike previous methods), either obtained through different imaging modalities, or consecutive DWI acquisitions. Then, the network presents a synergistic use of 3 distinct tricks :

- (1) the PINNs [14][24] architecture we previously introduced. By hard-coding some real-world constraints such as semi-definitive positiveness and using surrogate models, one is able to significantly speed up the resolution of PDEs without any labeled data. Originally, this architecture was designed to solve determinist PDEs such as steady-state fluid dynamics, by using as surrogate a flow-based generative model. Changing the pressure and viscosity by the fibers flow along a vector field, one can clearly make a usecase of this concept in tractography.
- (2) the DenseNET architecture [12], successor of ResNET. By concatenating every previous channel to the next one, this deep learning network is able to "learn" on the residuals and have a strong gradient flow. This compact model allows for a better efficiency as layers can be smaller and low-dimension. This improves the information flow of the network and allows for a more computationally efficient system.

- (3) and an Encoder-Decoder architecture, in order to build the features from a very high-dimensional space, and then decrease to a more computational-friendly space.

The output of the neural network is supposed to be directly the estimated Riemannian metric, which will be used to compute the exponential map implemented in the loss function. However, directly solving the metric was long and fastidious, and didn't make use of the fact that the output is supposed to be an SPD matrix. The authors preferred instead to output the Rodrigues eigendecomposition of the Riemannian metric. This allows them to directly compute the exponential map, and compare the estimated geodesics with the ground truth integral curves. Thus, they report a 5-fold change in training efficiency.

The Riemannian metric having the best fit with the DWI data is then supposed to represent the data more accurately, and can be used to register the corresponding tractography to an anatomical atlas.




Figure 2: Overall architecture of the network

4 EXPERIMENTAL METHOD

4.1 Ground truth

Authors use integral curves (denoted P) of a specific vector field as ground truth tracts for the experiments. In the braid experiment, this vector field is an artefact, while in the 2D and 3D experiments, it is derived from Human Connectome Project (HCP) datasets, using state of the art method provided by DSI Studio as mentioned above.

4.2 Evaluation metric

In order to assess the accuracy of the different tractography methods tested in the 2D Brain Slices experiment, the authors had to find a common metric that compares ground truth tracts to inferred geodesic tracts originating from the same seed. This metric can be interpreted as an error.

To calculate this error between curves Q (inferred geodesics) and curve P (ground truth), $\{P\}$ and $\{Q\}$ need to be seen as finite point sets, which allows to consider the mean minimal error between these two sets.

$$\text{Error}(P, Q) = \frac{1}{|P|} \sum_{p \in P} \min_{q \in Q} \|p - q\|_2^2 \quad (2)$$

where P denotes an integral curve and Q a geodesic tractography curve starting at the same seed point. Finally, authors plot this error for geodesics inferred with their method along with other Riemannian methods.

4.3 Results

4.3.1 2D tractography. The authors conducted a quantitative experiment focused on 2D tractography. They compared their method to three state-of-the-art geodesic tractography methods: the inverted diffusion tensor metric [16], the adjugate of the diffusion tensor [7], and the conformal metric [9]. See Figure 3. Using a boxplot, they illustrated the average error of each method, calculated from geodesics derived from 400 distinct seed points. Their analysis of the boxplot led them to conclude that their method outperforms the others in terms of ability to construct a Riemannian metric whose geodesic are close to ground truth fiber tracts.

However, given the high variance in the error calculations from each tract reconstruction, drawing definitive conclusions becomes challenging. This difficulty is compounded by their previous comparison in an earlier version of the article, which displayed significantly poorer results for the inverted and adjugate methods than for the conformal and proposed ones.

4.3.2 PINN vs CEDNN. Additionally, in order to demonstrate CEDNNs ability to perform crossing tracts reconstruction, the authors generate an artificial vector field with crossing-braid shape, and they compute the optimized Riemannian metric using their algorithm two times : one time with a PINN model and another with a CEDNN model. The point is to compare the resulting geodesics. Finally, along with the integral curve, they plot a geodesic associated to the CEDNN-architecture metric and the one associated to standard PINN-architecture metric originating from the same seed. While qualitative assessment on the resulting figure indicates that CEDNN yields a better geodesic, it remains unclear whether this discrepancy reflects a specific struggle of the PINN model in dealing with crossing fibers or if the PINN model is generally weaker than the CEDNN at this task. A valuable complementary experiment would involve comparing the models' abilities to infer accurate geodesics associated with non-crossing tracts. This additional comparison could provide a clearer understanding of the models' strengths and weaknesses in various tractography scenarios. We're currently working on this subject, and hope to present interesting results during the viva.

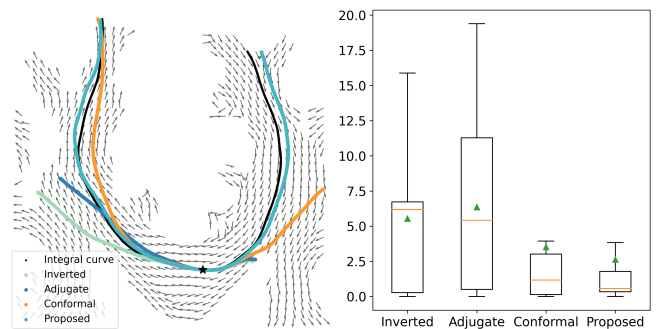


Figure 3: 2D experiment. On the left, detailed view of the geodesics and integral curve starting from the star. On the right, distances between integral curves and the different SOTA methods evaluated.

4.3.3 *Varying the activation function.* Additionally, the authors performed the log-scaled loss convergence comparison between both neural network architectures for solving the inverse problem of estimating a Riemannian metric from two synthetic vector fields in a "braid" pattern, using various activation functions. The resulting figure shows that the convolutional encoder-decoder neural network (CEDNN) with LeakyReLU activation function converges much faster and to a lower loss (by order of magnitudes) than the "vanilla" physics-informed neural network (PINN) with LeakyReLU, Siren, or Fourier embedding. Thus, they conclude that the CEDNN is more efficient and accurate for this problem than the PINN.

Here is how they proceeded :

- (1) Synthesize two vector fields in a "braid" pattern by translating two trigonometric functions across nine pixels horizontally and computing their tangent vectors at each point.
- (2) Configure the CEDNN and the PINN with approximately the same number of parameters and different activation functions or embeddings.
- (3) Train both models, using the synthetic vector fields as the input data.
- (4) Plot the log-scaled loss values at each iteration for each network architecture and compare their performance.

5 CODE

5.1 Struggle

One of the strengths of the paper lies in the authors' inclusion of their code, available on their [GitHub](#) repository. Initially, we intended to review this code, execute it, and ideally offer complementary insights into the data. However, despite investing hours of effort, we encountered difficulties replicating the authors' environment. There were primarily two reasons for this:

- (1) Firstly, both Anaconda and pip failed to resolve conflicts between libraries: scikit_image 0.18.3, skimage 0.0, and torch 1.10.2.
- (2) According to the authors' experience, the model required approximately an hour to train on an Nvidia Titan RTX GPU. Consequently, running it on a CPU was unfeasible. Moreover, to execute the code on our available GPU, installing CUDA was necessary, a process known for its intricacies. Regrettably, despite our attempts, we were unable to properly define the torch optimizer (neither Adadelta nor Adagrad), encountering errors such as "ValueError: networkx.__spec__ is None" and "ValueError: pandas.__spec__ is None." Despite seeking help from both ChatGPT and online forums, these issues remained unresolved.

We attempted to execute a simpler experiment using braid surrogate data but faced the same issues. Training the authors' model would have required a deep understanding of PyTorch intricacies and hours to adapt their code to PyTorch new versions.

Ultimately, the hours invested in reviewing the authors' code were not futile. They shed light on the complexity of ensuring code reproducibility when publishing a computer science paper. Managing dependencies, versions, and configurations across diverse platforms or environments often introduces intricate challenges. Additionally, accommodating hardware disparities or unforeseen edge

cases further complicates this task, turning it into a multifaceted endeavor that demands meticulousness. However, it's important to note that achieving reproducibility is not the primary objective for researchers when presenting their work.

5.2 Breaking news

Ultimately, failing to build a proper GPU setup, we adapted the code such that the model trains on our CPU. We wanted, firstly, to reproduce the experiment in which authors compare PINN's and CEDNN's ability to deal with crossing fibers and, secondly, to perform the same experiment on non crossing fibers, to check whether PINN is inherently weaker than CEDNN, or if it is specifically bad at dealing with crossing fibers on sinusoidal curves.

Thus, we trained the CEDNN model on the synthetic braid data discussed earlier, which required only 10 minutes using a laptop CPU. We inferred the optimal Riemannian metric (see left plot of Figure 4). We observe that we get results significantly different from the authors : in our metric, the ellipses are almost perfectly circular in the crossing regions, while ellipses obtained by the authors are flat in these regions and circular at sinus and cosinus extrema. Thus, we expected the geodesic reconstruction to be very different from the geodesic obtained by the authors, as our estimated metric seems broadly different.

Indeed, we shot geodesics originating from two points (see the right plot of Figure 4) and compared them to the associated integral curves, which stand for the ground truth. We observed that our geodesics are far from what we expected them to be, which reveal weaknesses of the training procedure, but we failed at elucidating this divergence.

Finally, we were unable to train the PINN in the same way, as we didn't find any trace of its implementation in the authors' GitHub (both 2D and 3D repositories).




Figure 4: Results from training the CEDNN model on braid data. Left : estimated metric. Right : reconstructed vector fields and geodesics shot from the starred points - black : integral curves, blue : reconstructed geodesics.

6 CRITICISM & LITERATURE

6.1 Two classical approaches

To date, most tractography techniques rely solely on the local vector field, and thus don't really allow for comparison between subjects or even timepoints. Two main methods are used starting from the vector fields: deterministic- and probabilistic tractography.

Both methods can start from a point in the brain and trace white matter connections. However, in general, deterministic methods generate a single fiber connection from the start point, while probabilistic methods aim to detect many possible connections from the start point. Often, deterministic methods are visualized as curved lines (streamlines), while probabilistic methods may output a map of connection probabilities.

Usually, these methods allow for a better single-voxel match, but are really dependant on the estimation of the vector field at a particular voxel, and are thus not really robust to noise. The deterministic pathway might accumulate several errors throughout the direction estimation, while the probabilistic one, not being any sparse, will bring much less information.

That's why global methods have seen lights in the past decades.

Global approaches such as the estimation of a Riemannian metric allow one to get rid of this voxel-dependency, and have a better estimation of the tractography on a larger scale. Going from a local to a global approach obviously affects the estimation of the reconstructed fiber paths, and leads to a higher connection density than their local probabilistic counterparts at the same percentile thresholds. The idea of this approach is to improve the sensitivity/specificity trade-off (we will come back on that).

This change of paradigm, from local to global and discrete to continuous, is not new and mostly led by the Riemannian manifold approach [16]. Several methods were proposed, but none seems to have imposed in the scientific community, as most of current tractography methods still rely on probabilistic tractography [15].

6.2 State of the Art in the Riemannian geometry approach

As discussed previously, the authors evaluated 3 other Riemannian-based methods against their proposed deep learning toolbox.

In the original paper [16], the inverse matrix of the Diffusion Tensor was proposed as an estimation of the metric. However, this rather "naive" approach is intuitive, but lacks robustness regarding negative eigenvalues, and is often not able to estimate the metric in noisy regions and high-curvature areas. When using this simple method, one will often find "holes" in the connectome where the tensor is ill-defined.

To address this issue, the authors themselves [9] proposed a conformal local rescaling of the metric to follow more closely the diffusion tensor principal eigenvalues along any curvature.

Lastly, a rather recent paper [7] proposed to improve the inverse metric by considering the brownian motion in the diffusion analysis not in the Euclidian space, but rather in the space induced by the local Riemannian metric. This allows to have a locally isotrope brownian motion, easier to describe and biologically more relevant. This is achieved simply by estimating the metric as the adjugate of tensor rather than the inverse.

Moreover, this metric can be "sharpened", *i.e.* elevated to some power, in order to increase the global anisotropy. This is applicable to any DTI estimation, but seems particularly powerful with this "adjugate" metric.

However, all of these methods are based on a single DTI methods and can't handle various imaging modalities, thus giving poor

results regarding fibers crossing. We can see this issue has been successfully addressed by their novel deep learning architecture. But several issues remain unanswered.

6.3 Criticism

Most of the critics proposed here are inherent to the choice of a Riemannian geometry and more broadly, of a global approach.

First, it is assumed that the geodesics are the one and only representation of the fiber tracts, but this is not always verified, and is, in our opinion, more of a framework proposal to study the brain. It is important in our eyes to realize that the Riemannian approach must not stand as a replacement of streamline tractography, but rather a genuine generalization that incidentally raises new problems: any two points in the brain are connected by at least one geodesic - *i.e.* geodesic completeness, which makes sense from a mathematical point of view, but might seem weaker from a medical standpoint.

Some researchers in this rather small community are conscient of this limitation [19], but this doesn't really seem to be the case for the authors here. From the papers we've read, the geodesics potential misfit with ground truth data is never adressed.

In conjunction to that, the authors seem to insist on the fact that their method remarkably handles fibers crossings. From what we've seen, nearly every paper presenting a new methods has the very same claim [1][21][20], and their results are not far from this paper's. Moreover, this might not seem very relevant from a medical point of view, as DWI is limited by its resolution anyway (1.5mm resolution at its absolute best), and pre-operative explorations aim at targeting a more global fiber tract or ax, and not a single fiber. In that sense, crossing fibers don't really matter to know whether a neural pathway goes through a tumor or not.

To moderate this critic, one has to keep in mind that in a DWI image, nearly 60% of voxels are supposed to contain one or several fibers crossings [13], and this number is likely to be underestimated. Moreover, this deep learning tool should not be regarded as a method for solely individual connectomes, but rather a new robust way of analyzing DTI data on a statistical level. When it comes to computing an average connectome atlas, one is really interested in the multiple fiber crossings. With this framework in mind, a substantial number of analyses will be greatly impacted in a cascade effect, from functional studies to proper structural knowledge of the spine.

Moreover, an individual tractography is really nice on the paper, and as proposed in introduction, might have a lot of clinical applications. But this might not be the case in real life, as pointed out by [23]. Indeed, the very nature of DTI makes it very noisy, and even the best individual tractography method will present several false positive fibers (low specificity). They found that state of the art methods presented more invalidated fibers tracts than valid ones, despite a high sensitivity (> 90%). Therefore, they might not be really reliable to date.

Rather, clinical studies mostly use normative connectomes, computed from a very large cohort size on specialized hardware, which are often more accurate and predictive than individual ones from clinical-grade MRIs. These data have indeed been successfully leveraged recently to study the mechanisms of DBS action [11] and also

to explore how data from disparate lesion studies can be integrated to understand the role of brain networks in disease [6], which was supposed to be the very role of individual connectomes.

The authors are conscient of this limitation, and that's why they insist on the fact that this paper is more of a method in a broader framework aiming at robust populational statistical analysis [2]. Again (to hammer the nail), using the metric proposed here would allow for an easy statistical analysis by simply representing the "common" atlas as a Fréchet mean of the individual connectomes. Thus, the specific brain shapes would be well-conserved, and one can compare "diseased" connectomics to healthy ones.

With that goal in mind (statistical atlas construction, principal geodesic analysis...), their deep learning method presented here is welcome, but won't be enough to achieve the ambitious goals from the introduction. It would hugely facilitate the construction of a large human connectome database, robust to registrations and deformations. In synergy with reliable anatomical priors from ex vivo histology, high-resolution post-mortem, or complementary electrophysiology for optimal guidance, tractography methods might become truly reliable.

To that end, a nice multi-modal geometric approach relying on graphs was recently proposed [8]; the conjunction of these 2 methods is really promising, and we might manage to overcome current challenges.

6.4 Discussion and perspectives

Rather than locally computing a fiber path, the authors propose here to find a global Riemannian metric on individuals brains by leveraging the capabilities of Encoder-Decoders, and represent them as points on a manifold, thus allowing for statistical comparison.

Despite several limitations, this paper is a nice method for a more robust analysis of brain connectomes, presenting a successful use of a broad range of deep learning techniques to efficiently analyze the computationally expensive DWI data. They make use of a really interesting mathematical method which is already bulletproof in other computational anatomy frameworks.

Some authors even go further in that matter [19], and propose to get rid of the "restrictive" Riemannian geometry in order to work with more generalistic tools, such as Finsler geometry. While they don't actively present breakthrough results in this paper, it would be interesting in our opinion to see this approach combined with the deep learning tools we presented and a variational approach on metric geometry.

We hope that this kind of approaches generalizes to the whole community, allowing for better statistical comparison, atlas registration, and in the end a strong clinical impact.

REFERENCES

- [1] Temesgen Bihonegn et al. "Geodesic fiber tracking in white matter using activation function". In: *Computer Methods and Programs in Biomedicine* 208 (Sept. 2021), p. 106283. ISSN: 0169-2607. DOI: [10.1016/j.cmpb.2021.106283](https://doi.org/10.1016/j.cmpb.2021.106283). URL: <http://dx.doi.org/10.1016/j.cmpb.2021.106283>.
- [2] Kristen M Campbell et al. "Structural connectome atlas construction in the space of Riemannian metrics". In: *Lecture Notes in Computer Science*. Cham: Springer International Publishing, 2021, pp. 291–303.
- [3] Olga Ciccarelli et al. "Diffusion-based tractography in neurological disorders: concepts, applications, and future developments". en. In: *Lancet Neurol.* 7.8 (2008), pp. 715–727.
- [4] Raúl R. Cruces et al. "Micapipe: A pipeline for multimodal neuroimaging and connectome analysis". In: *NeuroImage* 263 (Nov. 2022), p. 119612. ISSN: 1053-8119. DOI: [10.1016/j.neuroimage.2022.119612](https://doi.org/10.1016/j.neuroimage.2022.119612). URL: <http://dx.doi.org/10.1016/j.neuroimage.2022.119612>.
- [5] Haocheng Dai et al. *Modeling the Shape of the Brain Connectome via Deep Neural Networks*. 2023. arXiv: [2203.06122](https://arxiv.org/abs/2203.06122) [q-bio.NC].
- [6] Michael D Fox. "Mapping symptoms to brain networks with the human connectome". In: *New England Journal of Medicine* 379.23 (2018), pp. 2237–2245.
- [7] Andrea Fuster et al. "Adjugate Diffusion Tensors for Geodesic Tractography in White Matter". In: *Journal of Mathematical Imaging and Vision* 54.1 (June 2015), pp. 1–14. ISSN: 1573-7683. DOI: [10.1007/s10851-015-0586-8](https://doi.org/10.1007/s10851-015-0586-8). URL: <http://dx.doi.org/10.1007/s10851-015-0586-8>.
- [8] Mustafa Burak Gurbuz and Islem Rezik Id. *Deep graph normalizer: A geometric deep learning approach for estimating connective brain templates*. <https://arxiv.org/abs/2012.14131>. Accessed: 2023-12-1.
- [9] Xiang Hao, Ross T. Whitaker, and P. Thomas Fletcher. "Adaptive Riemannian Metrics for Improved Geodesic Tracking of White Matter". In: *Information Processing in Medical Imaging*. Springer Berlin Heidelberg, 2011, pp. 13–24. ISBN: 9783642220920. DOI: [10.1007/978-3-642-22092-0_2](https://doi.org/10.1007/978-3-642-22092-0_2). URL: http://dx.doi.org/10.1007/978-3-642-22092-0_2.
- [10] Jaimie M Henderson. "Connectomic surgery": diffusion tensor imaging (DTI) tractography as a targeting modality for surgical modulation of neural networks". In: *Front. Integr. Neurosci.* 6 (2012).
- [11] Andreas Horn et al. "Connectivity predicts deep brain stimulation outcome in Parkinson disease". In: *Annals of neurology* 82.1 (2017), pp. 67–78.
- [12] Gao Huang et al. *Densely Connected Convolutional Networks*. 2016. DOI: [10.48550/ARXIV.1608.06993](https://arxiv.org/abs/1608.06993). URL: <https://arxiv.org/abs/1608.06993>.
- [13] Ben Jeurissen et al. "Investigating the prevalence of complex fiber configurations in white matter tissue with diffusion magnetic resonance imaging". In: *Human brain mapping* 34.11 (2013), pp. 2747–2766.
- [14] George Em Karniadakis et al. "Physics-informed machine learning". In: *Nature Reviews Physics* 3.6 (May 2021), pp. 422–440. ISSN: 2522-5820. DOI: [10.1038/s42254-021-00314-5](https://doi.org/10.1038/s42254-021-00314-5). URL: <http://dx.doi.org/10.1038/s42254-021-00314-5>.
- [15] Klaus H Maier-Hein et al. "The challenge of mapping the human connectome based on diffusion tractography". en. In: *Nat. Commun.* 8.1 (2017), pp. 1–13.
- [16] Lauren O'Donnell, Steven Haker, and Carl-Fredrik Westin. "New Approaches to Estimation of White Matter Connectivity in Diffusion Tensor MRI: Elliptic PDEs and Geodesics in a Tensor-Warped Space". In: *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2002*. Ed. by Takeyoshi Dohi and Ron Kikinis. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 459–466. ISBN: 978-3-540-45786-2.
- [17] Ashwini Oswal et al. "Neural signatures of hyperdirect pathway activity in Parkinson's disease". In: *Nature Communications* 12 (1 Dec. 2021). ISSN: 20411723.
- [18] Thomas Schult et al. "Fiber visualization for preoperative glioma assessment: Tractography versus local connectivity mapping". en. In: *PLoS One* 14.12 (2019), e0226153.
- [19] Rick Sengers, Luc Florack, and Andrea Fuster. "Geodesic Uncertainty in Diffusion MRI". In: *Frontiers in Computer Science* 3 (Oct. 2021). ISSN: 2624-9898. DOI: [10.3389/fcomp.2021.718131](https://doi.org/10.3389/fcomp.2021.718131). URL: <http://dx.doi.org/10.3389/fcomp.2021.718131>.
- [20] P. Staempfli et al. "Resolving fiber crossing using advanced fast marching tractography based on diffusion tensor imaging". In: *NeuroImage* 30.1 (Mar. 2006), pp. 110–120. ISSN: 1053-8119. DOI: [10.1016/j.neuroimage.2005.09.027](https://doi.org/10.1016/j.neuroimage.2005.09.027). URL: <http://dx.doi.org/10.1016/j.neuroimage.2005.09.027>.
- [21] V.J. Wedeen et al. "Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers". In: *NeuroImage* 41.4 (July 2008), pp. 1267–1277. ISSN: 1053-8119. DOI: [10.1016/j.neuroimage.2008.03.036](https://doi.org/10.1016/j.neuroimage.2008.03.036). URL: <http://dx.doi.org/10.1016/j.neuroimage.2008.03.036>.
- [22] Fang-Cheng Yeh, Pei-Fang Tang, and Wen-Yih Isaac Tseng. "Diffusion MRI connectometry automatically reveals affected fiber pathways in individuals with chronic stroke". In: *NeuroImage: Clinical* 2 (2013), pp. 912–921. ISSN: 2213-1582. DOI: [10.1016/j.nicl.2013.06.014](https://doi.org/10.1016/j.nicl.2013.06.014). URL: <http://dx.doi.org/10.1016/j.nicl.2013.06.014>.
- [23] Fan Zhang et al. "Quantitative mapping of the brain's structural connectivity using diffusion MRI tractography: A review". In: *NeuroImage* 249 (2022), p. 118870. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2021.118870>. URL: <https://www.sciencedirect.com/science/article/pii/S1053811921011411>.
- [24] Yin hao Zhu et al. "Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data". In: *Journal of Computational Physics* 394 (Oct. 2019), pp. 56–81. ISSN: 0021-9991. DOI: [10.1016/j.jcp.2019.05.024](https://doi.org/10.1016/j.jcp.2019.05.024). URL: <http://dx.doi.org/10.1016/j.jcp.2019.05.024>.