Age Conditional Longitudinal Forecasting of Adolescent Functional Connectivity via Brownian Bridge Diffusion Models

Rongye Zhang[1], Shasha Xu[1], Xiaobo Liu[2], Jing Bian[1], Rui Cao[1*], and Xin Wen[1 [0000−0003−2363−1190]*]

[1] Taiyuan University of Technology, China
xwen@tyut.edu.cn
[2] McGill University, Canada

Abstract. To predict and generate developmental changes in brain functional connectivity (FC), we propose a framework combining a Graph Attention Autoencoder with Brownian Bridge diffusion. The autoencoder encodes FC into low dimensional features that preserve topology, while the diffusion model anchors the source and target domains and leverages age conditioning for stable trajectory modeling. This approach enables direct prediction and generation of the FC of target-age from FC of initial-age. Validated on longitudinal fMRI data from the CCNP cohort (ages 6–18), the framework outperforms existing generative models on key metrics, including mean squared error (MSE), Pearson correlation coefficient (PCC), and structural similarity index (SSIM), while capturing biologically significant developmental brain regions, providing a reliable tool for predictive FC modeling. Code is available at https://github.com/rongyezhang18-eng/BBdiffusion-FC.

Keywords: Adolescent brain development · Functional connectivity · Graph representation learning · Brownian bridge diffusion · Resting-state fMRI.

1 Introduction

Functional magnetic resonance imaging (fMRI) constructs functional connectivity (FC) from blood oxygen level dependent (BOLD) signals, capturing functional associations between atlas-defined regions of interest (ROIs)[1]. Adolescence is a critical period of rapid brain network reorganization, and systematically characterizing FC remodeling in this stage is key to understanding neurodevelopment and individual variability[2, 3]. Using longitudinal FC data, we aim to learn a generative model that, conditioned on early-stage FC, predicts subject-specific FC at arbitrary target ages, thereby addressing temporal sparsity and irregular sampling in longitudinal studies and providing a computational framework for continuous developmental trajectory prediction.

Recently, interest has grown in modeling the connectome's two fundamental representations—structural connectivity (SC) and functional connectivity

(FC)[4]. FC-to-FC modeling focuses on how FC patterns change over time and across states and ages, capturing network plasticity and individual developmental trajectories[5]. However, most existing work either uses cross-sectional analyses or SC–FC mapping and does not explicitly incorporate physiological factors such as age[4, 6, 9]. Longitudinal FC-to-FC prediction remains rare, and current approaches, typically based on state estimation or traditional statistical models, struggle to forecast FC in a way that both matches the target state and preserves intrinsic network geometry.

In contrast, deep learning–based generative models learn data distributions rather than single input–output mappings. Variational Autoencoders (VAEs)[8] enable latent variable modeling but often produce overly smooth outputs due to the trade-off between reconstruction accuracy and regularization[10]. Generative Adversarial Networks (GANs)[7] address this loss of detail via adversarial training, but suffer from instability and mode collapse[10]. More recently, diffusion models abandon adversarial training in favor of iterative denoising, yielding stable, high-fidelity generations[11, 12]. However, standard diffusion models sample from pure noise, leading to costly sampling and complex conditioning schemes in paired translation tasks, which severely undermines training stability[13].

To tackle these challenges, we propose a framework that combines ROI level representation learning and Brownian Bridge conditional diffusion. First, a graph attention network compresses high-dimensional functional connectivity into node-level latent variables, while a structured decoder preserves key geometric and topological properties of brain networks, simplifying generation. Second, Brownian Bridge diffusion in a unified latent space models the transformation from early to late FC, with trajectories naturally constrained by both endpoints rather than pure noise injection, making it well-suited for source-to-target longitudinal brain development modeling. Finally, conditioning on age as a guiding variable enables more precise and developmentally coherent FC generation.

2 Methods

2.1 Graph Attention Autoencoder for Representation Learning

To support latent space diffusion while preserving ROI interactions, we train a GAT-based autoencoder [14, 15] that encodes functional connectivity into node-level latent representations for each brain region. The encoder comprises three GATConv layers with eight attention heads each. From $\mathrm{FC} \in \mathbb{R}^{116\times116}$, a sparse graph is constructed by treating each ROI as a node and retaining the top-k connections (We use $k = 30$ and ablate $k = 20/60$, where $k = 20$ is comparable and $k = 60$ underperforms, indicating a better sparsity–performance trade-off). Node features are then updated through multi-layer attention as:

$$\mathbf{h}_i^{(l+1)} = o\left(\sum_{j\in\mathcal{N}(i)} a_{ij}^{(l)} W^{(l)} \mathbf{h}_j^{(l)}\right) \tag{1}$$

Ultimately, the encoder generates a low dimensional latent representation for each brain region, yielding Z $\in \mathbb{R}^{116\times d}$, d=64. The decoder employs an outer

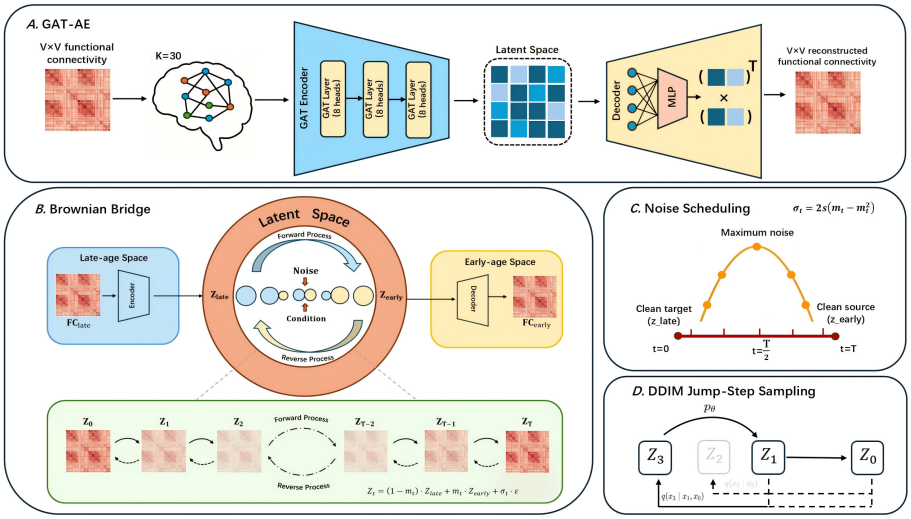

Fig. 1. Overall model architecture.(A) GAT-AE framework with graph attention encoder and outer product decoder.(B) Diffusion Bridge. Source and target domains form the two endpoints of the Brownian bridge. The forward process injects latent-space noise along a fixed interpolation path; the reverse process predicts the bridge residual with U-Net to recover the target latent.(C) Noise scheduling strategy.(D) DDIM jump-step sampling.

product mechanism to reconstruct the functional connectivity matrix, a design that inherently enforces the symmetry of the resulting output matrix. Specifically, the decoder first refines node features through an MLP, then reconstructs functional connectivity via an outer product:

$$\mathrm{FC}_{\mathrm{recon}} = \mathbf{Z}' \cdot (\mathbf{Z}')^{\mathsf{T}} \tag{2}$$

This approach naturally ensures that the FC matrix is symmetric and that the diagonal elements are fixed at 1.0 to represent regional autocorrelation.

2.2 Brownian Bridge Diffusion for Generation

Brownian bridge is a continuous stochastic process conditioned on two endpoints. Given an initial state x_0 and a terminal state x_T, the marginal distribution of an intermediate state x_t can be expressed as:

$$q\left(x_t \mid x_0, x_T\right) = \mathcal{N}\left(x_t; \left(1 - \frac{t}{T}\right) x_0 + \frac{t}{T} x_T, \frac{t(T-t)}{T}\mathbf{I}\right) \tag{3}$$

Recently, Brownian bridge processes have been extended to image-to-image translation [13, 19, 20]. To model FC matrices, we perform Brownian bridge diffusion in the latent space of the autoencoder, where $x_0 = z_{\text{late}}$ and $y = z_{\text{early}}$ denote late and early-stages of FC latents, respectively. The conditional distribution is:

$$q\left(x_t \mid x_0, y\right) = \mathcal{N}\left(x_t; (1 - m_t)x_0 + m_t y, \delta_t \mathbf{I}\right) \tag{4}$$

where $0 \leq t \leq T$ and $m_t = \frac{t}{T} \in [0, 1]$. The forward bridge maps $x_0 = z_{\text{late}}$ to $y = z_{\text{early}}$, while the reverse process predicts x_0 from y, corresponding to the early to late prediction setting.

In noise scheduling, the bridge variance is required to vanish at both endpoints and peak at an intermediate step. However, the original Brownian bridge variance $\frac{t(T-t)}{T}$ reaches $\frac{T}{4}$ at $t = \frac{T}{2}$, which increases with the number of diffusion steps and may destabilize training. Following the bounded variance principle in DDPM [16] and VPSDE [17], we adopt a quadratic variance schedule:

$$\delta_t = 2S(m_t - m_t^2) \tag{5}$$

where S controls the bridge noise scale. This variance peaks at $(m_t = \frac{1}{2})$ while vanishing at both endpoints.

Conditional design: Variables related to age are jointly encoded as $c = (\text{age}_{\text{early}}, \text{age}_{\text{late}}, \text{age}_{\text{diff}}, \text{age}_{\text{ratio}})$ for efficient multidimensional conditioning.

Training: Given $x_0 = z_{\text{late}}$ and $y = z_{\text{early}}$, the intermediate state x_t is sampled from the Brownian bridge distribution as:

$$x_t = (1 - m_t)x_0 + m_t y + \sqrt{\delta_t}\varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \mathbf{I}) \tag{6}$$

The neural network f_θ is trained to estimate the Brownian bridge residual $x_t - x_0$. Therefore, the diffusion training objective is:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{t,\varepsilon}\left[\ell\left(f_\theta(x_t, t, c), m_t(y - x_0) + \sqrt{\delta_t}\varepsilon\right)\right] \tag{7}$$

Inference: Given the observed early stage latent $y = z_{\text{early}}$, the target late stage latent is recovered by subtracting the predicted Brownian bridge residual from x_t:

$$\hat{x}_0 = x_t - f_\theta(x_t, t, c) \tag{8}$$

The reverse transition is formulated as:

$$p_\theta(x_{t-1}|x_t, y, c) = \mathcal{N}\left(x_{t-1}; \mu_\theta(x_t, t, y, c), \sigma^2_{t\to t-1}\mathbf{I}\right) \tag{9}$$

The transition variance term is defined as:

$$\delta_{t|t-1} = \delta_t - \delta_{t-1}\left(\frac{1 - m_t}{1 - m_{t-1}}\right)^2 \tag{10}$$

The posterior sampling variance is:

$$\sigma^2_{t\to t-1} = \frac{\delta_{t|t-1}\delta_{t-1}}{\delta_t} \tag{11}$$

Unlike DDPM, the Brownian bridge starts generation from the early FC latent rather than pure noise, with age covariates guiding the reverse process. Specifically, x_{t-1} is obtained through probabilistic sampling as follows:

$$x_{t-1} = \lambda_1(t)x_t + \lambda_2(t)y - \lambda_3(t)f_\theta(x_t, t, c) + \sigma_{t\to t-1}\eta \tag{12}$$

For compact notation, we define the posterior correction coefficient as:

$$\rho_{t,t-1} = \frac{\delta_{t-1}}{\delta_t} \cdot \frac{1 - m_t}{1 - m_{t-1}}. \tag{13}$$

The corresponding time-dependent coefficients are:

$$\begin{aligned} \lambda_1(t) &= 1 - m_{t-1} + m_t\rho_{t,t-1}, \\ \lambda_2(t) &= m_{t-1} - m_t\rho_{t,t-1}, \\ \lambda_3(t) &= (1 - m_{t-1}) \cdot \frac{\delta_{t|t-1}}{\delta_t} \end{aligned} \tag{14}$$

The three weights $\lambda_1(t)$, $\lambda_2(t)$, and $\lambda_3(t)$ control the contributions of the current state x_t, the starting point y, and the Brownian bridge residual predicted by $f_\theta(x_t, t, c)$ to the next state x_{t-1}. This update preserves the endpoint-constrained bridge structure while allowing both age conditioning and model predictions to guide the reverse generation process.

Furthermore, to improve inference efficiency, we adopt a non-Markovian sampling strategy [18] that allows the Brownian bridge process to be inferred using only a subset$\{x_{t_1}, \ldots, x_{t_n}\}$of time steps.

3 Experiments and Discussions

3.1 Dataset

This study analyzed longitudinal data from the Chinese Color Nest Project (CCNP) [21, 22], including 207 healthy participants aged 6–18 years. After quality control, 393 resting-state fMRI scans were preprocessed using standard procedures. Longitudinal training data were generated by pairing all within-subject scans across sessions, resulting in 1,380 sample pairs from 159 subjects, each consisting of functional connectivity matrices from two time points and their corresponding ages.

3.2 Comparative Experiments

In the generation phase, we compared the proposed method with three representative generative models, including DDPM, GAN, and RealNVP, under identical settings. Paired t-tests across 5-fold cross-validation assessed between-method differences, with p-values reported in Table 1. Our method achieved the best overall performance across all metrics. It obtained the lowest MSE (0.29), representing a 35.6% reduction compared to DDPM, and achieved the highest PCC (0.63) and SSIM (0.50), demonstrating superior reconstruction accuracy and structural consistency.

Table 1. Comparison with existing generative methods.

Models	MSE	PCC	SSIM	Cos-Sim
DDPM [12]	$0.45 \pm 0.03^{***}$	$0.60 \pm 0.01^{**}$	$0.42 \pm 0.01^{***}$	$0.95 \pm 0.002^{**}$
GAN [7]	$0.40 \pm 0.02^{***}$	$0.59 \pm 0.01^{*}$	$0.47 \pm 0.01^{**}$	$0.94 \pm 0.002^{**}$
RealNVP [23]	$0.42 \pm 0.003^{**}$	$0.61 \pm 0.003^{*}$	$0.43 \pm 0.004^{***}$	$0.95 \pm 0.001^{*}$
Ours	**0.29 ± 0.02**	**0.63 ± 0.01**	**0.50 ± 0.01**	**0.96 ± 0.003**

*: $p < 0.05$; **: $p < 0.01$; ***: $p < 0.001$

3.3 Ablation Experiments

We analyze the contribution of the AE by comparing two variants: without AE and with VAE. Removing AE causes a large performance drop (MSE 0.57, SSIM 0.23), indicating the necessity of low-dimensional representations for diffusion. Replacing AE with a VAE improves results (MSE 0.31, SSIM 0.45) but remains inferior to ours (MSE 0.29, SSIM 0.50), likely because stochastic sampling in VAEs degrades structural fidelity, while AEs better preserve patterns.

Table 2. Ablation Study of Model Components and Key Design Choices

Items		MSE	PCC	SSIM	Cos-Sim
Sampling Steps	50 steps	0.31 ± 0.02	0.62 ± 0.01	0.47 ± 0.01	0.95 ± 0.003
	100 steps	0.33 ± 0.06	0.60 ± 0.03	0.47 ± 0.03	0.95 ± 0.003
	200 steps	**0.29 ± 0.02**	**0.63 ± 0.01**	**0.50 ± 0.01**	**0.96 ± 0.003**
	1000 steps	0.31 ± 0.02	0.62 ± 0.01	0.48 ± 0.01	0.95 ± 0.002
Latent Dimensionality	Z=16	0.32 ± 0.02	0.59 ± 0.01	0.45 ± 0.02	0.95 ± 0.003
	Z=32	0.31 ± 0.03	0.62 ± 0.01	0.47 ± 0.001	0.95 ± 0.002
	Z=64	**0.29 ± 0.02**	**0.63 ± 0.01**	**0.50 ± 0.01**	**0.96 ± 0.003**
Noise Schedule Strategy	S=0.5	0.30 ± 0.02	0.62 ± 0.01	0.48 ± 0.01	0.95 ± 0.002
	S=1.0	**0.29 ± 0.02**	**0.63 ± 0.01**	**0.50 ± 0.01**	**0.96 ± 0.003**
	S=2.0	0.35 ± 0.05	0.60 ± 0.02	0.46 ± 0.02	0.95 ± 0.003
	S=4.0	0.33 ± 0.33	0.60 ± 0.01	0.46 ± 0.01	0.95 ± 0.003
Module Ablation	w/o AE	0.57 ± 0.04	0.39 ± 0.01	0.23 ± 0.002	0.86 ± 0.008
	AE→VAE	0.31 ± 0.04	0.62 ± 0.006	0.45 ± 0.01	0.95 ± 0.002
	Ours	**0.29 ± 0.02**	**0.63 ± 0.01**	**0.50 ± 0.01**	**0.96 ± 0.003**

We further study the impact of sampling steps (50, 100, 200, 1000). As the number of sampling steps increases, generation quality steadily improves and reaches its optimum at 200 steps. When the number of steps exceeds 200, performance remains nearly unchanged, while increasing to 1000 brings little additional gain but substantially increases sampling time.

We evaluate the scaling factor s in the Brownian Bridge noise schedule (Eq. 5) with $s \in \{0.5, 1, 2, 4\}$. As s increases, all four metrics deteriorate. This observation agrees with our discussion in the Methods, where adopting the original variance formulation of the Brownian Bridge yields an overly large peak variance and thus prevents meaningful sample generation.

Finally, we evaluate three latent dimensionalities (16, 32, 64). As the latent dimension increases, reconstruction performance consistently improves: MSE decreases from 0.32 to 0.29, PCC rises from 0.59 to 0.63, and SSIM from 0.45 to 0.50, suggesting that higher dimensional latent spaces provide greater representational capacity and are more suitable for stable diffusion modeling.

3.4 Performance Comparison Across Brain Atlases

We fixed the latent dimension to 64 and compared three atlases (AAL116, CC200, CC400). This experiment evaluates the parcellation most suitable for our method and examines the effect of atlas granularity on performance. AAL116 achieved the best overall performance (Table 3), whereas both CC atlases yielded lower overall results. Finer atlases fragment coherent regions, increase noise and

Table 3. Performance Comparison of the Model Across Different Brain Region Atlases

Atlases	MSE	PCC	SSIM	Cos-Sim
AAL116	0.29 ± 0.02	0.63 ± 0.01	0.50 ± 0.01	0.96 ± 0.003
CC200	0.37 ± 0.03	0.47 ± 0.01	0.37 ± 0.01	0.93 ± 0.002
CC400	0.33 ± 0.03	0.49 ± 0.01	0.39 ± 0.01	0.91 ± 0.002

inter-individual variability, and and the fixed 64 dimensional latent space may not adequately represent the increased number of brain regions; in contrast, the anatomically defined AAL116 offers more stable regional boundaries than the functionally defined CC atlases.

3.5 Longitudinal Comparative Validation

To assess whether the predicted developmental trajectories exhibit physiologically plausible changes, we used the similarity between empirical early-age and late-age FC as a baseline, reflecting the natural differences in functional connectivity observed across developmental stages. We then compared this baseline with the similarity between empirical early-age FC and model-generated late-age FC. Across all three atlases, the generated late-age FC showed similarity levels comparable to the empirical baseline (Table 4). For instance, under the AAL116 atlas, the PCC was 0.61 for generated results versus 0.63 for the baseline, with identical cosine similarity values (0.93). These results indicate that the model captures a realistic magnitude of age-related change without introducing exaggerated or attenuated deviations.

To examine early-FC copying, we treat Raw vs. Raw as an early-copy baseline, where early FC is directly used as predicted late FC, i.e., $\widehat{\mathrm{FC}}_{\mathrm{late}} = \mathrm{FC}_{\mathrm{early}}$. This strong baseline reflects subject-specific FC stability, yielding a PCC of 0.63 under AAL116. Compared with it, our model retains comparable PCC while improving MSE, SSIM, and Cos-Sim from 0.56/0.46/0.93 to 0.29/0.50/0.96, better matching true late-age FC rather than reproducing the input.

Table 4. Longitudinal Validation Results of Model-Generated Functional Connectivity

Atlases		PCC	SSIM	Cos-Sim
AAL116	Raw vs Raw (Baseline)	0.63 ± 0.13	0.47 ± 0.15	0.93 ± 0.04
	Raw vs Model-Generated	0.61 ± 0.14	0.45 ± 0.15	0.93 ± 0.04
CC200	Raw vs Raw (Baseline)	0.56 ± 0.13	0.44 ± 0.15	0.90 ± 0.06
	Raw vs Model-Generated	0.50 ± 0.17	0.40 ± 0.18	0.90 ± 0.06
CC400	Raw vs Raw (Baseline)	0.56 ± 0.12	0.43 ± 0.15	0.88 ± 0.07
	Raw vs Model-Generated	0.53 ± 0.17	0.42 ± 0.18	0.88 ± 0.07

3.6 Identification of Developmentally Critical Brain Regions

We assessed developmental plausibility by comparing age–connection correlations between ground-truth and reconstructed AAL116 FCs (FDR-corrected, $q < 0.05$) [27]. Based on these connections, developmental hub regions were

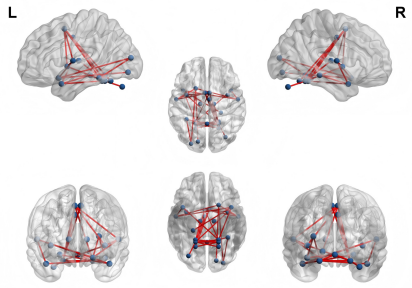

Fig. 2. Identified key brain regions[26]. Blue nodes represent brain regions, and red edges denote their age-related functional connections.

identified using an involvement rate, defined as the proportion of age-related connections per region. Notably, the identified hubs were predominantly located in higher-order cognitive networks, including the posterior cerebellum (e.g., Crus I/II and lobules VIII–X [34]), default mode–related temporal regions, and sensorimotor association areas [28–30], consistent with previous neurodevelopmental findings. In particular, the prominence of posterior cerebellar regions supports

emerging evidence for the cerebellum's role in cognitive and social maturation beyond motor control [31–33]. Overall, these results are consistent with prior evidence of selective functional reorganization during adolescence, suggesting that the proposed model learns biologically meaningful developmental patterns.

4 Conclusion

This study integrates GAT-AE with Brownian Bridge diffusion for developmental FC prediction. GAT-AE preserves FC topology in a low dimensional space, while Brownian Bridge diffusion enables age-conditioned prediction between source and target domains. Comparative and longitudinal analyses support its effectiveness and its ability to capture biologically plausible developmental patterns. Future work will expand the dataset through augmentation strategies (e.g., noise perturbation, masked occlusion) to further enhance FC prediction accuracy.

Acknowledgments. This work was supported in part by the National Natural Science Foundation of China under Grant [62206196], in part by the Natural Science Foundation of Shanxi Province under Grant [202303021221001], in part by the Natural Science Foundation for Young Scientists of Shanxi Province under Grant [202203021212216].

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

References

1. Biswal, B. et al.: Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magnetic Resonance in Medicine, 34(4), 537–541 (1995).
2. Luo, A. C. et al.: Functional connectivity development along the sensorimotor-association axis enhances the cortical hierarchy. Nature Communications, 15: 3511 (2024).
3. Calabro, F. J. et al.: Developmental decorrelation of local cortical activity through adolescence supports high-dimensional encoding and working memory. Developmental Cognitive Neuroscience, 73: 101541 (2025).
4. Jamison, Keith W., et al.: Krakencoder: a Unified Brain Connectome Translation and Fusion Tool. Nature Methods, 22: 1583–1592 (2025).
5. Yin, Weiyan, et al.: Charting Brain Functional Development from Birth to 6 Years of Age. Nature Human Behaviour, 9: 1246–1259 (2025).
6. Jain, Chirag, et al.: Diffusion Wavelets on Connectome: Localizing the Sources of Diffusion Mediating Structure-Function Mapping Using Graph Diffusion Wavelets. Network Neuroscience, 9(2): 777–797 (2025).
7. Goodfellow, Ian, and others: Generative adversarial networks. Communications of the ACM, 63(11): 139–144 (2020).
8. Kingma, Diederik P. and Welling, Max: Auto-encoding variational Bayes. In: International Conference on Learning Representations (ICLR). OpenReview.net (2014).
9. Tan, Yee-Fan, et al.: SFC-GAN: A Generative Adversarial Network for Brain Functional and Structural Connectome Translation. Computing Research Repository, abs/2501.07055 (2025).

10. Bond-Taylor, Samuel, and others: Deep generative modelling: A comparative review of VAEs, GANs, normalizing flows, energy-based and autoregressive models. IEEE Transactions on Pattern Analysis and Machine Intelligence, 44(11): 7327–7347 (2022).
11. Croitoru, Fabian-Alexandru, and others: Diffusion models in vision: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(9): 10850–10869 (2023).
12. Dhariwal, Prafulla and Nichol, Alexander Quinn: Diffusion Models Beat GANs on Image Synthesis. In: Advances in Neural Information Processing Systems, 34, pp. 8780–8794 (2021).
13. Li, Bo, and others: BBDM: Image-to-image translation with brownian bridge diffusion models. In: Conference on Computer Vision and Pattern Recognition, pp. 1952–1961 (2023).
14. Vincent, Pascal and Larochelle, Hugo and Lajoie, Isabelle and Bengio, Yoshua and Manzagol, Pierre-Antoine: Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion. Journal of Machine Learning Research, 11: 3371–3408 (2010).
15. Veličković, Petar and others: Graph Attention Networks. In: International Conference on Learning Representations (2018).
16. Ho, Jonathan and Jain, Ajay and Abbeel, Pieter: Denoising diffusion probabilistic models. In: Advances in Neural Information Processing Systems, 33, pp. 6840–6851. Curran Associates, Inc. (2020).
17. Song, Yang and Sohl-Dickstein, Jascha and Kingma, Diederik P. and Kumar, Abhishek and Ermon, Stefano and Poole, Ben: Score-based generative modeling through stochastic differential equations. arXiv:2011.13456 (2020).
18. Song, Jiaming and Meng, Chenlin and Ermon, Stefano: Denoising Diffusion Implicit Models. In: International Conference on Learning Representations (ICLR) (2021).
19. Tian, Jilong, et al.: MapGen-Diff: an End-to-End Remote Sensing Image to Map Generator Via Denoising Diffusion Bridge Model. Remote Sensing, 16(19): 3716 (2024).
20. Stoyanov, Ivan, Bongratz, Fabian, Wachinger, Christian: Spherical Brownian Bridge Diffusion Models for Conditional Cortical Thickness Forecasting. In: Shape in Medical Imaging. ShapeMI 2025. Lecture Notes in Computer Science, vol. 16171, pp. 150–162. Springer (2025).
21. Dong, Hao-Ming and Zuo, Xi-Nian: Age-specific brain functional connectivity matrices for children and adolescents at school age, https://www.scidb.cn/, last accessed 2025/01/22.
22. Yang, Ning, et al.: Chinese Color Nest Project: Growing Up in China. Chinese Science Bulletin, 62(26): 3008–3022 (2017). doi: 10.1360/N972017-00362.
23. Dinh, Laurent and Sohl-Dickstein, Jascha and Bengio, Samy: Density Estimation Using Real NVP. In: International Conference on Learning Representations (2017).
24. Tzourio-Mazoyer, Nathalie, et al.: Automated Anatomical Labeling of Activations in SPM Using a Macroscopic Anatomical Parcellation of the MNI MRI Single-Subject Brain. NeuroImage, 15(1): 273–289 (2002).
25. Craddock, Robert Cameron, et al.: A Whole Brain Fmri Atlas Generated Via Spatially Constrained Spectral Clustering. Human brain mapping, 33(8): 1914–1928 (2012).
26. Xia, Mingrui, et al.: Brainnet Viewer: A Network Visualization Tool for Human Brain Connectomics. PLoS ONE, 8(7): e68910 (2013).

27. Benjamini, Yoav and Hochberg, Yosef: Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal Statistical Society: Series B (Methodological), 57(1): 289–300 (1995).
28. Paquola, Casey, et al.: The Architecture of the Human Default Mode Network Explored Through Cytoarchitecture, Wiring and Signal Flow. Nature Neuroscience, 28: 654–664 (2025).
29. Fan, Lingzhong, et al.: Connectivity-Based Parcellation of the Human Temporal Pole Using Diffusion Tensor Imaging. Cerebral Cortex, 24(12): 3365–3378 (2014).
30. Sherman, Lauren E. and Rudie, Jeffrey D. and Pfeifer, Jennifer H. and Masten, Carrie L. and McNealy, Kristin and Dapretto, Mirella: Development of the default mode and central executive networks across early adolescence: a longitudinal study. Developmental Cognitive Neuroscience, 10: 148–159 (2014). doi: 10.1016/j.dcn.2014.08.002.
31. Buckner, Randy L.: The Cerebellum and Cognitive Function: 25 Years of Insight from Anatomy and Neuroimaging. Neuron, 80(3): 807–815 (2013).
32. Stoodley, Catherine J. and Schmahmann, J.: Functional Topography in the Human Cerebellum: A Meta-Analysis of Neuroimaging Studies. NeuroImage, 44: 489–501 (2009).
33. Strick, Peter L. and Dum, Richard P. and Fiez, Julie A.: Cerebellum and nonmotor function. Annual Review of Neuroscience, 32: 413–434 (2009). Publisher: Annual Reviews, doi: 10.1146/annurev.neuro.31.060407.125606.
34. Van Overwalle, Frank, et al.: The Posterior Crus II Cerebellum is Specialized for Social Mentalizing and Emotional Self-Experiences: a Meta-Analysis. Social cognitive and affective neuroscience, 15(9): 905–928 (2020).