Deep Learning Multi-channel Structural and Diffusion Tensor Neonatal Image Registration

Irina Grigorescu[1(✉)], Alena Uus[1], Daan Christiaens[1,2], Lucilio Cordero-Grande[1,4], Jana Hutter[1], Dafnis Batalle[1,3], A. David Edwards[1], Joseph V. Hajnal[1], Marc Modat[1], and Maria Deprez[1]

[1] School of Biomedical Engineering and Imaging Sciences, King's College London, London, UK
irina.grigorescu@kcl.ac.uk
[2] Departments of Electrical Engineering, ESAT/PSI, KU Leuven, Leuven, Belgium
[3] Department of Forensic and Neurodevelopmental Science, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK
[4] Biomedical Image Technologies, ETSI Telecomunicación, Universidad Politécnica de Madrid and CIBER-BNN, ISCIII, Madrid, Spain

Abstract. Investigating the maturation process of the brain through clinical studies is crucial for advancing our understanding of the developing brain. Neuroimaging studies employing magnetic resonance images (MRI) often necessitate pre-alignment of images to a standardized space. Registration of diffusion tensor images (DTI) has the potential to better align white matter (WM) structures than using structural MRI only, as it enables alignment of fiber orientation at each voxel. However, microstructural data are not well suited to accurately register the cortical gray matter (cGM) ribbon. In spite of this, in many studies the alignment of subjects is primarily guided by a single modality. In this work, we propose a multi-channel attention-based deep learning registration approach that selects the most salient features from multiple image modalities to improve alignment of individual MR images to a common atlas space. We apply the technique to align multi-channel datasets composed of structural T_2-weighted (T_2w) MRI and DTI maps into atlas space. The quantitative and qualitative evaluation confirmed that when we use the two modalities we obtain good alignment of anatomical structures, while also improving the alignment of the underlying white matter tracts. Moreover, we show that learning a spatially-varying attention map for weighting the two modalities obtains superior performance than the baseline multi-channel method that does not incorporate attention.

Keywords: multi-channel registration · attention · deep learning registration · diffusion tensor imaging

M. Modat et al. (Eds.): WBIR 2024, LNCS 15249, pp. 192–204, 2024.
https://doi.org/10.1007/978-3-031-73480-9_15

1 Introduction

The neonatal brain undergoes dramatic changes during early life, such as cortical folding, development of white matter (WM) tracts and myelination [8]. Magnetic resonance imaging (MRI) provides snapshots of the evolving morphology and tissue properties in the developing brain across multiple subjects and time-points [23]. However, neuroimaging studies often require that the images are pre-aligned to a common space so that development or pathologies can be studied longitudinally. Accurate anatomical alignment remains one of the key challenges for image registration algorithms [8].

Structural and microstructural MRI modalities offer complementary information about morphology and tissue properties of the developing neonatal brain, but often inter-subject alignment is driven by a single modality, such as structural [4], or diffusion [31] MRI. Moreover, the alignment of both structural T_2-weighted (T_2w) scans and diffusion tensor imaging (DTI) maps to a common reference space holds significant importance in neuroimaging, as T_2w scans are valuable for morphometric purposes, such as brain segmentation, while microstructural data can provide information about the underlying tissue properties.

Multi-channel registration aims to find inter-subject anatomical correspondences using sets of images of multiple modalities for each subject. Most such approaches rely on minimizing the weighted sum of dissimilarity measures calculated for each channel individually [3,6,7,12,18]. One drawback of this approach is that each channel is weighted on a global scale, thus not accounting for local variations within the images. To address this, studies have attempted to weight the deformation fields with certainty maps calculated from normalised spatial gradients correlated to structural content [11,29]. Rohde *et al.* [21] proposed a framework for multi-channel registration using a multivariate correlation approach, and applied it to DTI data. Guimond *et al.* [12] and Park *et al.* [18] introduced the multi-channel Demons algorithm [26], while Peyrat *et al.* [20] further incorporated trajectory constraints for registering 4D time-series cardiac computed tomography (CT) images. Finally, multi-channel registration has also been helpful in multi-modal applications where the target and the moving images consisted of different modalities, and in these cases [5,6] proposed the use of image synthesis [13,15] to generate synthetic images of the missing channels.

In this work, we propose a multi-channel deep learning attention-based image registration framework which combines T_2w images with DTI maps (see Fig. 1). Our proposed solution is based on a diffeomorphic non-rigid registration with stationary velocity field (SVF) representation [14]. Our method aims to optimise the multi-channel alignment through learned spatial attention maps which provide an optimal spatial weighting for the predicted modality-specific velocity fields. More specifically, we propose a lightweight convolutional neural network (CNN) termed **attention velocity network** which produces attention maps which identify salient features in the individual channels in order to drive the multi-channel registration process. To evaluate the proposed framework, we compare networks trained on single modality to multi-channel models, either with or without attention. Our results show that by using attention in a multi-channel

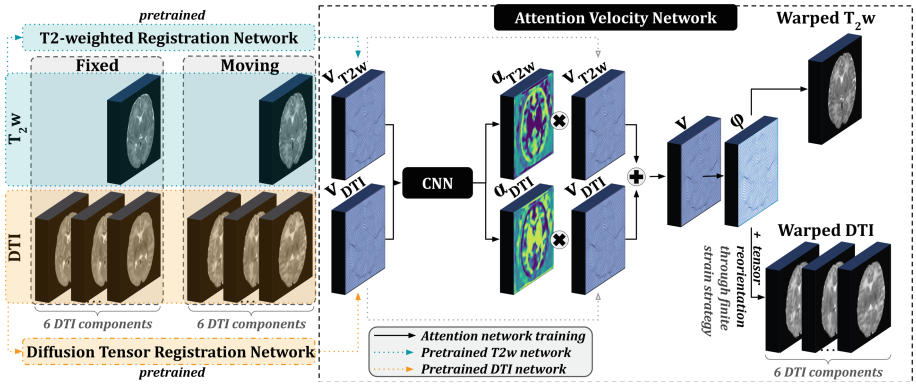

Fig. 1. Pipeline of our proposed multi-channel structural and microstructural **attention**-driven image registration framework which takes as input the modality-specific velocity fields v_{T2w} and v_{DTI}, and outputs the attention maps α_{T2w} and α_{DTI}. The black arrows represent the training and inference flow of data through the proposed attention velocity network, where the final deformation field φ is used to interpolate the moving images into their warped counterparts. Note that the data also undergoes tensor reorientation (*finite strain* strategy [2]). Network architecture details can be found in Fig. 2.

image registration setting we achieve superior overall alignment considering brain morphology as well as the WM microstructure.

2 Methodology

Our proposed **attention**-driven multi-channel deep learning image registration framework aims to combine information from T_2w neonatal scans with diffusion-weighted imaging (DWI)-derived microstructural maps. The overall pipeline is presented in Fig. 1, where the two channels are T_2w and DTI. The **attention velocity network** aims to learn the necessary spatial attention maps which select the most salient features from both structural and microstructural channels. More specifically, given a pair of fixed and moving T_2w and DTI volumes, the network produces two 3D modality-specific spatial attention maps α_{T2w} and α_{DTI} which are used to locally weigh the velocity fields v_{T2w} and v_{DTI}. The latter are obtained from pre-trained T_2w and DTI single-modality image registration networks which are based on a conditional variational autoencoder (CVAE) as proposed by [14]. Note that in Fig. 1 we show our proposed pipeline for T_2w and DTI data, however, these can be replaced with different modalities, such as tissue segmentation maps, or scalar-valued fractional anisotropy (FA) maps.

Throughout this work we use 3D structural (T_2w) and diffusion (DTI) MRI brain scans acquired as part of the developing Human Connectome Project (dHCP) [10] for the moving images, and a 36 weeks old neonatal atlas [29] of

the same modalities as the fixed images. Details about the aquisition parameters are found in [10].

2.1 Data Preprocessing and Selection

Data Preprocessing. We affinely pre-registered all data to a common 36 weeks gestational age atlas space [29] using MIRTK [24], and then we resampled both structural and microstructural volumes to be 1 mm^3 isotropic resolution. To obtain the diffusion tensor (DT) maps, we used the DWI2TENSOR MRtrix3 [27] command, and we performed skull-stripping using the available dHCP brain masks [17]. All images were cropped to a $128 \times 128 \times 128$ size.

Data Selection. To train our image registration models, we divided our dataset using an 80%–10%–10% split (see Table 1), using stratified sampling in order to keep the distribution of post-menstrual age (PMAs) at scan and the male-to-female ratio close to the original distribution. We report all of our results on the test set.

Table 1. Number of scans used for training, validation and testing the models.

Dataset	Subjects	GA (weeks)	PMA (weeks)
Train	350 (164♀+ 186♂)	38.0 (±3.8)	40.6 (±1.9)
Validate	34 (14♀ + 20♂)	39.7 (±1.4)	40.7 (±1.7)
Test	30 (12♀ + 18♂)	39.8 (±1.5)	40.6 (±1.9)

2.2 Model Training

Let F, M represent the fixed (target) and the moving (source) magnetic resonance (MR) volumes, respectively, and let φ be the deformation field. In this work, the focus is on scalar-valued structural data (F^{T2w} and M^{T2w}) and tensor-valued microstructural DTI volumes (F^{DTI} and M^{DTI}) which are 6-component data. The moving images (M^{T2w} and M^{DTI}, co-registered) are acquired from the same subject, while the fixed images (F^{T2w} and F^{DTI}, co-registered) are a 36 weeks old neonatal atlas [29].

Scalar-Valued. The **baseline** image registration network [14] is trained with T_2w data only, and aims to minimize the scalar-valued image registration loss $\mathcal{L}^{\mathbf{S}}_{reg} = \mathcal{L}_{\mathbf{KLD}} + \lambda_{\mathbf{BE}}\mathcal{L}_{\mathbf{BE}}(\varphi) + \lambda\mathcal{L}^{\mathbf{S}}_{dsim}$, where $\lambda_{\mathbf{BE}} = 0.01$ and $\lambda = 5000$, as proposed in [14]. The three terms are: the Kullback-Leibler (KL) divergence $\mathcal{L}_{\mathbf{KLD}}$ [9,14] (to reduce the gap between the prior and the encoded distributions), the bending energy (BE) regularisation penalty $\mathcal{L}_{\mathbf{BE}}$ [22] which regularizes the predicted deformation field φ, and the scalar-valued image dissimilarity loss $\mathcal{L}^{\mathbf{S}}_{dsim} = \lambda_{T2w}\mathcal{L}_{\mathbf{NCC}}\left(\mathbf{F}^{T2w}, \mathbf{M}^{T2w}\right)$, where $\mathcal{L}_{\mathbf{NCC}}$ is the symmetric normalised cross correlation (NCC) dissimilarity measure.

Tensor-Valued. For the models which use DTI data, we develop layers capable of dealing with the higher-order microstructural data and introduce them as part of the **baseline** network. Registration of DT images is not as straightforward to perform as scalar-valued data. When transforming the latter, the intensities in the moving image are interpolated at the new locations determined by the deformation field φ and assigned to the corresponding location in the target image space. However, after interpolating DT images, the diffusion tensors need to be reoriented to remain anatomically correct [2]. This is done by applying a rotation matrix R to each resampled diffusion tensor $\mathbf{D}$, such that: $\mathbf{D}' = R\mathbf{D}R^T$. When the transformation is non-linear the reorientation matrix can be computed at each point in the deformation field φ through a polar decomposition of the local Jacobian matrix J. This factorisation transforms the non-singular matrix J into a unitary matrix R (the pure rotation) and a positive-semidefinite Hermitian matrix P, such that $J = RP$ [25]. The rotation matrices R are then used to reorient the tensors without changing the local microstructure. This is known as the *finite strain* strategy [2]. We train this model either single (DTI-only) or multi-channel (T_2w+DTI). In this case, the optimizer aims to minimize the scalar and tensor-valued image registration loss $\mathcal{L}^{\mathbf{D}}_{reg} = \mathcal{L}_{\mathbf{KLD}} + \lambda_{\mathbf{BE}}\,\mathcal{L}_{\mathbf{BE}}(\varphi) + \lambda\,\mathcal{L}^{\mathbf{D}}_{dsim}$, where $\mathcal{L}^{\mathbf{D}}_{dsim} = \lambda_{T2w}\,\mathcal{L}_{\mathbf{NCC}}\left(\mathbf{F}^{T2w}, \mathbf{M}^{T2w}\right) + \lambda_{DTI}\,\mathcal{L}_{\mathbf{EDS}}\left(\mathbf{F}^{DTI}, \mathbf{M}^{DTI}\right)$. Here, $\mathcal{L}_{\mathbf{EDS}}$ is the Euclidean distance squared (EDS) [31].

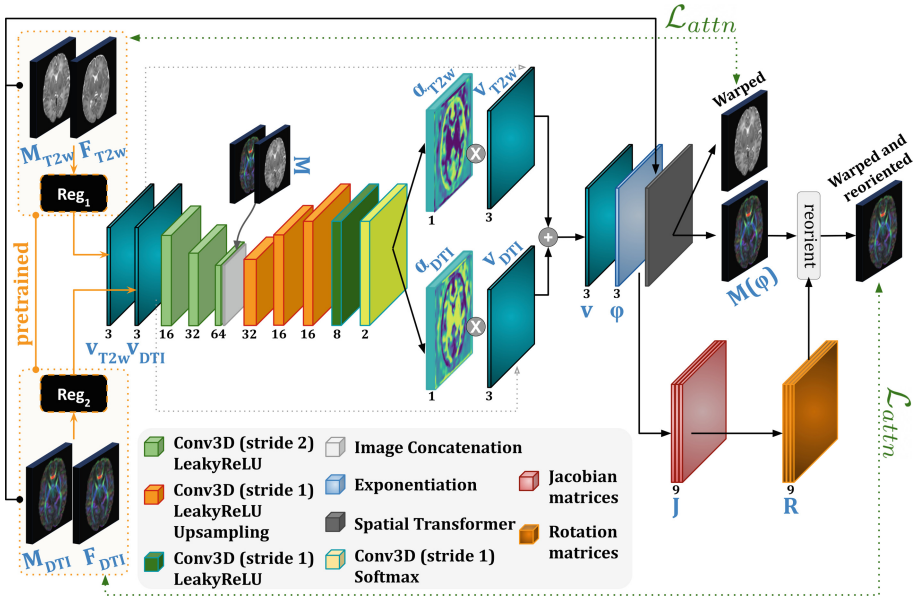

Fig. 2. Network architecture of our proposed **attention**-driven multi-channel model. Note that the same diffusion tensor reorientation layers are added to the **baseline** DTI-only image registration network.

Attention Velocity Network. For the **attention**-driven registration task, we construct a lightweight CNN which uses pairs of modality-specific velocity fields as input, and outputs a combined velocity field which aims to align both structural and microstructural data simultaneously. The proposed network learns the attention maps α_{T2w} and α_{DTI}, for which $\alpha_{T2w} + \alpha_{DTI} = \mathbf{1}$ at every voxel. The input velocity fields are weighted with the attention maps and combined to create a final velocity field $v = v_{T2w} \odot \alpha_{T2w} + v_{DTI} \odot \alpha_{DTI}$, where $\odot$ represents the element-wise multiplication, and $+$ represents the element-wise addition. The final dense deformation field φ is then obtained through *scaling and squaring* layers applied to the velocity field v. The architecture of our proposed attention-based network is presented in Fig. 2. The loss function for training this network is $\mathcal{L}_{attn} = \lambda_{T2w}\,\mathcal{L}_{\mathbf{NCC}}(\mathbf{F}^{T2w}, \mathbf{M}^{T2w}) + \lambda_{DTI}\,\mathcal{L}_{\mathbf{EDS}}(\mathbf{F}^{DTI}, \mathbf{M}^{DTI})$.

Training Details. We train 12 models: 2 single-modality (T_2w-only or DTI-only), and 10 multi-channel models (5 **baseline** and 5 **attention**-driven models) with the following set of hyperparameters $(\lambda_{T2w}, \lambda_{DTI}) = \{(1.0, 0.5), (1.0, 1.0), (1.0, 1.5), (1.0, 2.0), (1.0, 2.5)\}$. We train each model until convergence (using the Adam optimizer with its default parameters), with a maximum of 300 epochs, and use the models which performed best on the validation set. All networks were implemented in PyTorch (v1.10.2), with TorchIO (v0.18.73) [19] for data preprocessing (intensity normalisation) and loading, and training was performed on a 12 GB Titan XP. Average training times were: ∼5 h for the scalar-valued networks (T_2w-only), and ∼100 h for the tensor-valued networks due to their increased complexity (99 h for the DTI-only, 104 h for the **baseline** T_2w+DTI, and 99 h for the **attention** T_2w+DTI).

2.3 Evaluation of Image Registration Models

The quantitative evaluation is carried out on our test dataset of 30 subjects. Each subject and the atlas had the following tissue label segmentations obtained from T_2w images using the Draw-EM pipeline [16]: cortical gray matter (cGM), WM, deep gray matter (dGM) and ventricles. Additionaly, a WM structure called the internal capsule (IC) was manually segmented on the FA maps of all test subjects. These labels were propagated from each subject into the atlas space using the predicted deformation fields. A two-sample, two-sided paired t-test with a significance level of 5% was used to compare pairs of trained models.

Hyperparameter Tuning. To evaluate performance of the registration, we calculate Dice scores between the warped labels and the atlas labels, for all of our models. To choose the best set of hyperparameters for the multi-channel models, we look at the performance of our trained models on two structures: the cGM and the IC. We chose these two structures because the cGM ribbon is a difficult-to-align structure due to the inherent inter-individual variability in cortical folding across the cohort, while the IC is a white matter structure which is prominent in the microstructural data, and barely visible in T_2w structural

images. We also produce average α_{DTI} attention maps obtained from 10 neonatal subjects scanned around 40 weeks PMA.

Alignment of Labels After Registration. After choosing the best overall set of hyperparameters based on the two selected structures, we report Dice scores for the remaining labels.

Evaluation of White Matter Alignment. We also carry out an evaluation of the alignment of WM microstructure, using two metrics proposed in [1]: 1) the cross correlation (CC) between the trace of tensors, which looks at the overall diffusivity, and 2) the overlap of eigenvalue–eigenvector pairs (OVL), which looks at the directional components of the DTIs.

Diffusion Tensor Maps. Finally, we present a qualitative result showcasing warped DTI maps and T_2w images obtained after registering a subject scanned around 40 weeks gestational age (GA) to the atlas.

3 Results

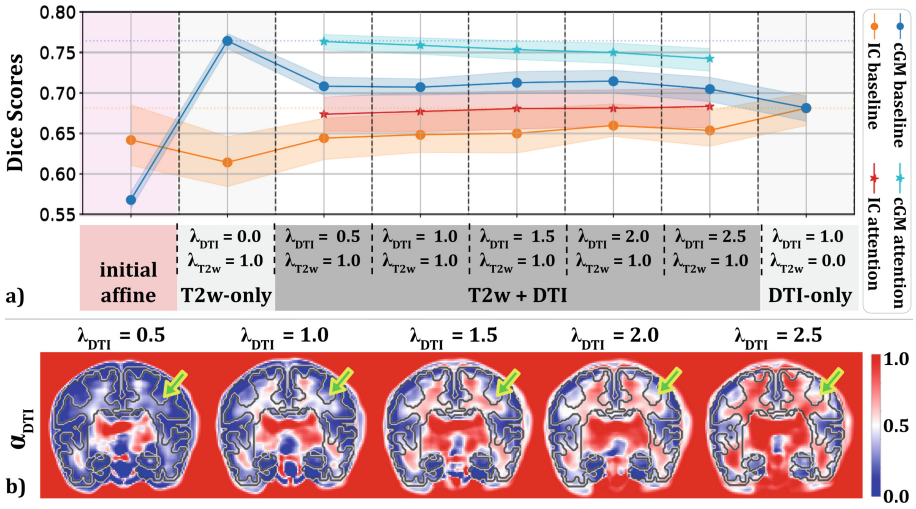

Fig. 3. Hyperparameter tuning results showing: a) mean Dice scores for cGM and IC structures where the shading around each line is the interquartile range (IQR), and b) average α_{DTI} attention maps.

Hyperparameter Tuning. Figure 3a summarizes our results in terms of Dice scores for the 7 **baseline** models (1 T_2w-only, 1 DTI-only and 5 T_2w+DTI models with different λ_{DTI} values), and the 5 proposed T_2w+DTI **attention**-based networks. The T_2w-only model (second column) obtains scores which are worse than the initial affine alignment (first column) for the IC, and best for the cGM label. On the other hand, the DTI-only model (last column) obtains the best overall IC alignment, and the worst for the cGM. The **attention** networks always outperform ($p - value < 0.05$) the **baseline** multi-channel models, for all values of λ_{DTI} (columns 3–7), for both the cGM and IC tissue types. Figure 3b shows the midbrain coronal slices of average α_{DTI} attention maps for increasing values of λ_{DTI}. It is evident from the figure that for $\lambda_{DTI} = 0.5$ the network uses the DTI channel mainly around the dGM region of the brain, while at the same time, for $\lambda_{DTI} > 1.0$, the α_{DTI} maps seem to *spill over* onto the cortical ribbon (as highlighted by the green arrows), *i.e.*, the models use some information from the DTI channel when aligning the cGM. We found the middle ground to be at $(\lambda_{T2w}, \lambda_{DTI}) = (1.0, 1.0)$, in order to not lose performance on cGM alignment, while also produced good results in terms of the IC label. All of our following results are based on this set of hyperparameters.

Table 2. Mean (±standard deviation) Dice scores of the single- and multi-channel ($\lambda_{DTI} = 1.0$) experiments, together with the initial affine alignment. Statistically significant differences ($p - value < 0.05$) are reported in terms of best/worst overall score (highlighted in bold/shaded in gray, respectively), and best amongst the multi-channel models (highlighted with a black box). ***A*** - *"attention"*.

Model	cGM	WM	Ventricles	dGM	IC
affine	.57±.02	.70±.03	.63±.05	.90±.02	.64±.07
T2w-only	**.76±.01**	**.84±.02**	.80±.02	**.93±.01**	.61±.04
DTI-only	.68±.02	.77±.02	.76±.03	.91±.01	**.68±.03**
T2w+DTI	.71±.01	.75±.03	.76±.03	.90±.01	.65±.03
A T2w+DTI	.76±.01	.83±.02	**.80±.02**	.93±.01	**.68±.03**

Alignment of Labels After Registration. Table 2 shows the Dice scores obtained by our trained models for all tissue types (cGM, WM, ventricles, dGM and IC). The proposed **attention**-based network systematically outperforms the **baseline** T_2w+DTI models for all structures, with the latter obtaining significantly worse scores ($p - value < 0.05$) than its **attention**-driven counterpart. These results also highlight the importance of adding the T_2w structural channel to DTI registration in order to improve structural alignment; specifically, the DTI-only model drops in performance when compared to the T_2w-only model in terms of aligning the cGM, WM, ventricles and dGM tissue types, but regains its performance when using both **attention** and the extra structural channel.

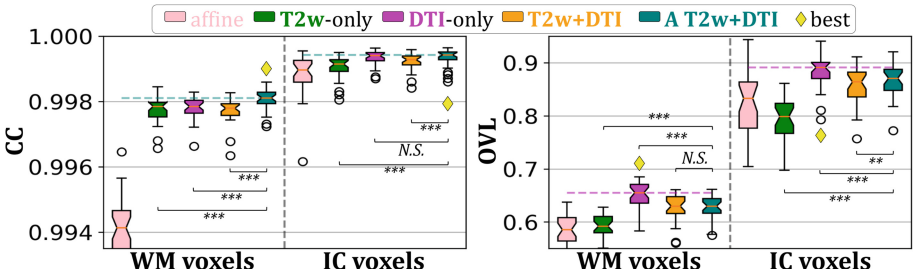

Fig. 4. Test dataset CC scores (left) and average OVL values (right) calculated between the warped DT images with reoriented tensors and the fixed DTI atlas. Best performing models ($p - value < 0.05$) are marked with a yellow diamond. $***$-p-$value < 0.001$, $**$-p-$value < 0.01$, $*-p$-$value < 0.05$, $N.S.$-"$not\ significant$". (Color figure online)

Evaluation of White Matter Alignment. Figure 4 summarizes our results in terms of CC scores (left) and average OVL values (right), either computed across the WM voxels (using the atlas WM label map), or within the IC voxels (using the atlas IC label map). In terms of the CC scores, our proposed **attention** T_2w+DTI model outperforms all other models ($p - value < 0.05$), followed closely by the DTI-only model which performs similarly well for the IC voxels. For the average OVL values, the best overall scores are obtained by the DTI-only model, followed by our proposed **attention** model, which outperforms its non-attention counterpart ($p - value < 0.05$) for the IC voxels. Both **baseline** and **attention** T_2w+DTI models perform similarly well for the WM voxels. On the other hand, the T_2w-only model consistently underperforms in these two metrics, showing that the DTI channel is essential for correct alignment of WM microstructure.

Diffusion Tensor Maps. Figure 5 shows an example of warped DTI maps (rows 1–2) and T_2w images (row 3) obtained after registering a subject scanned around 40 weeks GA to the atlas. The IC is strikingly less coherent in the T_2w-only model (yellow arrows, first row) when compared to the atlas, while the overall shape of the corpus callosum (white arrows, second row) is less preserved. This is also the case for the **baseline** T_2w+DTI model where the IC seems to be thinner than obtained by the **attention** T_2w+DTI and DTI-only models (yellow arrows, first row), while the corpus callosum exhibits an uneven shape (white arrows, second row). The cortical ribbon (last row) is not well aligned for the DTI-only model, it is generally "*overspilling*" into the WM for the **baseline** T_2w+DTI, while being similarly well aligned with the fixed (blue outline) cortex in the T_2w-only and our proposed **attention** T_2w+DTI models.

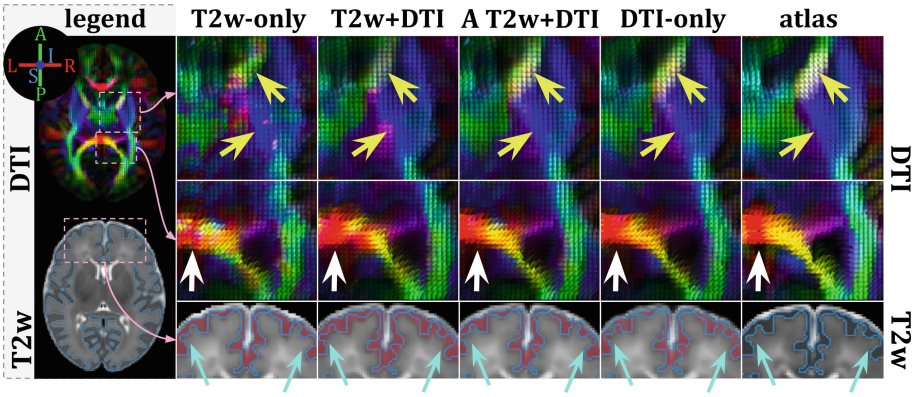

Fig. 5. Mid-brain axial slices of the warped DTI and T_2w maps of a 40 weeks GA neonate for 4 of our trained models: T_2w-only, **baseline** T_2w+DTI ($\lambda_{DTI} = 1.0$), **attention** T_2w+DTI ($\lambda_{DTI} = 1.0$) and DTI-only, together with the fixed DTI atlas. Yellow arrows (first row) point to the IC, while the white arrows (second row) point to the corpus callosum. The warped T_2w images (last row) are shown together with their corresponding highlighted cortical ribbon (in red), as well as an outline (in blue) of the fixed (target) cortex. Cyan arrows (last row) point to mis-registered regions of the cortical ribbon. (Color figure online)

4 Conclusion

In this work we presented a novel solution for multi-channel registration, which combines structural and microstructural tensor-valued MRI data based on learned spatially varying attention maps that optimise the multi-channel alignment. Our evaluation showed that the proposed **attention**-driven image registration network improves overall alignment when compared to models trained on multi-channel data without attention, while maintaining the performance of the single-channel registration for the structures salient on that channel. Key limitations of our work include the reliance on accurately co-registered structural and microstructural data, and the DTI model's inability to resolve crossing fibers [28,30]. This has the potential to be explored in future work.

Acknowledgements. This work was supported by the Academy of Medical Sciences Springboard Award [SBF004\1040], Medical Research Council (Grant no. [MR/K006355/1]), European Research Council under the European Union's Seventh Framework Programme [FP7/20072013]/ERC grant agreement no. 319456 dHCP project, the EPSRC Research Council as part of the EPSRC DTP (grant Ref: [EP/R513064/1]), the Wellcome/EPSRC Centre for Medical Engineering at King's College London [WT 203148/Z/16/Z], the NIHR Clinical Research Facility (CRF) at Guy's and St Thomas', and by the National Institute for Health Research Biomedical Research Centre based at Guy's and St Thomas' NHS Foundation Trust and King's College London.

References

1. Adluru, N., Zhang, H., Fox, A.S., Shelton, S.E., Ennis, C.M., Bartosic, A.M., Oler, J.A., Tromp, D.P., Zakszewski, E., Gee, J.C., et al.: A diffusion tensor brain template for rhesus macaques. Neuroimage **59**(1), 306–318 (2012)
2. Alexander, D.C., Pierpaoli, C., Basser, P.J., Gee, J.C.: Spatial transformations of diffusion tensor magnetic resonance images. IEEE transactions on medical imaging **20**(11), 1131–1139 (2001)
3. Avants, B., Duda, J.T., Zhang, H., Gee, J.C.: Multivariate normalization with symmetric diffeomorphisms for multivariate studies. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 359–366. Springer (2007)
4. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Medical image analysis **12**(1), 26–41 (2008)
5. Chen, M., Carass, A., Jog, A., Lee, J., Roy, S., Prince, J.L.: Cross contrast multichannel image registration using image synthesis for mr brain images. Medical image analysis **36**, 2–14 (2017)
6. Chen, M., Jog, A., Carass, A., Prince, J.L.: Using image synthesis for multi-channel registration of different image modalities. In: Medical Imaging 2015: Image Processing. vol. 9413, pp. 462–468. SPIE (2015)
7. Daga, P., Winston, G., Modat, M., Cardoso, M.J., Stretton, J., Symms, M., McEvoy, A.W., Hawkes, D., Duncan, J., Ourselin, S.: Integrating structural and diffusion mr information for optic radiation localisation in focal epilepsy patients. In: 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. pp. 353–356. IEEE (2011)
8. Dubois, J., Alison, M., Counsell, S.J., Hertz-Pannier, L., Hüppi, P.S., Benders, M.J.: Mri of the neonatal brain: a review of methodological challenges and neuroscientific advances. Journal of Magnetic Resonance Imaging **53**(5), 1318–1343 (2021)
9. Duchi, J.: Derivations for linear algebra and optimization. Berkeley, California **3**(1), 2325–5870 (2007)
10. Edwards, A.D., Rueckert, D., Smith, S.M., Seada, S.A., Alansary, A., Almalbis, J., Allsop, J., Andersson, J., Arichi, T., Arulkumaran, S., et al.: The developing human connectome project neonatal data release. Frontiers in Neuroscience **16** (2022)
11. Forsberg, D., Rathi, Y., Bouix, S., Wassermann, D., Knutsson, H., Westin, C.F.: Improving registration using multi-channel diffeomorphic demons combined with certainty maps. In: International Workshop on Multimodal Brain Image Analysis. pp. 19–26. Springer (2011)
12. Guimond, A., Guttmann, C.R., Warfield, S.K., Westin, C.F.: Deformable registration of dt-mri data based on transformation invariant tensor characteristics. In: Proceedings IEEE International Symposium on Biomedical Imaging. pp. 761–764. IEEE (2002)
13. Jog, A., Carass, A., Pham, D.L., Prince, J.L.: Random forest flair reconstruction from t 1, t 2, and p d-weighted mri. In: 2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI). pp. 1079–1082. IEEE (2014)
14. Krebs, J., Mansi, T., Mailhé, B., Ayache, N., Delingette, H.: Unsupervised probabilistic deformation modeling for robust diffeomorphic registration (2018)

15. Liu, X., Jiang, D., Wang, M., Song, Z.: Image synthesis-based multi-modal image registration framework by using deep fully convolutional networks. Medical & Biological Engineering & Computing **57**, 1037–1048 (2019)
16. Makropoulos, A., Gousias, I.S., Ledig, C., Aljabar, P., Serag, A., Hajnal, J.V., Edwards, A.D., Counsell, S.J., Rueckert, D.: Automatic whole brain MRI segmentation of the developing neonatal brain. IEEE transactions on medical imaging (2014)
17. Makropoulos, A., Robinson, E.C., Schuh, A., Wright, R., Fitzgibbon, S., Bozek, J., Counsell, S.J., Steinweg, J., Vecchiato, K., Passerat-Palmbach, J., et al.: The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. Neuroimage (2018)
18. Park, H.J., Kubicki, M., Shenton, M.E., Guimond, A., McCarley, R.W., Maier, S.E., Kikinis, R., Jolesz, F.A., Westin, C.F.: Spatial normalization of diffusion tensor mri using multiple channels. Neuroimage **20**(4), 1995–2009 (2003)
19. Pérez-García, F., Sparks, R., Ourselin, S.: TorchIO: a Python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning. arXiv:2003.04696 [cs, eess, stat] (Mar 2020), http://arxiv.org/abs/2003.04696, arXiv: 2003.04696
20. Peyrat, J.M., Delingette, H., Sermesant, M., Xu, C., Ayache, N.: Registration of 4d cardiac ct sequences under trajectory constraints with multichannel diffeomorphic demons. IEEE transactions on medical imaging **29**(7), 1351–1368 (2010)
21. Rohde, G.K., Pajevic, S., Pierpaoli, C., Basser, P.J.: A comprehensive approach for multi-channel image registration. In: International Workshop on Biomedical Image Registration. pp. 214–223. Springer (2003)
22. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: application to breast MR images. IEEE Transactions on Medical Imaging (1999)
23. Rutherford, M.: MRI of the neonatal brain. Magnetic resonance imaging of the brain in preterm infants: 24 weeks' gestation to term pp. 25–49 (2002)
24. Schuh, A., Makropoulos, A., Robinson, E.C., Cordero-Grande, L., Hughes, E., Hutter, J., Price, A.N., Murgasova, M., Teixeira, R.P.A., Tusor, N., et al.: Unbiased construction of a temporally consistent morphological atlas of neonatal brain development. BioRxiv p. 251512 (2018)
25. Shoemake, K., Duff, T.: Matrix animation and polar decomposition. In: Graphics Interface. vol. 92, pp. 258–264. Citeseer (1992)
26. Thirion, J.P.: Image matching as a diffusion process: an analogy with maxwell's demons. Medical image analysis **2**(3), 243–260 (1998)
27. Tournier, J.D., Smith, R., Raffelt, D., Tabbara, R., Dhollander, T., Pietsch, M., Christiaens, D., Jeurissen, B., Yeh, C.H., Connelly, A.: MRtrix3: A fast, flexible and open software framework for medical image processing and visualisation. NeuroImage **202**, 116137 (2019)
28. Tuch, D.S., Reese, T.G., Wiegell, M.R., Makris, N., Belliveau, J.W., Wedeen, V.J.: High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine **48**(4), 577–582 (2002)
29. Uus, A., Grigorescu, I., Pietsch, M., Batalle, D., Christiaens, D., Hughes, E., Hutter, J., Cordero Grande, L., Price, A.N., Tournier, J.D., Rutherford, M.A., Counsell, S.J., Hajnal, J.V., Edwards, A.D., Deprez, M.: Multi-channel 4D parametrized atlas of macro- and microstructural neonatal brain development. Frontiers in Neuroscience **15**, 721 (2021)

30. Wiegell, M.R., Larsson, H.B., Wedeen, V.J.: Fiber crossing in human brain depicted with diffusion tensor MR imaging. Radiology **217**(3), 897–903 (2000)
31. Zhang, H., Yushkevich, P.A., Alexander, D.C., Gee, J.C.: Deformable registration of diffusion tensor MR images with explicit orientation optimization. Medical image analysis **10**(5), 764–785 (2006)