



A Deep Generative Model of Neonatal Cortical Surface Development

Abdulah Fawaz¹(✉), Logan Z. J. Williams^{1,2}, A. David Edwards^{2,3,4},
and Emma C. Robinson^{1,2}

¹ Department of Biomedical Engineering, School of Biomedical Engineering and Imaging Sciences, King's College London, London, UK

abdulah.fawaz@kcl.ac.uk

² Centre for the Developing Brain, School of Biomedical Engineering and Imaging Sciences, King's College London, London, UK

³ Department for Forensic and Neurodevelopmental Sciences, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London SE5 8AF, UK

⁴ MRC Centre for Neurodevelopmental Disorders, King's College London, London SE1 1UL, UK

Abstract. The neonatal cortical surface is known to be affected by preterm birth, and the subsequent changes to cortical organisation have been associated with poorer neurodevelopmental outcomes. Deep Generative models have the potential to lead to clinically interpretable models of disease, but developing these on the cortical surface is challenging since established techniques for learning convolutional filters are inappropriate on non-flat topologies. To close this gap, we implement a surface-based CycleGAN using mixture model CNNs (MoNet) to translate sphericalised neonatal cortical surface features (curvature and T1w/T2w cortical myelin) between different stages of cortical maturity. Results show our method is able to reliably predict changes in individual patterns of cortical organisation at later stages of gestation, validated by comparison to longitudinal data; and translate appearance between preterm and term gestation (>37 weeks gestation), validated through comparison with a trained term/preterm classifier. Simulated differences in cortical maturation are consistent with observations in the literature.

Keywords: Geometric deep learning · Cortical surfaces · Neurodevelopment

1 Introduction

Deep Generative modelling presents enormous opportunities for medical imaging analysis: from image segmentation [8, 9, 13, 40], registration [39], motion correction and denoising [1, 23, 30, 42], to anomaly detection [4, 14, 45] and the development of clinically interpretable models of disease progression [2, 3, 25]. Image-to-image translation is a type of generative modelling problem where images

are transformed across domains, in a way that preserves their content. Common image-to-image translation tasks include transforming between imaging modalities [15, 36, 41] and to increase image resolution [11, 17, 44], though the methods used are general to any sets of imaging domains or classes, and can be applied to both 2D and 3D imaging data (see [43] for a full review of GANs in medical imaging).

The application of deep learning on surfaces, including medical surfaces, has been hindered by the mathematical incompatibility of data structures with non-flat topologies to conventional approaches for convolutional filtering. In response, the field of geometric deep learning (gDL) has emerged to extend deep learning to domains such as graphs, surfaces and meshes [7]. gDL is yet to produce a singular approach and numerous, often competing, tools and methods have been developed.

In this paper, we aim to develop surface-to-surface translation models of the developing cortex sensitive to individual changes in cortical maturation. Such difference maps could act as imaging biomarkers of interest to clinicians, since the negative impact of preterm birth on neurodevelopmental outcome remains an area of active research [5, 29]. This task is too challenging for traditional imaging analysis methods due to the significant variation in functional and structural organisation of human cortices, even amongst healthy populations. The heterogeneity is present in neonatal cortical surface patterns from an early stage, which then undergo rapid and complex cortical maturation during development, which are further obfuscated by the impact of preterm birth. These factors make this task a fitting domain for the use of gDL, and gDL models have already been shown to achieve state of the art performance for neurodevelopmental phenotype regression and on cortical segmentation, when benchmarked against classical methods [10, 12, 35, 37, 46]. To achieve this, we adapt the CycleGAN framework [47] with methods from gDL for use on surfaces.

2 Methods

The CycleGAN framework [47] is a well-established method for unsupervised bidirectional image-to-image translation across two domains. This is achieved with a pair of generators that learn mappings from a source domain to a target domain, trained adversarially to a corresponding pair of discriminators that differentiate between real and synthetic images. Traditionally, these generators and discriminators are Euclidean CNNs but in this work, we retain the architecture but replace the convolutions with a surface-compatible convolution operation from the MoNet model [28]. CycleGANs produce realistic images and do not require paired examples, which can be practically difficult to acquire in medical imaging. These fit in with our overall aim of generating realistic, subject-specific cross-domain difference maps that can identify individual deviations from normal morphological development.

2.1 Model Architecture

The following adaptations were made to the original CycleGAN [47] to facilitate application on surfaces:

Geometric Convolutions: All surface convolutions are implemented using MoNet Gaussian Mixture Model convolutions (GMMConv) [28]:

$$(f \star g)(x) = \sum_{j=1}^J g_j D_j(x) f \quad (1)$$

$$D_j(x) f = \sum_{y \in \mathcal{N}(x)} w_j(\mathbf{u}(x, y)) f(y), \quad j = 1, \dots, J$$

where f, g are the input and filter respectively, $D(x)f$ is a parameterisable patch operator that determines how the data is extracted from the surface before being weighted by the filter. Its general form is given in equation (1), with x a point on the surface, y a point in the defined neighbourhood of x , and w a kernel applied to the pseudo-coordinates $u(x, y)$ that define how the neighbourhood of an individual point is weighted relative to its neighbours on the surface. In MoNet, the weighting function w takes a Gaussian form:

$$w_j(\mathbf{u}) = \exp\left(-\frac{1}{2}(\mathbf{u} - \mu_j)^T \Sigma_j^{-1}(\mathbf{u} - \mu_j)\right) \quad (2)$$

where Σ_j and μ_j are learnable $d \times d$ and $d \times 1$ covariance matrix and mean vector of a Gaussian kernel, respectively. The covariances are further restricted to diagonal form. In this paper, Gaussian kernels are defined using a two dimensional ($J=2$) local pseudo-coordinate system defined as the vector difference between neighbouring points on the surface in polar coordinates. This is pre-computed.

Icospheric Pooling: Since GMMConv has no transposed convolution, and does not reduce the dimensionality of the data, we require pooling and unpooling layers adapted for the surface. This can be challenging for general surfaces, but icospheres of different resolutions may be generated from each other by iterative barycentric interpolation (Fig. 1). This allows us to define a native down/upsampling method on the icosphere; the former as a direct downsample of the data by including only points of the lower icosphere (e.g. the red points in Fig. 1), and the latter as a mean-unpool operation where the new points are the average of their direct neighbours during the upsampling process. Skip connections are also added to make the generator a U-net.

Training: The network is optimised following the protocol of a standard CycleGAN [47], which learns mappings $G_X : X \rightarrow Y$ and $G_Y : Y \rightarrow X$, through adversarial training. Unpaired images $\{x, y\}$ are input through each of the Generators, G, to produce synthetic images $\{\hat{x}_X, \hat{y}_X\}$ and $\{\hat{x}_Y, \hat{y}_Y\}$ from generators




Fig. 1. Icosahedrons can be efficiently up and downsampled to different resolutions. New points are generated by barycentric interpolation of existing points. Our down-sampling process simply keeps the points of the lower resolution icosphere while discarding the rest (e.g. to downsample from 4V to 2V, keep the points shown in red). Our up-sampling procedure mimics the up-sampling of the icospheres but with the new points added as an average of the existing points within its neighbourhood. (Color figure online)

G_X and G_Y respectively. The identity loss penalises each generator for changing an image already in its domain i.e. it forces $\hat{x}_X \approx x$, and $\hat{y}_Y \approx y$. The discriminators D_x and D_Y attempt to classify the true from synthetic images of their respective classes. Finally, the synthetic output of one generator is fed into the other $X \rightarrow Y \rightarrow X$ in order to generate a recovered image. The cycle loss penalises each generator for failing to reproduce the original image during this cycle.

We use an L1 loss for the cycle and identity losses, and an MSE loss for the adversarial loss. All networks are trained with an Adam optimizer with learning rate = 0.002, betas = [0.9, 0.99]. We set the real and false labels to 0.9 and 0.1 to help avoid discriminator overfitting. Networks were trained for 250 epochs.

The data is augmented with non-linear warps to represent realistic variations. These were generated offline by creating surface meshes warped by very small random displacements in the vertices, and resampling the original data onto these warped surface meshes by barycentric interpolation [12]. 100 warped counterparts are created for each image in our dataset. Each imaging modality was normalised to between 0 and 1 across the dataset.

3 Experimental Methods and Results

3.1 Data

All of the data used are part of the third release of the publically available Developing Human Connection Project (dHCP) dataset, consisting of cortical surface meshes and metrics, derived from T1-weighted and T2-weighted MRI images, processed and registered as described in [27] and references therein [16, 20, 21, 33]. We utilise left cortical hemispheres scanned between 25 and 45 weeks post-menstrual age (PMA) with myelination and cortical curvature

metrics initially registered to a 32k-vertex sphere using Multimodal Surface Matching (MSM) [31,32] and upsampled to a regular icosphere. The data set contains 419 neonates born at term (≥ 37 weeks gestational age (GA)) and 161 born preterm (< 37 weeks gestational age), where 45 preterm subjects were scanned twice, once at birth and once at term-equivalent age (TEA), giving a total of 625 scans.

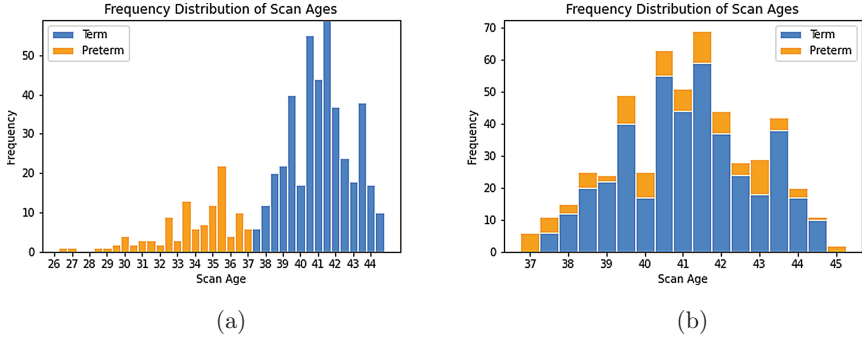


Fig. 2. A bar chart showing the distribution of subjects’ PMA at Scan in weeks for (a) Experiment 1 and (b) Experiment 2.

3.2 Experiment 1: Cortical Maturity

In this experiment, the objective was to simulate healthy cortical maturation; thus the two domains to be translated between were: A) preterm infants’ first scans (PMA < 37 weeks); and B) healthy term controls. We exclude the later scans of preterm subjects to avoid modelling atypical cortical maturation, giving a total of 530 subjects: 419 term/111 preterm (as shown in Fig. 2a).

A spatio-temporal atlas of cortical maturation maps is available for the dHCP data [6], and shows a general increase in myelination around the brain, but especially strongly along the central sulcus (somatosensory cortex), the posterior portion of the lateral fissure (auditory cortex), and at higher scan ages, around the calcarine sulcus (visual cortex). By 28 weeks, all the primary sulci are formed and there are already heterogeneities in cortical structure. These sulci are initially smooth but become significantly more folded during development, and the cortical surface further increases in complexity by the emergence of tertiary sulci. Our model must be able to preserve and extend the existing folding patterns of an immature brain in its predictions.


The presence of longitudinal data in the form of preterm subjects’ second scans allows us to quantitatively evaluate our model’s predictions by comparing the actual appearance of a preterm subject’s later scan to the model’s matured prediction based on the same subjects first scan. We analyse these both quantitatively, via the image similarity metric Peak Signal to Noise Ratio (PSNR), and qualitatively by visual comparison. We compare the images as a whole and

by individual modality (myelination/curvature). If the model represents cortical maturation, the model's synthetic predictions would be expected to be more similar to the later second scans than the original first scans. However, we would expect some differences due to variation around the age at scan and the impact of preterm birth on cortical development.


Figures 3a and b show the predicted surface to surface translation for low-to-high PMA of two different subjects' myelination and curvature respectively, displaying only the lateral view for clarity. The first, second and third columns show, in order, the original first scans, the true second scans and the synthetic scans produced by the model. The final column is a difference map between the model input and the model output. Each row shows a different subject and it can be seen that the model correctly predicts the broad structural changes expected in both myelination and curvature for both subjects. The relative homogeneity of myelination makes it easy to see that the model has correctly learned to simulate subject development.

Changes in curvature are more subtle and display greater heterogeneity across subjects and so are the better benchmark of performance. The preservation of cortical folding patterns is a key factor required for proper modelling of curvature and we see that our model does this well for both subjects despite significant topological differences between the two. Further, we observe that the model has added significant branching to most sulci, particularly the superior temporal sulcus which increases greatly in folding. The model has created new smaller folds in the temporal lobe, the parietal lobe and the frontal lobe and, extended existing sulci throughout. In the second row, the model even correctly predicts a change in the topology of a sulcus in the frontal lobe, which is split by a gyrus as it develops (shown by the red circle).


Figure 4a is a scatterplot showing the quantified similarity between the second scans, which acts as a fuzzy 'ground truth', and the model's synthetic aged scans as compared to the first original scan, plotted against PMA at second scan. The results show that, in almost every case, the model predictions demonstrate a significant increase in similarity to the second scan compared to the original. The exceptions are mostly confined to subjects with second scans at low PMA. The results are summarised in the boxplots in Fig. 4b, where the synthetic images are of greater similarity than the originals over both modalities individually and combined. We observe that, whilst synthetically generated curvature is consistently improved, there is a greater variation in myelination with a tail of synthetic images with low image similarity to the second scans. Visual inspection reveals that this discrepancy is responsible for the reduction in image similarity of synthetic images for low PMA second scans due to the model's predictions being significantly more mature than both the first and second scans in these cases, as seen in Fig. 3c where the predicted myelination is significantly greater than found in the second scan, but still consistent with expected development. This discrepancy is a consequence of the discontinuous nature of cycleGAN translations.



(a)



(b)



(c)

Fig. 3. Low to high PMA translation for the (a) myelination and (b) curvature modalities of two different subjects. The synthetic aged model predictions (column 3) are closer to the later second scans (column 2) than the original first scans (column 1). The fourth column gives the difference map between model input and model output. The red circles indicate a region of folding where the model correctly predicted a change in topology as the sulcus was split into two during development. (c) Myelination prediction for a subject with low PMA (39 weeks) at second scan showing the relative immaturity of the second scan myelination compared to the model prediction, resulting in a higher image similarity to the original first scan than to the model prediction. (Color figure online)

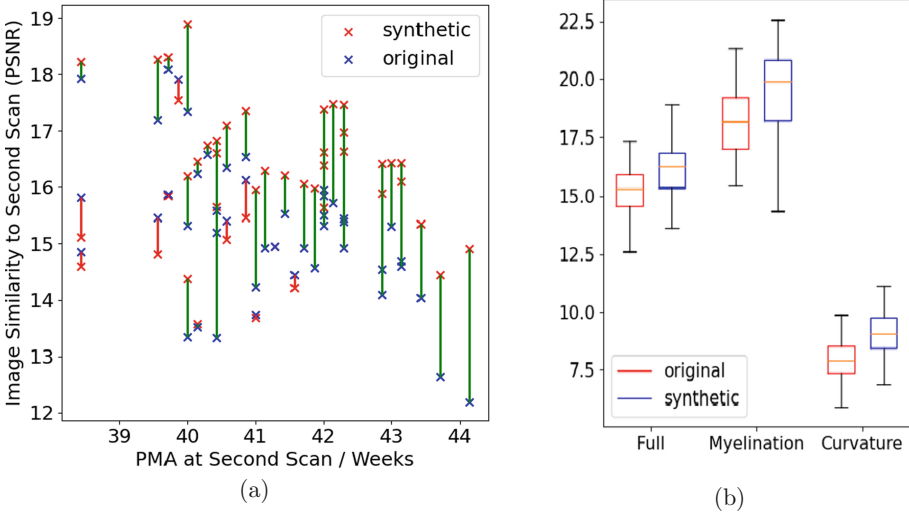


Fig. 4. (a) Image similarity of original first scans vs synthetic aged images compared to the same subject’s second scans, given by Peak Signal-to-Noise Ratio (PSNR). PSNR is plotted against PMA at second scan in weeks. Green lines indicate that the synthetic images are closer to the second scan than original first scans, and red lines indicate a reduction in image similarity. (b) Box plots comparing the range of image similarities as measured by PSNR of original first scans and synthetic aged scans to second scans, split by the full scan (both modalities), then by each modality individually. (Color figure online)

3.3 Experiment 2: Prematurity

In this experiment, the aim is to translate between preterm’s second scan and healthy term controls, in order to predict the impact of preterm birth on each individual’s cortical maturation. We exclude preterm first scans (PMA at scan <37 weeks) to only model the effect of preterm birth at term equivalent ages (TEA), leaving a total of 514 neonatal subjects (419 term/95 preterm). A plot of the cohort used is shown in Fig. 2b.

In this case, there can be no ground truth examples scanned at multiple birth ages, so to validate we train a classifier with same architecture as the discriminator to predict prematurity and apply it to our model’s predictions. Our classifier obtained an accuracy of 94% on the raw dataset, but our cycleGAN model’s synthetic images had an overall 78% success rate in fooling the classifier, and in the other 22% of cases reduced the classifiers confidence by a significant margin.

Figure 5 shows a subject (true GA at birth 40 weeks, PMA at scan 44 weeks) (column 1) that has been translated from term age to a synthetic preterm age (column 2). The image fools our classifier which predicts prematurity with $p = 0.96$ for the synthetic preterm and $p = 0.22$ for the original term subject. The difference map (column 3) shows that myelination is decreased overall but remains broadly unchanged in structure. The folding patterns are completely unchanged but there is a noticeable increase in overall curvature, and it can be seen from the difference map that increases in curvature are more pronounced along the gyri (folds), not the sulci. Comparing these observations to existing work, we note that increased overall cortical curvature has previously been associated with preterm birth when compared to age-equivalent fetuses [24], and when compared with normal term infants [26], with curvature especially high across a number of gyri. Increased curvature has also been implicated as a prognostic biomarker of adverse neurodevelopment [19]. Shimony et al. [34] also found increased overall curvature in preterm neonates compared to term neonates, with increased localised curvature around the gyri, but could not positively determine that the metric was predictive of preterm birth due to the differences in folding patterns between subjects obfuscating the measure - an issue circumvented here as we compare like-for-like subjects through our cycleGAN.




Fig. 5. Original term (columns 1–2) to synthetic premature (columns 3–4) image translation with difference map (columns 5–6). Subject’s true GA at birth is 40 weeks and was scanned at 44 weeks PMA. Classifier predicts $P_{\text{prem}} = 0.22$ for the original term image and $P_{\text{prem}} = 0.96$ for the synthetic premature image.

4 Conclusion

By integrating MoNet into a CycleGAN model, we developed a generative surface-to-surface translation model of cortical maturation that shows accurate, realistic, subject-specific predictions of future myelination and curvature, validated on longitudinal ground truth data. We also developed a model that transformed surfaces between term/preterm classes, producing outputs that fooled a pretrained classifier and showed structural differences in line with what has been

observed in the literature. The primary issue with our models was that the CycleGAN architecture is limited to transformation between discrete classes when this is fundamentally a problem of continuous interpolation. There are a number of different generative models that do allow continuous representation that would be able to interpolate smoothly between scan ages, with which gDL could be adapted. One class of methods utilise variational encodings such as variational autoencoders (VAEs) [18], and VAE-GANs [22], although the authors found that attempts to adapt these methods to surface domains were unsuccessful, with models unable to capture individual cortical folding variation and collapsing to group averages. A more powerful example of this applied to volumetric cortical data is the iCAM architecture [3], which encodes separate variational disentangled spaces for content and age, which may be more amenable to adaptation with gDL. There are further alternatives that utilise direct conditioning on latent variables such as conditional VAEs and conditional GANs, but these too are more commonly associated with conditioning on classes, not continuous variables. Again more complex variations of these models have been successfully applied to the volumetric brain [38] that may be successfully adapted to the surface.

References

1. Armanious, K., Gatidis, S., Nikolaou, K., Yang, B., Kustner, T.: Retrospective correction of rigid and non-rigid MR motion artifacts using GANs. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), pp. 1550–1554. IEEE (2019)
2. Bass, C., da Silva, M., Sudre, C., Tudosiu, P.D., Smith, S., Robinson, E.: ICAM: interpretable classification via disentangled representations and feature attribution mapping. arXiv preprint [arXiv:2006.08287](https://arxiv.org/abs/2006.08287) (2020)
3. Bass, C., et al.: ICAM-REG: interpretable classification and regression with feature attribution for mapping neurological phenotypes in individual scans. arXiv preprint [arXiv:2103.02561](https://arxiv.org/abs/2103.02561) (2021)
4. Benson, S., Beets-Tan, R.: GAN-based anomaly detection in multi-modal MRI images. bioRxiv (2020)
5. Boardman, J.P., Counsell, S.J.: Invited review: factors associated with atypical brain development in preterm infants: insights from magnetic resonance imaging. *Neuropathol. Appl. Neurobiol.* **46**(5), 413–421 (2020)
6. Bozek, J., et al.: Construction of a neonatal cortical surface atlas using multimodal surface matching in the developing human connectome project. *Neuroimage* **179**, 11–29 (2018)
7. Bronstein, M.M., Bruna, J., LeCun, Y., Szlam, A., Vandergheynst, P.: Geometric deep learning: going beyond Euclidean data. *IEEE Sig. Process. Mag.* **34**(4), 18–42 (2017). <https://doi.org/10.1109/MSP.2017.2693418>
8. Cirillo, M.D., Abramian, D., Eklund, A.: Vox2Vox: 3D-GAN for brain tumour segmentation. In: Crimi, A., Bakas, S. (eds.) *BrainLes 2020*. LNCS, vol. 12658, pp. 274–284. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-72084-1_25
9. Costa, P., et al.: End-to-end adversarial retinal image synthesis. *IEEE Trans. Med. Imaging* **37**(3), 781–791 (2017)

10. Dahan, S., Williams, L.Z.J., Rueckert, D., Robinson, E.C.: Improving phenotype prediction using long-range spatio-temporal dynamics of functional connectivity (2021)
11. Do, H., Helbert, D., Bourdon, P., Naudin, M., Guillevin, C., Guillevin, R.: MRI super-resolution using 3D cycle-consistent generative adversarial network. In: 2021 Sixth International Conference on Advances in Biomedical Engineering (ICABME), pp. 85–88. IEEE (2021)
12. Fawaz, A., et al.: Benchmarking geometric deep learning for cortical segmentation and neurodevelopmental phenotype prediction. *bioRxiv* (2021). <https://doi.org/10.1101/2021.12.01.470730>, <https://www.biorxiv.org/content/early/2021/12/02/2021.12.01.470730>
13. Gadermayr, M., et al.: Image-to-image translation for simplified MRI muscle segmentation. *Front. Radiol.* **1** (2021). <https://doi.org/10.3389/fradi.2021.664444>, <https://www.frontiersin.org/article/10.3389/fradi.2021.664444>
14. Han, C., et al.: MadGAN: unsupervised medical anomaly detection GAN using multiple adjacent brain MRI slice reconstruction. *BMC Bioinformatics* **22**(2), 1–20 (2021)
15. Hiasa, Y., et al.: Cross-modality image synthesis from unpaired data using CycleGAN. In: Gooya, A., Goksel, O., Oguz, I., Burgos, N. (eds.) SASHIMI 2018. LNCS, vol. 11037, pp. 31–41. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00536-8_4
16. Hughes, E.J., et al.: A dedicated neonatal brain imaging system. *Magn. Reson. Med.* **78**(2), 794–804 (2017)
17. Jiang, M., et al.: Fa-GAN: fused attentive generative adversarial networks for MRI image super-resolution. *Comput. Med. Imaging Graph.* **92**, 101969 (2021)
18. Kingma, D.P., Welling, M.: Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114* (2013)
19. Kline, J.E., Illapani, V.S.P., He, L., Altaye, M., Logan, J.W., Parikh, N.A.: Early cortical maturation predicts neurodevelopment in very preterm infants. *Arch. Dis. Child Fetal Neonatal.* Ed. **105**(5), 460–465 (2020)
20. Kuklisova-Murgasova, M., et al.: A dynamic 4D probabilistic atlas of the developing brain. *Neuroimage* **54**(4), 2750–2763 (2011)
21. Kuklisova-Murgasova, M., Quaghebeur, G., Rutherford, M.A., Hajnal, J.V., Schnabel, J.A.: Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Med. Image Anal.* **16**(8), 1550–1564 (2012)
22. Larsen, A.B.L., Sønderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. In: International Conference on Machine Learning, pp. 1558–1566. PMLR (2016)
23. Latif, S., Asim, M., Usman, M., Qadir, J., Rana, R.: Automating motion correction in multishot MRI using generative adversarial networks. *arXiv preprint arXiv:1811.09750* (2018)
24. Lefèvre, J., et al.: Are developmental trajectories of cortical folding comparable between cross-sectional datasets of fetuses and preterm newborns? *Cereb. Cortex* **26**(7), 3023–3035 (2015)
25. Li, M., Tang, H., Chan, M.D., Zhou, X., Qian, X.: DC-AL GAN: pseudoprogression and true tumor progression of glioblastoma multiform image classification based on DCGAN and AlexNet. *Med. Phys.* **47**(3), 1139–1150 (2020)
26. Makropoulos, A., et al.: Regional growth and atlasing of the developing human brain. *Neuroimage* **125**, 456–478 (2016)

27. Makropoulos, A., et al.: The developing human connectome project: a minimal processing pipeline for neonatal cortical surface reconstruction. *Neuroimage* **173**, 88–112 (2018)
28. Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., Bronstein, M.M.: Geometric deep learning on graphs and manifolds using mixture model CNNs. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017
29. Morel, B., et al.: Automated brain MRI metrics in the Epirmex cohort of preterm newborns: Correlation with the neurodevelopmental outcome at 2 years. *Diagn. Interv. Imaging* **102**(4), 225–232 (2021)
30. Ran, M., et al.: Denoising of 3D magnetic resonance images using a residual encoder-decoder Wasserstein generative adversarial network. *Med. Image Anal.* **55**, 165–180 (2019)
31. Robinson, E.C., et al.: Multimodal surface matching with higher-order smoothness constraints. *Neuroimage* **167**, 453–465 (2018)
32. Robinson, E.C., et al.: MSM: a new flexible framework for multimodal surface matching. *Neuroimage* **100**, 414–426 (2014)
33. Schuh, A., et al.: A deformable model for the reconstruction of the neonatal cortex. In: *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pp. 800–803. IEEE (2017)
34. Shimony, J.S., et al.: Comparison of cortical folding measures for evaluation of developing human brain. *Neuroimage* **125**, 780–790 (2016)
35. Vosylius, V., et al.: Geometric deep learning for post-menstrual age prediction based on the neonatal white matter cortical surface. In: Sudre, C.H., et al. (eds.) *UNSURE/GRAIL -2020*. LNCS, vol. 12443, pp. 174–186. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-60365-6_17
36. Welander, P., Karlsson, S., Eklund, A.: Generative adversarial networks for image-to-image translation on multi-contrast MR images—a comparison of CycleGAN and unit. *arXiv preprint arXiv:1806.07777* (2018)
37. Williams, L.Z., Fawaz, A., Glasser, M.F., Edwards, D., Robinson, E.C.: Geometric deep learning of the human connectome project multimodal cortical parcellation. *bioRxiv* (2021)
38. Xia, T., Chartsias, A., Wang, C., Tsaftaris, S.A., Initiative, A.D.N., et al.: Learning to synthesise the ageing brain without longitudinal data. *Med. Image Anal.* **73**, 102169 (2021)
39. Yan, P., Xu, S., Rastinehad, A.R., Wood, B.J.: Adversarial image registration with application for MR and TRUS image fusion. In: Shi, Y., Suk, H.-I., Liu, M. (eds.) *MLMI 2018*. LNCS, vol. 11046, pp. 197–204. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00919-9_23
40. Yan, W., et al.: The domain shift problem of medical image segmentation and vendor-adaptation by Unet-GAN. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11765, pp. 623–631. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_69
41. Yang, H., et al.: Unpaired brain MR-to-CT synthesis using a structure-constrained CycleGAN. In: Stoyanov, D., et al. (eds.) *DLMIA/ML-CDS -2018*. LNCS, vol. 11045, pp. 174–182. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00889-5_20
42. Yang, Q., et al.: Low-dose CT image denoising using a generative adversarial network with Wasserstein distance and perceptual loss. *IEEE Trans. Med. Imaging* **37**(6), 1348–1357 (2018)

43. Yi, X., Walia, E., Babyn, P.: Generative adversarial network in medical imaging: a review. *Med. Image Anal.* **58**, 101552 (2019)
44. You, C., et al.: Ct super-resolution GAN constrained by the identical, residual, and cycle learning ensemble (GAN-circle). *IEEE Trans. Med. Imaging* **39**(1), 188–203 (2019)
45. Zenati, H., Foo, C.S., Lecouat, B., Manek, G., Chandrasekhar, V.R.: Efficient GAN-based anomaly detection. arXiv preprint [arXiv:1802.06222](https://arxiv.org/abs/1802.06222) (2018)
46. Zhao, F., et al.: Spherical U-net on cortical surfaces: methods and applications. CoRR abs/1904.00906 (2019). <http://arxiv.org/abs/1904.00906>
47. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2223–2232 (2017)