Longitudinally Consistent Individualized Prediction of Infant Cortical Morphological Development

Xinrui Yuan[1], Jiale Cheng[1], Dan Hu[1], Zhengwang Wu[1], Li Wang[1], Weili Lin[1], Gang Li[1(✉)]

Department of Radiology and Biomedical Research Imaging Center, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA
gang_li@med.unc.edu

Abstract. Neurodevelopment is exceptionally dynamic and critical during infancy, as many neurodevelopmental disorders emerge from abnormal brain development during this stage. Obtaining a full trajectory of neurodevelopment from existing incomplete longitudinal data can enrich our limited understanding of normal early brain development and help identify neurodevelopmental disorders. Although many regression models and deep learning methods have been proposed for longitudinal prediction based on incomplete datasets, they have two major drawbacks. First, regression models suffered from the strict requirements of input and output time points, which is less useful in practical scenarios. Second, although existing deep learning methods could predict cortical development at multiple ages, they predicted missing data independently with each available scan, yielding inconsistent predictions for a target time point given multiple inputs, which ignores longitudinal dependencies and introduces ambiguity in practical applications. To this end, we emphasize temporal consistency and develop a novel, flexible framework named longitudinally consistent triplet disentanglement autoencoder to predict an individualized longitudinal cortical developmental trajectory based on each available input by encouraging the similarity among trajectories with a dynamic time-warping loss. Specifically, to achieve individualized prediction, we employ a surfaced-based autoencoder, which decomposes the encoded latent features into identity-related and age-related features with an age estimation task and identity similarity loss as supervisions. These identity-related features are further combined with age conditions in the latent space to generate longitudinal developmental trajectories with the decoder. Experiments on predicting longitudinal infant cortical property maps validate the superior longitudinal consistency and exactness of our results compared to baselines'.

Keywords: Cortical Prediction, Individualized Development, Longitudinal Consistency.

1 Introduction

The human brain undergoes exceptionally dynamic and critical development during infancy. Delineating a complete individualized developmental trajectory is essential to

understand individual uniqueness and identify abnormal development. However, it's impossible to obtain complete longitudinal scans in clinical scenarios due to the absence of patients at certain scan ages and limitations in imaging protocol designs. Therefore, predicting individualized developmental changes based on incomplete longitudinal scans is critical for comprehensively investigating individualized development and understanding the divergence between prediction and ground truth, which is essential for designing prompt individualized interventions.

Recently, several methods made success in predicting missing data during early brain development, e.g., the regression models [1,2,15] and learning-based models [18]. These methods either capitalized on the inherent structure and inter-relation between existing and missing data or leveraged the geometric and dynamic features of the cortical surface to model/predict the dynamic development. However, these computational methods imposed strict data requirements for training the models, which require scans at predefined time points or complete longitudinal scans. Consequently, many longitudinal data fail to meet these requirements, resulting in the discard of various irregularly distributed but valuable data. Alternatively, deep learning-based models have emerged as powerful tools in many scenarios, such as age prediction [13], cognitive prediction [4,5], development prediction [7,8,10,16,24], disease diagnosis [3,17], etc., leveraging their impressive interpretability and flexibility to capture profound latent features. For example, [7,8] have sought to investigate the developing cortex in cortical maturation, which used the generative models directly conditioned on the post-menstrual age and gestational age. [13, 23] used the feature disentanglement to predict development with incomplete longitudinal data. Although these deep learning-based methods demonstrate promising results and are devoid of stringent data requirements, they encounter an inherent drawback of predicting based on each input time point independently, ignoring their intrinsic relationship. These approaches lead to inconsistent and implausible results for the same time point when using input from different time points.

To address these issues, we propose a novel, flexible deep learning method to predict individualized full longitudinal cortical property maps within 24 months of age. Our approach builds on a surface-based autoencoder [25] and uses an attention-based strategy [22] to conduct the disentanglement of identity-related and age-related components. To fully leverage irregularly distributed and limited training samples in each age group, we adopted the triplet units followed by triplet autoencoder [12, 13, 23]. To reduce the ambiguity introduced by independent predictions and highlight the longitudinal consistency, we leveraged the dynamic warping loss [6] to constrain the predicted trajectory based on different input time points. To learn a more controllable and meaningful generative model, we enforced the cycle consistency between forward and backward mappings, inspired by CycleGAN [27]. The main contributions of this paper can be summarized as follows: 1) we formulated missing data prediction as a one-to-sequence prediction and proposed a novel and flexible framework capable of taking the cortical property map at any time point as input to generate a complete developmental trajectory. 2) We designed a novel feature disentanglement and an identity-preserved fusion block to take advantage of the irregularly distributed longitudinal dataset and effectively preserve identity-related patterns along the individualized cortical development. 3) We enforced temporal consistency by leveraging the dynamic time-warping

loss to guarantee consistent and meaningful results across different inputs. To the best of our knowledge, this is the first flexible individualized longitudinal prediction framework that overcomes limitations such as strict data requirements, the inability to preserve temporal consistency, and the lack of interpretability. Extensive experimental results validate the superiority of our proposed model over other state-of-the-art methods.

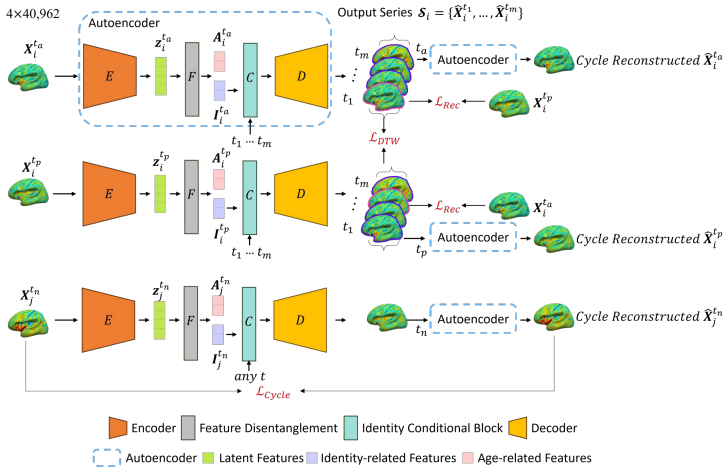

Fig. 1. The framework of the proposed longitudinally consistent triplet disentanglement autoencoder. Cortical thickness, sulcal depth, surface area, and myelin content were used as the input cortical property maps as well as the prediction target. For clarity, only the cortical thickness map was shown in the figure.

2 Method

2.1 Network Architecture

We developed our method based on [24], but in an autoencoder manner to leverage its simplicity and lightweight nature. We used four cortical property maps, i.e., cortical thickness, sulcal depth, surface area, and myelin content, as our input cortical property maps as well as prediction target. The framework constitutes four primary parts, i.e., an encoder, a feature disentanglement module, an identity conditional block, and a decoder as illustrated in **Fig. 1**.

The encoder (E) first extracts the latent features $\boldsymbol{z}$: $\boldsymbol{z} = E(\boldsymbol{X})$, incorporating the fundamental operations of spherical CNN proposed in Spherical U-Net [25] that extends convolution, pooling, and upsampling operations to the spherical surface. Concretely, the encoder comprises 5 repeated spherical 1-ring Convolution (Conv)+Batch Normalization (BN)+LeakyReLU layers with 4 spherical max pooling layers interspersed. E consists of five resolution steps (40,962, 10,242, 2,562, 642, and 162 vertices) and the feature channels at each resolution are set to 16, 32, 64, 128, and 256, respectively.

To detect and separate entangled identity information and age information from $\boldsymbol{z}$, we introduce a feature disentanglement module (F) for its reliability and simplicity [14,

24]. The feature disentanglement module comprises channel and spatial attentions in parallel to capture age-related information, $\boldsymbol{A} = \boldsymbol{z} \otimes \sigma(\boldsymbol{z})$, where $\otimes$ denotes element-wise multiplication and σ denotes attention operation, which can be supervised by an age prediction task. While the residual part, $\boldsymbol{I} = \boldsymbol{z} \otimes (1 - \sigma(\boldsymbol{z}))$, regarded as the identity-related information, will be supervised by an intensive triplet loss [12]. Thus, the identity and age information are well disentangled from latent features: $(\boldsymbol{A}, \boldsymbol{I}) = F(\boldsymbol{z})$. Herein, the channel attention in F contains two repeated 1-ringConv+BN+LeakyReLU following a sigmoid layer at the end, while the spatial attention contains one 1-ringConv+BN+LeakyReLU with a sigmoid layer at the end. The feature channel in F is 256.

To maintain identity invariance and achieve the identity-level evolving patterns, the identity conditional block (C) takes identity-related information and uses a weight-sharing strategy to smoothly obtain individualized developmental features. By specifying the target age t, we can filter out the individualized developmental features at t. Then the decoder (D) gradually upsamples the individualized developmental features and reconstructs the target age with the original size, which can be formulated as: $\hat{\boldsymbol{x}}^t = D(C(\boldsymbol{I}), t)$. Specifically, the decoder consists of 5 repeated spherical 1-ring transposed Convolutional+BN+LeakyReLU layers to deal with the concatenation and one spherical 1-ring Convolution (Conv)+Batch Normalization (BN)+LeakyReLU layer with a sigmoid layer at the end, where the resolutions are 162, 642, 2,562, 10,242, 40,962, and 40,962 vertices, and the corresponding feature channels are set to 256, 128, 64, 32, 16, and 4, respectively.

2.2 Loss Design

The objective of optimizing the whole model includes four types of losses: 1) a disentangling loss for decomposing the identity- ($\boldsymbol{I}$) and age-related features ($\boldsymbol{A}$) from latent features ($\boldsymbol{z}$); 2) A longitudinal consistency loss for guaranteeing the temporal consistency between the same predicted series from different inputs; 3) A reconstruction loss to directly constrain the vertex-wise similarity between the generated surface maps and the corresponding ground truth; 4) A cycle-consistent loss to prevent the decoder from producing irrelevant surface maps.

Disentangling Loss. To guarantee accurate and meaningful feature disentanglement, the disentangling loss includes two goals: 1) obtaining unique identity-related features; and 2) obtaining evolving age-related features. As illustrated in **Fig. 1**, the triplet units $(\boldsymbol{X}_i^{t_a}, \boldsymbol{X}_i^{t_p}, \boldsymbol{X}_j^{t_n})$, the first two surfaces are from the same subject i at different ages, i.e., t_a and t_p, while the last one is from a different subject j at any age t_n. We assume the identity-related features from the same subjects should be similar, while those from different subjects should be disparate. To construct abundant training units from limited scans and reinforce the relationship between triplet units, intensive triplet loss was employed to ensure the effectiveness of identity-related feature extraction:

$$\mathcal{L}_{Tri} = \sum_{i \neq j}^{N} \left(Corr\left(\boldsymbol{I}_i^{t_p}, \boldsymbol{I}_j^{t_n}\right) + Corr\left(\boldsymbol{I}_i^{t_a}, \boldsymbol{I}_j^{t_n}\right) - 2 \cdot Corr\left(\boldsymbol{I}_i^{t_a}, \boldsymbol{I}_i^{t_p}\right) \right), \quad (1)$$

where *Corr* denotes the Pearson's correlation coefficient and N denotes the total number of training samples.

To further ensure the age-related features contain developing information, the cross-entropy loss was applied to supervise the disentanglement of age-related features:

$$\mathcal{L}_{Age} = -\sum_{i=1}^{N}\sum_{t\in T_i} y_i^t \cdot \log\left(w(\boldsymbol{A}_i^t)\right), \tag{2}$$

where w refers to the age classifier using age-related feature $\boldsymbol{A}_i^t$, T_i refers to the collection of time points with available scans of subject i, and y_i^t refers to ground truth age one-hot label. The cross-entropy loss thus encourages $\boldsymbol{A}$ to capture age information same as the correct age label. w consists of two convolutional layers with BN+LeakyReLU in between.

Longitudinal Consistency Loss. To reinforce the connection and coherence between longitudinal scans and decrease the ambiguity, a dynamic time warping loss [6,19] was applied to ensure the predicted trajectories to be similar and match with each other. The trajectory $\boldsymbol{S}_i$ consists of a set of predicted cortical property maps at various time points, denotes as $\boldsymbol{S}_i = \{\widehat{\boldsymbol{X}}_i^{t_1}, \dots, \widehat{\boldsymbol{X}}_i^{t_m}\}$. Thus, the longitudinal consistency loss was calculated as:

$$\mathcal{L}_{DTW} = \sum_{i=1}^{N} \text{soft-DTW}_{\gamma}(\boldsymbol{S}_i^{t_a},\ \boldsymbol{S}_i^{t_p}), \tag{3}$$

where soft-DTW_{γ} is the dynamic warping distance calculated using [20], γ equals to 0.1, $\boldsymbol{S}_i^{t_a}$ and $\boldsymbol{S}_i^{t_p}$ are the predicted trajectories based on t_a and t_p, respectively.

Reconstruction Loss. Incorporating the longitudinal consistency loss compels two predicted trajectories to exhibit similarity; however, it is critical to avoid a scenario where these trajectories deviate significantly from the corresponding ground truths. Consequently, for direct prediction of cortical property map at t_2 from t_1: $\left(\boldsymbol{A}_i^{t_1}, \boldsymbol{I}_i^{t_1}\right) = F\left(E\left(\boldsymbol{x}_i^{t_1}\right)\right), \widehat{\boldsymbol{x}}_i^{t_2} = D\left(C(\boldsymbol{I}_i^{t_1}), t_2\right)$, where t_1 and t_2 can be any age in 24 months, then we have incorporated a reconstruction loss into our framework, serving the purpose of ensuring the meaningful and reliable generation of surface maps:

$$\mathcal{L}_{Rec} = \sum_{i=1}^{N}\left\|\widehat{\boldsymbol{X}}_i^{t_2} - \boldsymbol{X}_i^{t_2}\right\|_2^2. \tag{4}$$

Cycle-Consistent Loss. To achieve a generative model that is simultaneously more controllable and meaningful, we incorporated additional constraints applied to the backward mapping process during the training of the decoder. As shown in **Fig. 1**, after generating the trajectories from each input, we randomly selected one generated map at t_2 and backward mapped it to the input at t_1: $\left(\boldsymbol{A}_i^{t_1}, \boldsymbol{I}_i^{t_1}\right) = F\left(E\left(\boldsymbol{X}_i^{t_1}\right)\right), \widehat{\boldsymbol{X}}_i^{t_2} = D\left(C(\boldsymbol{I}_i^{t_1}), t_2\right), \left(\boldsymbol{A}_i^{t_2}, \boldsymbol{I}_i^{t_2}\right) = F\left(E\left(\widehat{\boldsymbol{X}}_i^{t_2}\right)\right), \widehat{\boldsymbol{X}}_i^{t_1} = D\left(C(\boldsymbol{I}_i^{t_2}), t_1\right)$, enforcing the cycle-

consistent loss to guarantee the meaningfulness between inputs maps and cycle generated maps:

$$\mathcal{L}_{Cycle} = \sum_{i=1}^{N} \left\| D\left(F\left(E\left(\widehat{\boldsymbol{X}}_i^{t_2}\right),\ t_1\right) - \boldsymbol{X}_i^{t_1}\right\|_1. \quad (5)$$

3 Experiments and Results

3.1 Dataset

All experiments were evaluated using neuroimaging data from the public Baby Connectome Project (BCP) dataset [11]. Cortical property maps were derived from T1- and T2-weighted MRI images, processed and registered using an infant-dedicated pipeline [21]. We employed left hemispheres with cortical thickness, sulcal depth, surface area, and myelin content, which were initially mapped onto the sphere, nonlinearly aligned, and further resampled with a regular spherical mesh with 40,962 vertices [9]. We used 437 longitudinal scans of 209 typically developing infants within 24 months of age. All subjects were stratified and split into training and testing set with a portion of 7:3, and a 5-fold cross-validation in training set was used for tuning the parameters.

3.2 Experimental Settings

The network was implemented based on the public Spherical U-Net code [25]. We used four cortical property maps, including cortical thickness, sulcal depth, surface area, and myelin content, as our input channels. Each cortical property map was normalized between 0 and 1 across the dataset. The network was trained with an Adam optimizer with a fixed learning rate 5e-4 in an end-to-end manner for 100 epochs. The weights of different loss terms were empirically set as 0.01, 1.0, 10.0, 1e-5, 1.0 for $\mathcal{L}_{Age}$, $\mathcal{L}_{Tri}$, $\mathcal{L}_{Rec}$, $\mathcal{L}_{DTW}$, and $\mathcal{L}_{Cycle}$, respectively.

We compare the proposed model with the following models: (1) the conditional Spherical U-Net based on vanilla Spherical U-Net [26] concatenates with one-hot encoding, specifying the target age, along the channel dimension in the upsampling blocks (CSUNet); (2) The lightweight DITSAA shares the same framework in [24] (DITSAA). These two models share the same feature channel setting as our model at each corresponding resolution for a fair comparison.

3.3 Results

Validation on Preserving Longitudinal Consistency. To evaluate if the longitudinal consistency is successfully preserved, we first illustrated the individual-level trajectories predicted by different inputs in **Fig. 2**. As shown, CSUNet exhibited very similar trajectories across different inputs. However, it always predicted the same value at any age, indicating a tendency towards averaged results and overlooking individualized patterns. DITSAA preserved more individualized patterns, but significant gaps were observed between results based on different inputs. It lacked consideration for temporal

consistency and introduced ambiguity in clinical scenarios. In contrast, various developmental trajectories forecasted by different inputs using our method demonstrated strong consistency among themselves and closely aligned with the ground truth.

Moreover, we also calculated the average distance in terms of DTW using Eq. (3), between all trajectories in **Table 1**. Our model outperformed the other competing methods by achieving the lowest average trajectory distance in predicting four cortical property maps. The quantitative results demonstrated the effectiveness of longitudinal consistency between predicted trajectories from different inputs.

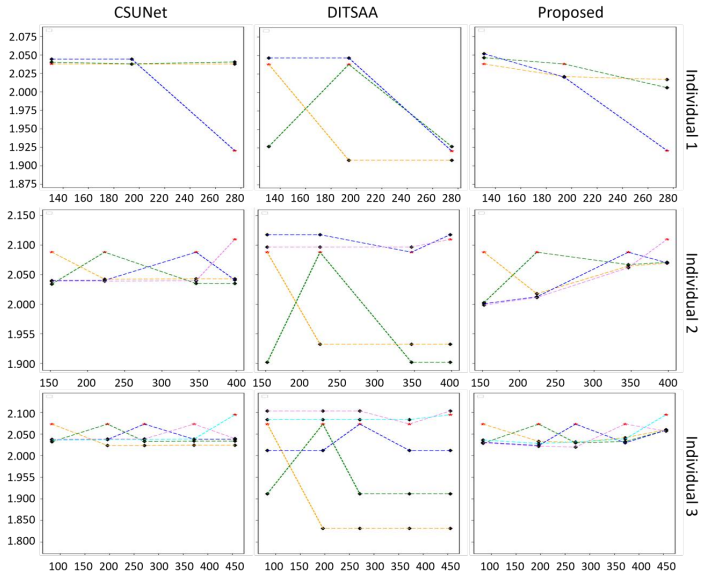

Fig. 2. Comparison of the predicted individual trajectories from different inputs with different methods. The x-axis represents ages in days, and the y-axis represents average cortical thickness (mm). The red stars denote ground truths at corresponding ages, while the black diamonds denote predictions at corresponding ages. Each color of the dashed line represents the predicted trajectory based on the source input at the red star.

Validation on Individualized Prediction. Ensuring longitudinal consistency among generated cortical property maps is of paramount importance; nevertheless, our primary objective is to acquire cortical property maps that are both meaningful and dependable. Consequently, an investigation into the predictive performance of our models was also undertaken. We computed the mean absolute error (MAE) and Pearson's correlation coefficient (PCC) between the predicted and expected cortical property maps. The average MAE and PCC of four cortical property maps are shown in **Table 2**. Our model showed superior quantitative results to other competing models, with lower MAE and higher PCC in cortical thickness, surface area, and myelin content. Besides, we also visualized average absolute error maps obtained by different methods, as illustrated in **Fig. 3**. The average absolute error maps were consistent with quantitative results.

Overall, based on the information depicted in **Table 1** and **Table 2**, we can conclude that our model excels not only in maintaining longitudinal consistency but also in achieving reliable individualized predictions.

Table 1. Performance comparison of different methods in terms of DTW distance (magnitude: 10^3) to evaluate the longitudinal consistency. Mean and standard deviation values (Mean ± Std) of the testing results based on 5-fold cross-validation are reported.

	Cortical Thickness	Sulcal Depth	Surface Area	Myelin Content
CSUNet	6.279±0.551	163.9±14.28	0.897±0.060	3.921±0.381
DITSAA	8.647±1.203	190.8±67.06	1.055±0.177	5.923±1.353
Proposed	6.071±0.300	171.5±7.149	0.845±0.041	3.584±0.209

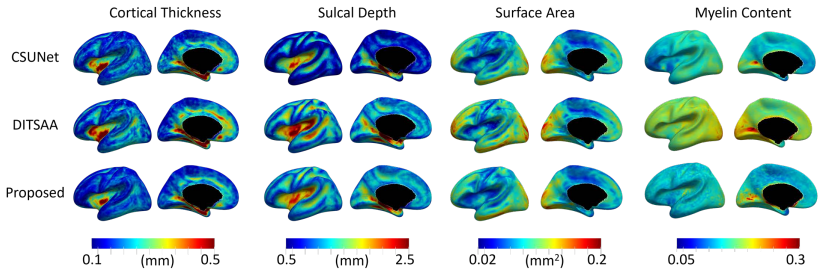

Fig. 3. Average absolute error maps of four cortical property maps by different methods.

Table 2. Prediction performance comparison between different methods. Mean and standard deviation values (Mean ± Std) of the testing results based on 5-fold cross validation are reported.

	Cortical Thickness (mm)		Sulcal Depth (mm)		Surface Area (mm^2)		Myelin Content	
Methods	MAE	PCC	MAE	PCC	MAE	PCC	MAE	PCC
CSUNet	0.247± 0.007	0.830± 0.006	0.964± 0.080	0.977± 0.003	0.085± 0.004	0.921± 0.003	0.163± 0.002	0.655± 0.026
DITSAA	0.238± 0.014	0.835± 0.008	1.091± 0.036	0.975± 0.002	0.091± 0.004	0.909± 0.015	0.207± 0.003	0.642± 0.013
Proposed	0.209± 0.005	0.865± 0.008	0.972± 0.073	0.977± 0.001	0.077± 0.001	0.925± 0.001	0.158± 0.002	0.671± 0.020

4 Conclusion

In this paper, to address the practical yet emphasized problem in individualized longitudinal developmental prediction, we proposed a novel, flexible deep learning method to predict individualized full longitudinal cortical property maps within 24 months of age. We leveraged the interpretability of the disentangling strategy, the efficacy of longitudinal consistency, and the enhanced meaningfulness and reliability afforded by cycle-consistency. This allows us to successfully address the intricate challenges of

individualized prediction using incomplete, irregularly distributed longitudinal datasets, which is cumbersome and remains unresolved for existing methods. Both visual and quantitative results substantiate the efficacy of our model, demonstrating its potential for precise individual-level development prediction and providing inspiration for personalized interventions.

Acknowledgments. This work was supported in part by NIH grants (MH116225, MH123202, ES033518, AG075582, NS128534, and NS135574).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Adeli, E., Meng, Y., Li, G., Lin, W., Shen, D.: Joint sparse and low-rank regularized multi-task multi-linear regression for prediction of infant brain development with incomplete data. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 40–48. Springer (2017)
2. Adeli, E., Meng, Y., Li, G., Lin, W., Shen, D.: Multi-task prediction of infant cognitive scores from longitudinal incomplete neuroimaging data. NeuroImage **185**, 783–792 (2019)
3. Bass, C., da Silva, M., Sudre, C., Tudosiu, P.D., Smith, S., Robinson, E.: Icam: interpretable classification via disentangled representations and feature attribution mapping. Advances in Neural Information Processing Systems **33**, 7697–7709 (2020)
4. Cheng, J., Zhang, X., Ni, H., Li, C., Xu, X., Wu, Z., Wang, L., Lin, W., Li, G.: Path signature neural network of cortical features for prediction of infant cognitive scores. IEEE Transactions on Medical Imaging **41**(7), 1665–1676 (2022)
5. Cheng, J., Zhang, X., Zhao, F., Wu, Z., Yuan, X., Wang, L., Lin, W., Li, G.: Prediction of infant cognitive development with cortical surface-based multimodal learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 618–627. Springer (2023)
6. Cuturi, M., Blondel, M.: Soft-dtw: a differentiable loss function for time-series. In: International Conference on Machine Learning. pp. 894–903. PMLR (2017)
7. Fawaz, A., Williams, L.Z.J., Dahan, S., Edwards, A.D., Robinson, E.C.: Continuous gcn-gans for modelling neonatal cortical surface development. In: Geometric Deep Learning in Medical Image Analysis (Extended abstracts) (2022)
8. Fawaz, A., Williams, L.Z., Edwards, A.D., Robinson, E.C.: A deep generative model of neonatal cortical surface development. In: Annual Conference on Medical Image Understanding and Analysis. pp. 469–481. Springer (2022)
9. Fischl, B.: Freesurfer. Neuroimage **62**(2), 774–781 (2012)
10. Hong, Y., Kim, J., Chen, G., Lin, W., Yap, P.T., Shen, D.: Longitudinal prediction of infant diffusion mri data via graph convolutional adversarial networks. IEEE Transactions on Medical Imaging **38**(12), 2717–2725 (2019)
11. Howell, B.R., Styner, M.A., Gao, W., Yap, P.T., Wang, L., Baluyot, K., Yacoub, E., Chen, G., Potts, T., Salzwedel, A., et al.: The unc/umn baby connectome project (bcp): An overview of the study design and protocol development. NeuroImage **185**, 891–905 (2019)
12. Hu, D., Wang, F., Zhang, H., Wu, Z., Wang, L., Lin, W., Li, G., Shen, D., Consortium, U.B.C.P.: Disentangled intensive triplet autoencoder for infant functional connectome fingerprinting. In: Medical Image Computing and Computer Assisted Intervention–MICCAI

2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VII 23. pp. 72–82. Springer (2020)
13. Hu, D., Zhang, H., Wu, Z., Wang, F., Wang, L., Smith, J.K., Lin, W., Li, G., Shen, D.: Disentangled-multimodal adversarial autoencoder: Application to infant age prediction with incomplete multimodal neuroimages. IEEE Transactions on Medical Imaging **39**(12), 4137–4149 (2020)
14. Huang, Z., Zhang, J., Shan, H.: When age-invariant face recognition meets face age synthesis: A multi-task learning framework and a new benchmark. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(6), 7917–7932 (2023)
15. Meng, Y., Li, G., Rekik, I., Zhang, H., Gao, Y., Lin, W., Shen, D.: Can we predict subject-specific dynamic cortical thickness maps during infancy from birth? Human Brain Mapping **38**(6), 2865–2874 (2017)
16. Pinaya, W.H., Tudosiu, P.D., Dafflon, J., Da Costa, P.F., Fernandez, V., Nachev, P., Ourselin, S., Cardoso, M.J.: Brain imaging generation with latent diffusion models. In: MICCAI Workshop on Deep Generative Models. pp. 117–126. Springer (2022)
17. Ravi, D., Alexander, D.C., Oxtoby, N.P., Initiative, A.D.N.: Degenerative adversarial neuroimage nets: generating images that mimic disease progression. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 164–172. Springer (2019)
18. Rekik, I., Li, G., Lin, W., Shen, D.: Prediction of longitudinal development of infant cortical surface shape using a 4d current-based learning framework. In: Information Processing in Medical Imaging: 24th International Conference, IPMI 2015, Sabhal Mor Ostaig, Isle of Skye, UK, June 28-July 3, 2015, Proceedings 24. pp. 576–587. Springer (2015)
19. Tavenard, R., Faouzi, J., Vandewiele, G., Divo, F., Androz, G., Holtz, C., Payne, M., Yurchak, R., Rußwurm, M., Kolar, K., Woods, E.: Tslearn, a machine learning toolkit for time series data. Journal of Machine Learning Research **21**(118), 1–6 (2020)
20. Lee, K.: Soft-dtw-loss. bluehttps://github.com/keonlee9420/Soft-DTW-Loss, (2021)
21. Wang, L., Wu, Z., Chen, L., Sun, Y., Lin, W., Li, G.: ibeat v2. 0: a multisite-applicable, deep learning-based pipeline for infant cerebral cortical surface reconstruction. Nature Protocols **18**(5), 1488–1509 (2023)
22. Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: Cbam: Convolutional block attention module. In: Proceedings of the European conference on computer vision (ECCV). pp. 3–19 (2018)
23. Xia, T., Chartsias, A., Tsaftaris, S.A., Initiative, A.D.N.: Consistent brain ageing synthesis. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV 22. pp. 750–758. Springer (2019)
24. Yuan, X., Cheng, J., Zhao, F., Wu, Z., Wang, L., Lin, W., Zhang, Y., Li, G.: Multi-task joint prediction of infant cortical morphological and cognitive development. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 545–554. Springer (2023)
25. Zhao, F., Wu, Z., Wang, L., Lin, W., Gilmore, J.H., Xia, S., Shen, D., Li, G.: Spherical deformable u-net: Application to cortical surface parcellation and development prediction. IEEE Transactions on Medical Imaging **40**(4), 1217–1228 (2021)
26. Zhao, F., Xia, S., Wu, Z., Duan, D., Wang, L., Lin, W., Gilmore, J.H., Shen, D., Li, G.: Spherical u-net on cortical surfaces: methods and applications. In: Information Processing in Medical Imaging: 26th International Conference, IPMI 2019, Hong Kong, China, June 2–7, 2019, Proceedings 26. pp. 855–866. Springer (2019)

27. 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)