

DTI-DeformIt: GENERATING GROUND-TRUTH VALIDATION DATA FOR DIFFUSION TENSOR IMAGE ANALYSIS TASKS

Brian G. Booth and Ghassan Hamarneh

Medical Image Analysis Lab, School of Computing Science, Simon Fraser University, Canada
{bgb2, hamarneh}@sfu.ca

ABSTRACT

We propose DTI-DeformIt: a framework to generate realistic synthetic datasets from a smaller number of, or even one, annotated image(s). Our approach extends the DeformIt technique of Hamarneh et al. [1] to handle the deformations and noise conditions of diffusion tensor images. An implementation of our proposed framework is also provided as a free download. We further show that DTI-DeformIt generates images that, according to eigenvector distance, are no different from real images than other real images, making them suitable for machine learning and validation.

Index Terms— Diffusion Tensor Imaging, Validation, Machine Learning, Image Generation

1. INTRODUCTION

Expertly-annotated image data plays a key role in medical image analysis for two main reasons. First, expert-provided annotations are used to validate automated image analysis algorithms. Notable examples of this include using manually-drawn segmentations as ground truth to validate segmentation and registration algorithms (e.g. via Dice similarity coefficient). Second, expertly-provided annotations can be used as training data for (semi-)supervised image analysis algorithms. The recent popularity of these learning techniques has further increased the desire for annotated training data.

While expertly-annotated data is becoming more and more available in many imaging modalities, there has been little in terms of publicly-available, expertly-annotated data for diffusion tensor MRI. Those that do exist are typically atlases (e.g. ICBM DTI-81 atlas [2]) and only few papers have used manually annotated datasets (e.g., [3]).

To date, validation of DTI analysis algorithms has been limited to physical [4, 5] and computational phantoms [6]. Yet, these phantoms pale in comparison to the complexity seen in real images. SimDTI [7] also provides synthetic datasets but is limited to brain DTI examples. Reproducibility tests using the Kirby21 dataset further provide means to examine an algorithm’s ability to provide consistent results for different images of the same brain [8]. However, this testing does not directly assess the accuracy of the DTI analysis.

The lack of annotated DTI data may be due to the fact that DTI is a relatively new imaging modality that produces high-dimensional, manifold-valued image data. Such data is difficult to visually interpret and manually annotate. Unfortunately, this dearth of annotated image data limits our ability to accurately validate DTI analysis algorithms and our ability to apply machine learning techniques to DTI analysis.

To address this problem, we propose DTI-DeformIt, a framework through which an arbitrarily large set of annotated DTI images can be generated from a small set (even a single dataset) of annotated DTI scans. The framework is based on the original DeformIt algorithm presented in [1] that generates annotated scalar images using a joint physical-statistical generative model. We extend this framework to DTI datasets and show quantitatively, using statistics of the tensors’ primary eigenvectors, that the generated images resemble real images. We further provide a free implementation of DTI-DeformIt and show how the framework’s parameters allow for the generation of datasets of different levels of variability (thereby controlling segmentation difficulty).

2. BACKGROUND: DeformIt

The initial DeformIt framework generates novel images from applying statistically- and physically-based spatial deformations, followed by non-uniform intensity warps and additive noise, to given images [1]. The annotations for the given images are also deformed by the same hybrid statistical-physical deformations. It is these deformations that form the basis of our DTI-DeformIt approach.

The deformations generated by DeformIt are defined on a uniform lattice of $L \times M \times N$ control points $\mathbf{x} = (cx_{ijk}, cy_{ijk}, cz_{ijk}); i \in \{1, \dots, L\}, j \in \{1, \dots, M\}, k \in \{1, \dots, N\}$. These control points are displaced by the statistical and physical models described below and the displacement of any other 3D point interpolated from these control point displacements.

Generative Statistical Model: Given a training set of control points that correspond across the small collection of given images, we can learn the variability of the control point locations. Each lattice of control points is represented by a

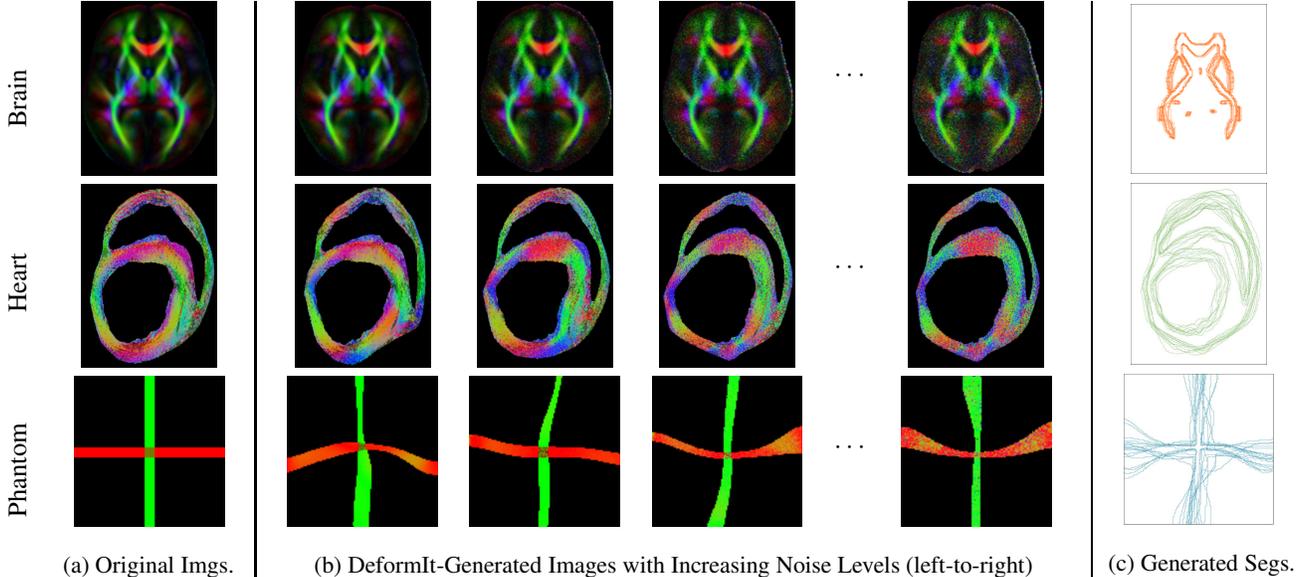


Fig. 1: Sample DTI-DeformIt results for various DT images. The original images in (a) are deformed using a statistical-physical deformation model, followed by tensor reorientation and the introduction of additive noise, to obtain the novel synthetic images in (b). The corresponding segmentations for the novel images are shown as overlaid contours in (c). Note that DTI-DeformIt allows for the ability to control the amount of variability and noise in the generated data.

$3LMN$ -dimensional vector and we compute the sample covariance matrix S of these lattices. Using principal component analysis (PCA), we approximate the known control point locations by the point distribution model $\mathbf{x} = \bar{\mathbf{x}} + \mathbf{P}\mathbf{b}$, where $\bar{\mathbf{x}}$ is the sample mean lattice and \mathbf{P} are the given lattice’s main modes of variation. We randomly select different weights for the modes of variation \mathbf{b} (within ± 3 std. dev.) in order to generate new control point locations for our generated synthetic images.

Generative Physical Model: Using the work of Cootes and Taylor [9], we generate elastic-based modes of variation using a finite element model (FEM). The grid of control points are modelled as masses connected by springs to their spatially-adjacent neighbours. The springs between points are set to be at rest at the given location of the control points, $\hat{\mathbf{x}}$. New locations for the control points can be generated using $\mathbf{x} = \hat{\mathbf{x}} + \Phi\mathbf{u}$, where Φ are the modes of elastic deformation and weights \mathbf{u} are randomly selected to vary the effect of these modes. The modes of variation, Φ , are generated from the solution to the generalized eigensystem that encodes the FEM mass-spring model described above. The details of this computation are provided in [9].

Combined Statistical-Physical Model: A combined generative model for the deformation of the control points is then obtained by a weighted combination of the deformations provided by each model:

$$\mathbf{x} = \hat{\mathbf{x}} + \alpha\mathbf{P}\mathbf{b} + \beta\Phi\mathbf{u}. \quad (1)$$

The weights α, β control the magnitude of each contribu-

tion. Under this hybrid model, the control point grid locations $\hat{\mathbf{x}}$ are set to the mean of the statistical model, $\bar{\mathbf{x}}$.

Once control point deformations are computed using (1), the image deformation field ϕ is generated using thin-plate spline interpolation. After applying ϕ to the image, DeformIt also allows image intensities to be non-uniformly transformed, and noise added, as discussed in [1].

3. METHODS: DTI-DeformIt

Extending the concepts of DeformIt to diffusion tensor data requires addressing three key points: 1) ensuring the generated deformations are diffeomorphic, 2) applying tensor reorientation to the generated images that is in line with the deformation applied to the given image, and 3) adding noise to the generated images appropriately. We contribute these three additions to DeformIt and discuss each below.

Ensuring Diffeomorphic Deformations: A diffeomorphic deformation is one that is smooth and invertible, ensuring that the image does not tear or fold when being deformed. A diffeomorphic deformation is necessary for warping diffusion tensor images as it ensures that the rotations required for tensor reorientation can be numerically computed.

To ensure that the deformations generated by (1) are diffeomorphic, we employ the same post-processing step as diffeomorphic demons: the scaling and squaring approach described in [10]. In brief, the algorithm scales the magnitude of the deformation field down to the point where it is small enough not to fold or tear the image space. Once scaled down,

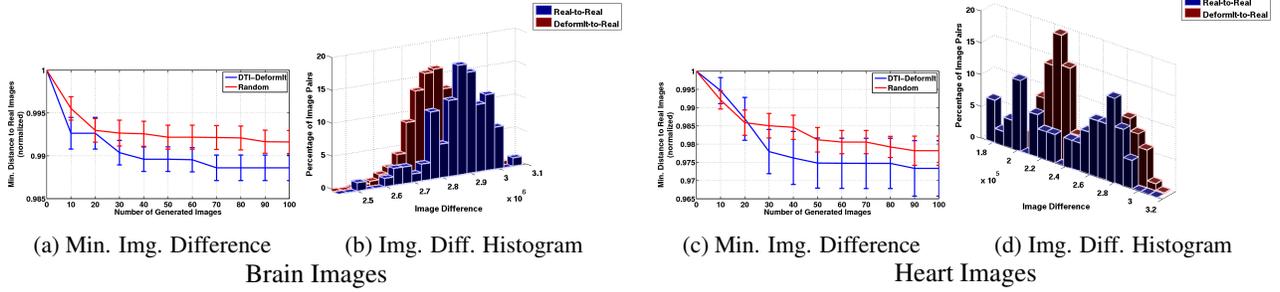


Fig. 2: Comparison of image differences between real image pairs and between those same images and DTI-DeformIt-generated images. (a,c) show the image difference between the real images and their closest DTI-DeformIt-generated image decreases more rapidly than if the applied deformations were randomly generated. (b,d) show the histograms of the image differences between real image pairs and image differences between 100 DTI-DeformIt-Generated images and each real image. Both histograms cover a similar range, suggesting that the images we generate differ from real data no more than other real images.

the deformation is concatenated with itself until its original magnitude is restored. Note that the concatenation of diffeomorphisms results in a diffeomorphism [10].

Tensor Reorientation: As has been well noted in DTI registration, tensors need to be reoriented in step with the image deformation in order to preserve tract connectivity in the DT image. To accomplish this task, we employ the preservation of principal direction (PPD) algorithm of Alexander et al. [11]. The PPD algorithm iteratively realigns the tensor’s eigenvectors by applying a sequence of rotations that aligns each eigenvector e_i to $\mathbf{F}e_i$, where $\mathbf{F} = \mathbf{I} + \mathbf{J}_\phi$ and \mathbf{J}_ϕ is the Jacobian of the deformation vector field ϕ . A full description of the algorithm is given in [11].

Adding Noise to DT Images: Diffusion tensor images are generated by fitting a Gaussian to the diffusion signals obtained using diffusion weighted MR imaging. The Gaussian fitting is performed using the Skejskal-Tanner equation: $S(\mathbf{g}) = S_0 \exp(-b\mathbf{g}^T \mathbf{D} \mathbf{g})$. As imaging noise presents itself in the diffusion-weighted images $S(\mathbf{g})$, we use the Skejskal-Tanner equation to decompose the diffusion tensor into its diffusion-weighted images and add Rician or Gaussian noise at that level. The B_0 image S_0 , the diffusion gradients \mathbf{g} , and the diffusion weighting b are parameters that can be provided by the user if known; otherwise, we provide default values based on the work of Frindel et al. [12] ($S_0 = 1$, $b = 1000$, the six gradient imaging scheme for \mathbf{g}). Once noise has been added to the computed diffusion weighted images, we re-fit the tensors using a linear least squares fitting defined by the logarithm of the Skejskal-Tanner equation.

With these three contributions now in DeformIt, we can take a DTI-segmentation pair and generate synthetic datasets by (a) generating a diffeomorphism, (b) applying it to the image-segmentation pair, (c) performing tensor reorientation, and (d) adding Gaussian or Rician noise. A full MATLAB implementation of DTI-DeformIt is freely provided¹.

4. EXPERIMENTAL ANALYSIS

Fig. 1 shows sample DTI-DeformIt-generated image-segmentation pairs for the ICBM DTI-81 brain atlas, a DTI scan of a canine heart, and a simple crossing phantom. Note that the amount of deformation is one of many parameters that can be controlled so as to be appropriate for the different anatomical structures (e.g. the heart image is deformed more than the brain image). Further note that tensor reorientation is handled as can be seen on the phantom dataset.

Realism of DTI-DeformIt Data: To evaluate the quality of the synthetic images generated by DTI-DeformIt, we compared image differences between the synthetic images we generate and real images from a test set. We hypothesize that these image differences will be on par with those between a pair of real images. For this evaluation, we used both brain and heart DTI data. For the brain data, the ICBM DTI-81 atlas [2] was used as input to DTI-DeformIt and 13 images from the IXI database² were used as the test set. For the heart data, we used 10 healthy canine heart images from the CCBM database³. The first heart image was used as input into DTI-DeformIt while the remaining 9 images composed the test set.

To compare our generated images to the test sets, we employ the same primary eigenvector difference metric over the image space Ω as used by van Hecke et al. [13]:

$$d(e_1^G, e_1^T) = \sum_{x \in \Omega} \cos^{-1}(e_1^G(x) \cdot e_1^T(x)) \quad (2)$$

where e_1^G and e_1^T are the primary eigenvector images from the generated dataset and the test set, respectively. We employ this metric as it allows us to compare fiber structure, as given by the primary eigenvector, without being concerned with differences in diffusion properties that are known to exist between the ICBM DTI-81 atlas and real data (the atlas

¹www.cs.sfu.ca/~hamarneh/software/DeformIt

²<http://brain-development.org>

³www.ccbm.jhu.edu/research/DTMRIDS.php

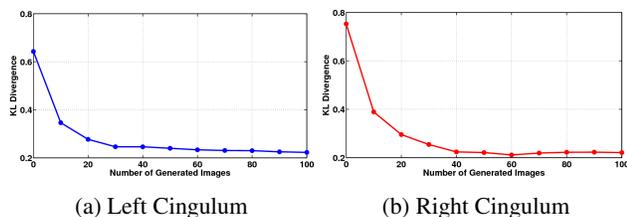


Fig. 3: KL-Divergence between the primary eigenvector distributions learned from ground truth data and DTI-DeformIt-generated data. Note that as the number of DTI-DeformIt-generated images increases, the better we are able to model the primary eigenvector distributions of each region.

has lower FA and higher MD than real data due to image misalignment in the atlas creation step).

Figs. 2a and 2c show that as we generate more images, the image differences between the test images and their closest DTI-DeformIt-generated image decreases more rapidly than if the applied deformations were randomly generated. This suggests that our hybrid statistical-physical deformation model performs better at generating realistic looking images than a random deformation model. Figs. 2b and 2d show the histograms of image differences between images in our test set (blue) and between those same images and 100 DTI-DeformIt-generated images (red). Note that the image differences for both cases cover the same range of values, suggesting that the synthetic images we create are not substantially different than other real images.

Suitability for Machine Learning: To test whether we can use DTI-DeformIt-generated data as training data in a machine learning algorithm for learning regional appearance priors, we compare the data distribution in segmented regions of the generated images to the data distribution in the same, manually-segmented, regions in the test set. In this experiment, we used the generated brain data and its corresponding test set from the IXI database. The left and right cingulum we both segmented manually from the 13 images in our test set. Once again, we use the primary eigenvector for comparison purposes and we compared the distributions of the primary eigenvectors using KL-divergence.

Fig. 3 shows the KL-divergence between the real and synthetic distributions versus the number of DTI-DeformIt-generated images. Note that as we increase the number of generated images, we are better able to match the regional primary eigenvectors distributions, suggesting that DTI-DeformIt-generated data can be used to train machine learning algorithms.

5. CONCLUSION

We have presented here an extension to the DeformIt framework [1] to generate realistic annotated DTI data from a small number of annotated images. The proposed extension en-

ures that the generated hybrid statistical-physical deformations are diffeomorphic, that tensor reorientation is properly performed, and that imaging noise is added at the appropriate diffusion-weighted image level. We have also made a MATLAB implementation of this algorithm freely available and in doing so, have provided the ability to improve validation of DTI analysis techniques and the generation of learning-based DTI analysis algorithms. Future work may include introducing noise models for parallel MRI, adding variability in both diffusion and topology in order to simulate pathology, and encoding spatially varying elasticity parameters reflecting properties of underlying heterogeneous tissue.

6. ACKNOWLEDGEMENTS

We acknowledge Drs. Patrick A. Helm and Raimond L. Winslow at the Center for Cardiovascular Bioinformatics and Modeling, and Dr. Elliot McVeigh at the National Institute of Health, for provision of the heart data. GH is partially supported by NSERC.

7. REFERENCES

- [1] G. Hamarneh, P. Jassi, and L. Tang, "Simulation of ground-truth validation data via physically- and statistically-based warps," in *Proc. of MICCAI*, 2008, pp. 459–467.
- [2] S. Mori et al., "Stereotaxic white matter atlas based on diffusion tensor imaging in an ICBM template," *NeuroImage*, vol. 40, pp. 570–582, 2008.
- [3] B. G. Booth and G. Hamarneh, "A cross-sectional piecewise constant model for segmenting highly curved fiber tracts in diffusion MR images," in *Proc. of MICCAI*, 2013, pp. 469–476.
- [4] J.S.W. Campbell et al., "Flow-based fiber tracking with diffusion tensor and q-ball data: Validation and comparison to principal diffusion direction techniques," *NeuroImage*, vol. 27, pp. 725–736, 2005.
- [5] P. Fillard et al., "Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom," *NeuroImage*, vol. 56, no. 1, pp. 220–234, 2011.
- [6] T. G. Close et al., "A software tool to generate simulated white matter structures for the assessment of fibre-tracking algorithms," *NeuroImage*, vol. 47, no. 4, pp. 1288–1300, 2009.
- [7] W. Van Hecke et al., "On the construction of a ground truth framework for evaluating voxel-based diffusion tensor MRI analysis methods," *NeuroImage*, vol. 46, pp. 692–707, 2009.
- [8] B. A. Landman et al., "Multi-parametric neuroimaging reproducibility: A 3T resource study," *NeuroImage*, vol. 54, no. 4, pp. 2854–2866, 2011.
- [9] T. Cootes and C. Taylor, "Combining point distribution models with shape models based on finite element analysis," *Image and Vision Computing*, vol. 13, no. 5, pp. 403–409, 1995.
- [10] T. Vercauteren et al., "Diffeomorphic demons: Efficient non-parametric image registration," *NeuroImage*, vol. 45, pp. S61–S72, 2009.
- [11] D.C. Alexander, C. Pierpaoli, P.J. Basser, and J.C. Gee, "Spatial transformations of diffusion tensor magnetic resonance images," *IEEE Trans. on Med. Imag.*, vol. 20, no. 11, pp. 1131–1139, 2001.
- [12] C. Frindel et al., "Comparison of regularization methods for human cardiac diffusion tensor MRI," *Med. Imag. Anal.*, vol. 13, pp. 405–418, 2009.
- [13] W. Van Hecke et al., "Nonrigid coregistration of diffusion tensor images using a viscous fluid model and mutual information," *IEEE Trans. Med. Imag.*, vol. 26, no. 11, pp. 1598–1612, 2007.