Continuous GCN-GANs for Modelling Neonatal Cortical Surface Development

Abdulah Fawaz[1] ABDULAH.FAWAZ@KCL.AC.UK
Logan Z. J. Williams[1,2]
Simon Dahan[1]
A. David Edwards[2,3,4]
Emma C. Robinson[1,2] EMMA.ROBINSON@KCL.AC.UK
[1] *Department of Biomedical Engineering, School of Biomedical Engineering and Imaging Sciences, King's College London, London, United Kingdom*
[2] *Centre for the Developing Brain, Department of Biomedical Engineering, School of Biomedical Engineering and Imaging Sciences, King's College London, London, United Kingdom*
[3] *Department for Forensic and Neurodevelopmental Sciences, Institute of Psychiatry,Psychology and Neuroscience, King's College London, London, SE5 8AF, United Kingdom*
[4] *MRC Centre for Neurodevelopmental Disorders, King's College London, London, SE1 1UL, United Kingdom*

Editor: Editor's name

Abstract

Structural MRI-derived features of the cortical surface are known to correlate to phenotypes such as age, sex and cognitive outcomes. Deep generative modelling of cortical neurodevelopment can lead to clinically interpretable models of disease or identify atypical cases for clinical intervention, but deep modelling of non-Euclidean domains, such as surfaces, poses additional challenges. In this work, we adapt a graph convolutional network (GCN) to model the neonatal cortical surface, and synthesise realistic, age-conditioned images of myelination and sulcal depth cortical surface maps. We train our models without longitudinal data, using randomised aging cycles of varying length, which we validate by ablation and with comparison to a CycleGAN. An independently trained deep regression model evaluates the accuracy of the generated images as the difference between their apparent post-menstrual age (PMA) and their respective target ages, obtaining a mean absolute error (MAE) of 1.02 ± 0.28 weeks (baseline accuracy 0.6 weeks).

Keywords: Generative Modelling, GAN, GCN, Graph Convolutional Networks, Cortical Surface, MRI

1. Introduction

The human cortex undergoes significant microstructural and functional development around 24-44 weeks post menstrual age (PMA). This process is known to be affected by multiple factors, and its disruption is associated with poorer neurodevelopmental outcomes. Deep generative models have the potential to lead to clinically interpretable models of disease, which can be used to identify atypical cortical neurodevelopment for early clinical intervention. Building such models is complicated by the fact that traditional techniques are unsuitable for use on non-Euclidean domains such as surfaces, so adapted methods from the field of geometric deep learning are necessary. Further, cortical surface development is confounded by two continuous and related variables: PMA at scan and the gestational age (GA) at birth. We address these challenges in this paper by using a Graph Convolutional Network (GCN) trained with a continuous 3-point extension to the regular CycleGAN.

2. Method

We construct discriminator and generator models using a graph convolution operator (Morris et al., 2019). Our training scheme involves generating a closed cycle of n synthetic images at randomly chosen PMA as shown in Figure 1(a), regularised with a cycle loss. Our dataset consists of samples where GA at birth $\approx$ PMA at scan (correct to an average 1.4 weeks), so the two will be used interchangeably. An age-dependent discriminator is used to train the model to classify images as realistically belonging to a given PMA (Figure 1(b)). In this paper, we focus on a model where $n = 3$ cycles but this is explored via an ablation study, and compared against a CycleGAN without continuous age-conditioning (Zhu et al., 2017; Fawaz et al., 2022). Quantitatively, we validate the models based on their accuracy in generating images of the target PMA, and the subject specificity. The former is validated using a pre-trained regressor model based on MoNet (Monti et al., 2017; Fawaz et al., 2021), with baseline accuracy 0.6 weeks mean absolute error (MAE) on predicting PMA. For the latter we use image similarity metrics to measure the similarity of the generated images to the original image. The metrics used include the Mean Square Error (MSE), the Peak-Signal-to-Noise Ratio (PSNR) and a novel adaptation of the Structural Similarity Index Measure (SSIM) to the surface.

3. Results

Quantitative results are shown in Table 1, with our model generating accurately aged images whilst retaining subject specificity. Generated image examples for a single neonate (original PMA 41 weeks) at different ages are shown in Figure 2. The results match well with expected neurodevelopmental changes documented in the literature (Williams et al., 2021).

Table 1: A Table comparing generative models on age generation accuracy and image specificity. n denotes the length of the generative cycle used in training (see Figure 1(a))

Model	**Accuracy (MAE / Weeks)**	**Structural Similarity**		
		SSIM	**PSNR**	**MSE**
CycleGAN	2.92 ± 2.1	**0.82**	16.8	**163.1**
GCN ($n = 2$)	2.70 ± 0.70	0.74	17.3	203
Our GCN ($n = 3$)	1.02 ± 0.20	**0.82**	**17.9**	167
GCN ($n = 4$)	$\mathbf{0.8 \pm 0.38}$	0.70	16.1	231
GCN ($n = 5$)	1.55 ± 0.50	0.74	16.1	208

4. Conclusion

Our GCN model can generate realistic, continuous, age progression/regression images of a single subject, accurate to a week of the target age on average, whilst retaining subject specificity.

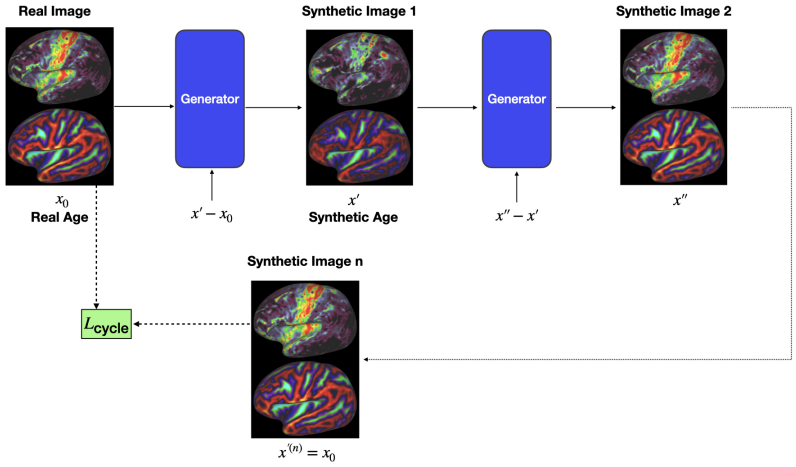

(*a*) A Generative Cycle

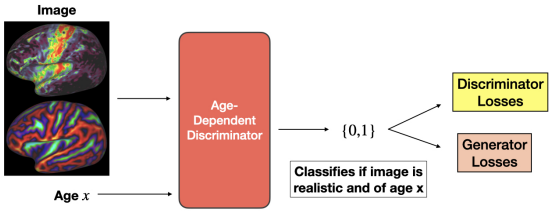

(*b*) Age-Dependent Discriminator

Figure 1: A figure demonstrating our training scheme. (a) a generative cycle involves a cycle of n synthetic images generated at randomly determined ages $x'^{\cdots}$ ending at the original image age x_0 and regularised with a L1 Cycle loss. (b) Generator and Discriminator Losses are determined with an age-dependent discriminator

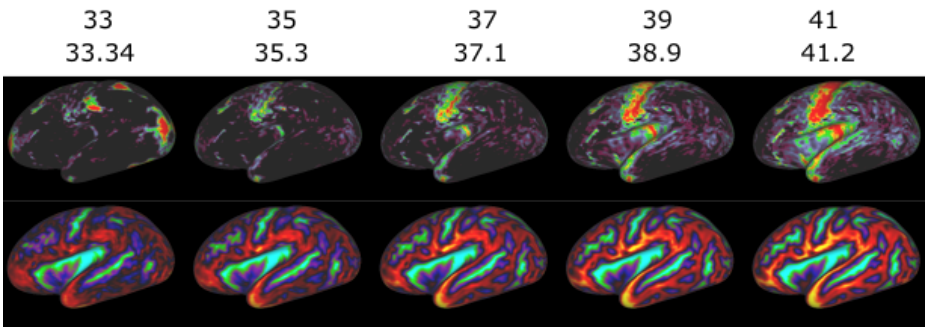

Figure 2: Synthetic images of a neonatal subject (original PMA 41.57 weeks) aged with our model. The regressor predicted ages are shown below the confounded ages. Only the lateral view is shown

References

Abdulah Fawaz, Logan ZJ Williams, Amir Alansary, Cher Bass, Karthik Gopinath, Mariana da Silva, Simon Dahan, Chris Adamson, Bonnie Alexander, Deanne Thompson, et al. Benchmarking geometric deep learning for cortical segmentation and neurodevelopmental phenotype prediction. *bioRxiv*, 2021.

Abdulah Fawaz, Logan ZJ Williams, A David Edwards, and Emma C Robinson. A deep generative model of neonatal cortical surface development. In *Annual Conference on Medical Image Understanding and Analysis*, pages 469–481. Springer, 2022.

Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5115–5124, 2017.

Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 4602–4609, 2019.

Logan Zane John Williams, Sean Patrick Fitzgibbon, Jelena Bozek, Anderson M Winkler, Ralica Dimitrova, Tanya Poppe, Andreas Schuh, Antonios Makropoulos, John Cupitt, Jonathan O'Muircheartaigh, et al. Structural and functional asymmetry of the neonatal cerebral cortex. *bioRxiv*, 2021.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.