

T2I-Diff: fMRI Signal Generation via Time-Frequency Image Transform and Classifier-Free Denoising Diffusion Models

Hwa Hui Tew¹, Junn Yong Loo^{1*}, Yee-Fan Tan¹, Xinyu Tang¹, Hernando Ombao², Fuad Noman¹, Raphaël C.-W. Phan¹, and Chee-Ming Ting^{1*}

¹ School of Information Technology, Monash University Malaysia

² Statistics Program, King Abdullah University of Science and Technology
{loo.junnyong, ting.cheeming}@monash.edu

Abstract. Functional Magnetic Resonance Imaging (fMRI) is an advanced neuroimaging method that enables in-depth analysis of brain activity by measuring dynamic changes in the blood oxygenation level-dependent (BOLD) signals. However, the resource-intensive nature of fMRI data acquisition limits the availability of high-fidelity samples required for data-driven brain analysis models. While modern generative models can synthesize fMRI data, they often underperform because they overlook the complex non-stationarity and nonlinear BOLD dynamics. To address these challenges, we introduce T2I-Diff, an fMRI generation framework that leverages time-frequency representation of BOLD signals and classifier-free denoising diffusion. Specifically, our framework first converts BOLD signals into windowed spectrograms via a time-dependent Fourier transform, capturing both the underlying temporal dynamics and spectral evolution. Subsequently, a classifier-free diffusion model is trained to generate class-conditioned frequency spectrograms, which are then reverted to BOLD signals via inverse Fourier transforms. Finally, we validate the efficacy of our approach by demonstrating improved accuracy and generalization in downstream fMRI-based brain network classification. The code is available at repository

Keywords: fMRI · Time-Frequency Image · Diffusion Models.

1 Introduction

With the rapid advancement of artificial intelligence, deep generative modeling has shown promising capability in generating realistic variations of neuroimaging data [26]. Although the generation of brain connectivity has been extensively studied, research on the direct generation of fMRI signals remains limited [29,1]. fMRI signals are particularly valuable as they provide critical insights into neural activity, enabling more precise assessment and identification of neuropsychiatric and neurodevelopmental disorders [14,13]. Despite their significance, acquiring

* Corresponding authors

fMRI signals is challenging, often resulting in small sample sizes that limit the performance of data-driven brain analysis models [17]. These limitations can lead to data imbalances and a lack of temporal dynamics in acquired brain signals, ultimately affecting the accuracy of predictive and diagnostic models for neurological and psychiatric conditions [2,22]. To address these challenges, generative techniques have been explored for synthesizing fMRI signals, enabling data augmentation to enhance various downstream applications [15,19].

Generative models, such as Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs), have been widely used to synthesize realistic time series data [3]. While GANs are known for generating high-quality samples, VAEs are favored for their faster generation speed. However, both models face challenges in achieving stable training, often struggling with issues like mode collapse and optimization difficulties. For instance, Yoon et al. proposes TimeGAN, which jointly trains adversarial and supervised loss components within a learned embedding space, successfully preserving both the static and dynamic characteristics of synthetic time-series data [27]. Apart from that, TimeVAE incorporates time-series components into its encoder-decoder network, enhancing interpretability in time-series generation, where it has demonstrated success in reducing overall training time compared to adversarial methods [5].

Recently, denoising diffusion probabilistic models (DDPMs) have gained traction in time-series applications due to their robustness against model collapse [10]. However, their potential for capturing the complex intrinsic properties of brain signals remains underexplored [29,1]. For example, Coletta et al. propose different types of constraints to improve time-series generation: hard constraints enforce fixed points and global minima; soft constraints introduce penalties to guide the model towards desired temporal trends [4]. Furthermore, ImagenTime leveraged the time-frequency images to generate synthetic data for general time-series benchmarks, demonstrating the feasibility of capturing intricate spectral and temporal patterns [12]. The recently proposed Diffusion-TS model introduces a novel non-autoregressive diffusion model for time-series data, including fMRI signals. This framework explicitly captures temporal dynamics through transformer-based model architecture and disentangled seasonal-trend representations, achieving success across various generation tasks [28].

To address the challenge of insufficient inductive bias in dynamic brain characterization, this paper introduces T2I-Diff, an fMRI generation framework that leverages time-frequency image representations of BOLD signals and classifier-free denoising diffusion. This approach effectively captures both spectral and dynamic features underlying fMRI BOLD signals, which are essential for downstream brain disorder classification. Additionally, our framework incorporates classifier-free denoising diffusion and EDM sampling for high-fidelity generation of time-windowed fMRI spectrograms conditioned on subject classes. The generated time-frequency images are then reverted to BOLD signals via image-to-time series transforms. Our method eliminates the need for additional classifiers to guide brain profiles, thereby reducing training redundancy and simplifying model optimization. We summarized our main contributions as follows:

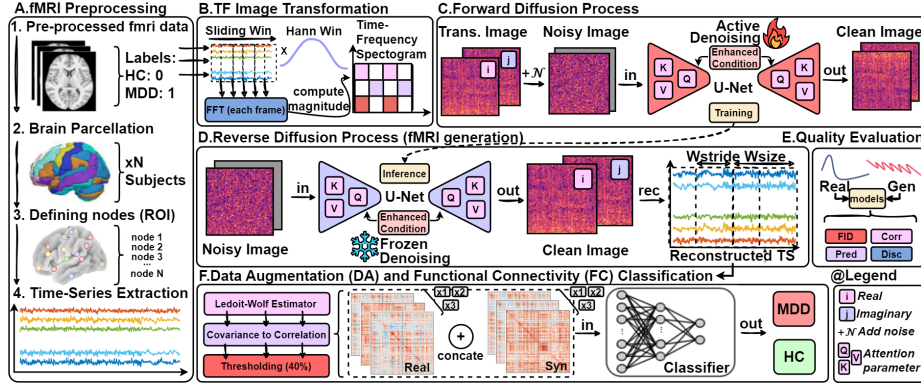


Fig. 1. Our framework T2I-Diff begins by (A.) preprocessing raw fMRI data in time domain and transforming into (B.) time-frequency spectrogram (image) using windowed Fourier transform. Next, the images are fed into the diffusion model executes both (C.) forward and (D.) reverse processes. Finally, we perform (E.) quality evaluation of the generation and apply (F.) DA on top of the real fMRI signals for MDD classification.

1. Our proposed T2I-Diff framework is the first to integrate a time-frequency image transform to capture complex spatiotemporal and spectral features for fMRI BOLD signal generation and brain disorder classification.
2. We introduce classifier-free denoising diffusion and EDM sampling to generate conditioned time-windowed fMRI spectrograms without requiring an additional classifier. This enables T2I-Diff to dynamically modulates the diffusion process to guide the model in learning different brain signal profiles.
3. Our results shows that T2I-Diff demonstrates competitive performance on time-frequency image generation and brain disorder (MDD) classification complemented by class-conditioned synthetic fMRI BOLD signals.

2 Methods

Time-Frequency Image Transformation. Fig. 1 and 2 provides an overview of our proposed framework. Given high-dimensional fMRI signals from S subjects, denoted as $\mathcal{X}_S = \{\mathbf{x}_s\}_{s=1}^S$, where each subject $\mathbf{x}_s \in \mathbb{R}^{D \times T}$ consists of D regions of interest (ROIs) recorded over T time points, our objective is to learn the underlying real data distribution $p_{\text{data}}(\mathcal{X}_S)$ and generate a synthetic distribution $p_{\theta}(\mathcal{X}_S)$ that is statistically indistinguishable from the real data.

Unlike conventional time-series generative tasks that operate exclusively in the time domain, our approach transforms fMRI time series into a time-frequency spectrogram image using the windowed Fourier transform (WFT), defined as:

$$\mathbf{X}_s(m, k) = \sum_{n=0}^{N-1} x_s[n + m h] w[n] e^{-j2\pi \frac{k}{N} n}, \quad (1)$$




Fig. 2. Original vs synthetic BOLD signals and generated normalized spectrograms. Our framework directly generates BOLD signals as opposed to correlation matrices.

which serves as the input to the generative model. Here, $n \in \{1, 2, \dots, N\}$ is the local time index within an WFT window of length N ; $m \in \{1, 2, \dots, M\}$ is the window index, where the total number of windows are computed as $M = \lfloor (T - N)/h \rfloor + 1$ with signal length T and hop size h ; and $k \in \{1, 2, \dots, K\}$ is the frequency index, with the number of frequency bins defined as $K = N/2 + 1$ since only the first half of the spectrum is retained for real-valued signals.

Given an univariate input signal $x_s \in \mathbb{R}^{1 \times T}$, the WFT produces a spectrogram image representation $\mathbf{X}_s \in \mathbb{R}^{2 \times K \times M}$, where the channels are doubled (2 channels) to store both the real \mathbf{X}_s^{Re} and imaginary \mathbf{X}_s^{Im} components of the WFT output, and the spatial dimensions (height and width) of the image are determined by the frequency bins K and time frames M , respectively. We further perform component-wise normalization to accentuate between high and low frequencies. This structured representation allows generative models to process time-series data as images while preserving both spectral and temporal characteristics, enabling robust generative modeling of the fMRI BOLD signals.

To reconstruct the original fMRI time series spectrogram representation, we first unnormalize the real and imaginary components, we merge them to form a complete complex spectrogram $\hat{\mathbf{X}}_s(m, k) = \mathbf{X}_s^{\text{Re}}(m, k) + j \mathbf{X}_s^{\text{Im}}(m, k)$. For each window indexed by m and for local time indices n , we compute the inverse fourier transform to reconstruct the fMRI BOLD signals as follows:

$$x_s[n + m h] = \frac{1}{N} \sum_{k=0}^{K-1} \hat{\mathbf{X}}_s(m, k) e^{j 2\pi \frac{k}{N} n}. \quad (2)$$

Finally, these segments are recombined using the overlapping method with Hann window to obtain the reconstructed time-domain signal \tilde{x}_s . This process ensures that both the spectral and temporal characteristics of the original fMRI BOLD signals are preserved.

Classifier-Free Denoising Diffusion. Rather than incrementally adding noise through a Markov chain, we employ EDM framework [7], which perturb the ground-truth fMRI spectrogram images using a single noise level σ_t drawn from a lognormal distribution. At each diffusion step $t \in \{0, 1, \dots, T\}$, the images are corrupted according to the perturbation kernel $p_\theta(x_t | x_0)$ which corresponds to

$x_t = x_0 + n_t$, where the total noise $n_t = \sigma_t \cdot \epsilon_t$ with $\sigma_t = \exp(\eta_t \cdot P_{\text{std}} + P_{\text{mean}})$, $\eta_t \sim \mathcal{N}(0, 1)$ and $\epsilon_t \sim \mathcal{N}(0, I)$. Here, n_t represents a sample from a Gaussian distribution whose variance is modulated by the lognormal factor σ_t . This process gradually transforms the fMRI spectrogram images into non-informative Gaussian noise. The reverse process aims to learn the conditional distribution $p_\theta(x_{t+1} | x_t)$ to reconstruct the original fMRI spectrogram images. In particular, a deep neural network denoiser $\epsilon_\theta(x_t, \sigma_t)$ is trained to predict the clean image x_0 from its noisy version x_t at noise level σ_t across time steps. Specifically, the noise scheduling is defined as $\sigma_t = \left(\sigma_{\text{max}}^{1/\rho} + \frac{t}{T-1} (\sigma_{\text{min}}^{1/\rho} - \sigma_{\text{max}}^{1/\rho}) \right)^\rho$. Using Heun’s second-order method, the reverse-diffusion update is given by:

$$x_{t+1} = x'_t + (\sigma_{t+1} - \sigma'_t) \cdot \frac{1}{2} \left[\frac{x'_t - \epsilon_\theta(x'_t, \sigma'_t)}{\sigma'_t} + \frac{x'_{t+1} - \epsilon_\theta(x_{t+1}, \sigma_{t+1})}{\sigma_{t+1}} \right], \quad (3)$$

where $\sigma'_t = \sigma_t + \gamma_t \sigma_t$ and $x'_t = x_t + \sqrt{\sigma'^2_t - \sigma_t^2} \epsilon'_t$ is perturbed by noise $\epsilon_t \sim \mathcal{N}(0, S_{\text{noise}}^2 I)$, where both γ_t and S_{noise} are EDM hyperparameters. By iteratively applying this update from $t = T$ down to $t = 0$, a synthetic fMRI spectrogram image $x_0 \sim p_\theta(x)$ is sampled from the marginal data distribution.

To achieve class-conditioned generation, we apply classifier-free guidance and parameterize the conditional distribution $p_\theta(x_t | c)$ via a conditional noise-prediction model $\epsilon_\theta(x_t, t, c)$, where c is the subject class. For unconstrained generation, we then input a null token \emptyset for the unconditional model $\epsilon_\theta(x_t, t) = \epsilon_\theta(x_t, t, \emptyset)$. Concurrently, these noise-prediction models are trained to minimize the discrepancy between the true noise ϵ and the predicted noise ϵ_θ across timesteps, via the noise-matching objective $\mathcal{L}(\theta) = \mathbb{E}_{x_t, \epsilon, t, c} [\|\epsilon - \epsilon_\theta\|^2]$. The unconditional and conditional models are jointly trained by randomly setting c to the unconditional class identifier \emptyset with some discrete probability p_\emptyset .

3 Experiments and Results

Data Acquisition and Pre-processing. We preprocessed the dataset from REST-meta-MDD Consortium database [24] using Data Processing Assistant for Resting-State fMRI (DPARSF) [25]. This resting-state fMRI dataset consist of 250 healthy controls (HC) and 227 Major Depressive Disorder (MDD) patients. The data were acquired using a Siemens (Tim Trio 3T) scanner (TR/TE = 2000/30 ms, 3mm slice thickness). The brain was parcellated into 116 region-of-interests (ROIs), including cortical and subcortical areas, and the mean time series of 232 time points for each ROI was extracted using the Automated Anatomical Labeling (AAL) atlas.

State-of-the-art Baselines. i) Quality Evaluation: Our proposed T2I-Diff model is compared against six time-series generative model baselines. These baselines include generative adversarial networks (GANs) and diffusion models such as CoT-GAN [23], DiffTime [4], DiffWave [9], TimeVAE [5] and Diffusion-TS [28]. ii) Classification Score: We further evaluate the performance of our

Table 1. Comparison of fMRI signal generation quality with SOTA time-series generation and our proposed T2I-Diff.

	CoT-GAN	DiffTime	DiffWave	TimeVAE	TimeGAN	Diffusion-TS	T2I-Diff
Context-FID	7.813±.550	0.340±.015	0.244±.018	14.449±.969	0.126±.002	0.105±.006	1.384±.107
Correlational	26.824±.449	1.501±.048	3.927±.049	17.296±.526	23.502±.039	1.411±.042	4.121±.094
Discriminative	0.492±.018	0.245±.051	0.402±.029	0.476±.044	0.484±.042	0.167±.023	0.400±.059
Predictive	0.185±.003	0.100±.000	0.101±.000	0.113±.003	0.126±.002	0.099±.000	0.102±.001

Table 2. Ablation study of short-term and long-term fMRI signal generations across sequence lengths 24, 64, 128 and 256.

24	64	128	256		24	64	128	256
Context-FID					Correlation			
1.384±0.107	7.141±0.789	10.655±1.001	1.661±0.146		4.121±0.094	4.737±0.136	3.525±0.101	4.124±0.243
Discriminative					Predictive			
0.400±0.059	0.280±0.268	0.176±0.090	0.360±0.164		0.102±0.001	0.100±0.001	0.098±0.001	0.095±0.002

proposed model with a special designed classifier for brain connectivity [8]. The baselines include 1D-Deep Convolutional GAN (1D-DCGAN) [16] and Wasserstein GAN Gradient Policy(WGAN-GP) [6].

Implementation Details. i) Connectivity Network Construction: The functional connectivity of brain networks is constructed using Ledoit-Wolf (LDW) regularized shrinkage covariance estimator to keep the strongest $\tau = 40\%$ connections and set other connections to zeros. This results in FCs of size 116×116 for each subject. ii) T2I-Diff Training: The proposed T2I-Diff framework generates the fMRI signals corresponding to the subjects' condition (HC and MDD). The BrainNetCNN classifier then discriminates between the HC and MDD subjects. We train the T2I-Diff via an Adam optimizer using a learning rate of $3e^{-4}$ for 1000 epochs. For all experiments, we use 18 diffusion sampling steps and experiment in image sizes of 8×8 , 16×16 , 32×32 , and 64×64 . iii) Data Augmentation and Classifier Training: The trained T2I-Diff is used to augment fMRI signals of $1 \times$, $2 \times$ and $3 \times$ on top of the real fMRI signals. For our classifier, the L2 regularization weight decay from 10^{-8} to 10^{-2} , scheduler learning rate reduce factor from 0.1 to 0.9, batch size from 5 to 16 same as in [18]. These model hyperparameters are selected based on a 5-fold stratified cross-validation.

3.1 Overall Performance

We first trained our diffusion models unconstrained to produce similar outputs. Then, we followed the standard setting for the evaluation of time series generation from Diffusion-TS [28]. The purpose of our experiments is to demonstrate the effectiveness of the proposed time-frequency representation and classifier-free diffusion models in achieving a trade-off between FID and predictive scores,

Table 3. Classification performance of different classifiers trained on the ground-truth data and increasing amount of augmented time series data using our proposed model.

Method	Train Set	Accuracy	Recall	Precision	F1-Score	ROC
W/O Augmentation	Real	58.90 \pm 2.98	58.90 \pm 2.98	59.56 \pm 2.74	58.39 \pm 3.09	59.00 \pm 2.56
1D-DCGAN	Real + Synth 1 \times	62.94 \pm 2.01	62.94 \pm 2.01	63.43 \pm 2.20	62.23 \pm 2.68	62.71 \pm 2.26
	Real + Synth 2 \times	65.04 \pm 2.02	65.04 \pm 2.02	66.35 \pm 2.13	64.12 \pm 2.10	64.74 \pm 2.08
	Real + Synth 3 \times	58.21 \pm 2.98	58.21 \pm 2.98	55.70 \pm 6.58	52.86 \pm 4.14	57.38 \pm 3.11
WGAN-GP	Real + Synth 1 \times	66.02 \pm 4.25	66.02 \pm 4.25	66.22 \pm 4.24	65.93 \pm 4.20	65.95 \pm 4.13
	Real + Synth 2 \times	64.76 \pm 4.25	64.76 \pm 4.25	65.67 \pm 4.08	64.23 \pm 4.52	64.73 \pm 4.14
	Real + Synth 3 \times	64.56 \pm 3.18	64.56 \pm 3.18	64.78 \pm 3.17	64.38 \pm 3.15	64.41 \pm 3.08
T2I-Diff (Ours)	Real + Synth 1 \times	66.87 \pm 3.22	66.87 \pm 3.22	67.06 \pm 3.34	66.83 \pm 3.21	67.26 \pm 6.00
	Real + Synth 2 \times	65.41 \pm 2.37	65.41 \pm 2.37	66.30 \pm 1.67	64.73 \pm 2.80	65.75 \pm 3.22
	Real + Synth 3 \times	66.03 \pm 1.75	66.03 \pm 1.75	66.50 \pm 1.32	65.85 \pm 1.82	66.58 \pm 5.33

comparable to other time-series generative models. Rather than solely aiming for state-of-the-art sample quality metrics on these benchmarks, our primary goal is to validate the ability of our approach to model complex spatiotemporal patterns and excel in capturing conditional distribution over time.

fMRI Signal Generation Quality. Table 1 compares our T2I-Diff framework with state-of-the-art time-series generative models across four evaluation metrics for a time-series generation length of 24. T2I-Diff demonstrates competitive overall performance, surpassing strong time-series generation models such as CoT-GAN, DiffWave, TimeGAN, and TimeVAE. Notably, the low discriminative and predictive scores suggest that T2I-Diff effectively generates synthetic fMRI BOLD signals with in-distribution temporal patterns and spectral features that may not be fully captured in the time domain. By leveraging the WFT image transform, the diffusion model not only preserves global contextual information but also models the complex spatiotemporal dynamics inherent in fMRI signals. While T2I-Diff does not achieve the highest overall scores, our results demonstrate that the time-series-to-image generative framework is well-suited for modeling high-dimensional, complex fMRI time series.

Ablation Study. Table 2 assess the impact of sequence length on model performance, we conducted ablation studies on time-series generation lengths of 64, 128, and 256 to contrast the sequence length of 24. Shorter sequences exhibit a low context-fid score and high predictive score, but their lower predictive score suggests a limited ability to capture temporal dynamics due to insufficient contextual information. Mid-length sequences (64 and 128) show improved context-fid scores but struggle with increased complexity. The longest sequence (256) achieves the best context-fid and predictive scores, effectively capturing both global and fine-grained features.

Classification Score. To validate the fidelity of the generated samples, we evaluate the classification performance of BrainNetCNN and compare it against

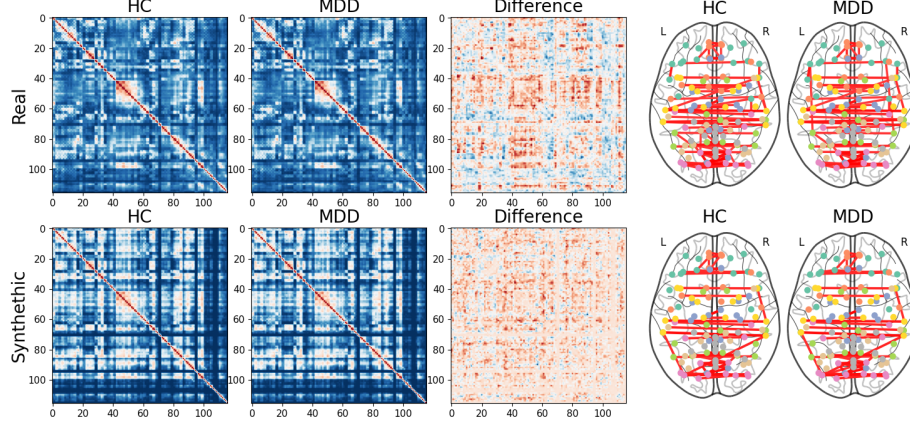


Fig. 3. Real and synthetic (left) connectivity pattern and (right) brain networks.

ID-DCGAN and WGAN-GP on our proposed fMRI dataset. Here, we maintain the original sequence length rather than truncating, as supported by the performance of long sequence generation in Table 1. Table 3 presents the classification results on the 5-fold cross-validation test set. Notably, T2I-Diff achieves the highest accuracy in the $1\times$ data augmentation setting. Furthermore, the results indicate that as the augmentation range increases, our model exhibits lower variance, suggesting that the synthetic data generalizes well across varying augmentation levels while maintaining consistent classification performance. The results highlight that T2I-Diff not only enhances the diversity of synthetic fMRI samples but also effectively captures and preserves critical structural and functional patterns for downstream tasks such as brain disorder classification.

3.2 Functional Connectivity (FC) Visualization

To further evaluate the quality of generated data, we derived functional connectivity (FC) measures from the synthetic fMRI BOLD signals to assess differences between healthy controls (HC) and major depressive disorder (MDD) patients. We computed the average connectivity using a threshold of 0.6 to highlight significant connections. Our analysis reveals that the synthetic FC closely aligns with the functional changes observed in the real FC distribution. Furthermore, connectogram comparisons between HC and MDD in both real and synthetic FC data indicate a reduction in connectivity within the left superior frontal gyrus (FrontalSupL) and decreased connectivity between the left middle frontal gyrus (FrontalMidL) and the anterior cingulate cortex (CingulumAntL). The results suggest impaired cognitive functions, such as difficulties in decision-making and emotion regulation, reinforcing the biological plausibility of the generated data.

4 Conclusions and Future Work

In this paper, we propose T2I-Diff, which effectively captures both temporal dynamics and spectral evolution underlying the ground-truth data distribution for accurate brain signal generation. For future work, we aim to further validate MDD classification using graph-based deep learning models [21,20]. Furthermore, we plan to incorporate energy-based models to identify out-of-distribution (OOD) patterns in brain spectrogram associated with neurological disorders [11].

Acknowledgments. This work was supported by the Advanced Computing Platform (ACP), Monash University Malaysia. The work of Junn Yong Loo is supported by Monash University under the SIT Collaborative Research Seed Grants 2024 I-M010-SED-000242.

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

References

1. Avberšek, L.K., Repovš, G.: Deep learning in neuroimaging data analysis: Applications, challenges, and solutions. *Frontiers in neuroimaging* **1**, 981642 (2022)
2. Bollmann, S., Barth, M.: New acquisition techniques and their prospects for the achievable resolution of fmri. *Progress in neurobiology* **207**, 101936 (2021)
3. Bond-Taylor, S., Leach, A., Long, Y., Willcocks, C.G.: 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 (2021)
4. Coletta, A., Gopalakrishnan, S., Borrajo, D., Vyetenko, S.: On the constrained time-series generation problem. *Advances in Neural Information Processing Systems* **36**, 61048–61059 (2023)
5. Desai, A., Freeman, C., Wang, Z., Beaver, I.: Timevae: A variational auto-encoder for multivariate time series generation. *arXiv preprint arXiv:2111.08095* (2021)
6. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein gans. *Advances in neural information processing systems* **30** (2017)
7. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems* **35**, 26565–26577 (2022)
8. Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., Zwicker, J.G., Hamarneh, G.: Brainnetcnn: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage* **146**, 1038–1049 (2017)
9. Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761* (2020)
10. Lin, L., Li, Z., Li, R., Li, X., Gao, J.: Diffusion models for time-series applications: a survey. *Frontiers of Information Technology & Electronic Engineering* **25**(1), 19–41 (2024)

11. Loo, J.Y., Yu, L.F., Adeline, M., Lau, J.K., Tew, H.H., Pal, A., Baskaran, V.M., Ting, C.M., Phan, R.C.: Learning energy-based generative models via potential flow: A variational principle approach to probability density homotopy matching. *Transactions on Machine Learning Research* (2025)
12. Naiman, I., Berman, N., Pemper, I., Arbiv, I., Fadlon, G., Azencot, O.: Utilizing image transforms and diffusion models for generative modeling of short and long time series. *Advances in Neural Information Processing Systems* **37**, 121699–121730 (2024)
13. Noman, F., Ting, C.M., Kang, H., Phan, R.C.W., Ombao, H.: Graph autoencoders for embedding learning in brain networks and major depressive disorder identification. *IEEE Journal of Biomedical and Health Informatics* **28**(3), 1644–1655 (2024)
14. Noman, F., Yap, S.Y., Phan, R.C.W., Ombao, H., Ting, C.M.: Graph autoencoder-based embedded learning in dynamic brain networks for autism spectrum disorder identification. In: 2022 IEEE International Conference on Image Processing (ICIP). pp. 2891–2895 (2022)
15. Power, J.D., Mitra, A., Laumann, T.O., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E.: Methods to detect, characterize, and remove motion artifact in resting state fmri. *neuroimage* **84**, 320–341 (2014)
16. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* (2015)
17. Tan, Y.F., Loo, J.Y., Ting, C.M., Noman, F., Phan, R.C.W., Ombao, H.: Brainfc-gan: A conditional generative adversarial network for brain functional connectivity augmentation and aging synthesis. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1511–1515. IEEE (2024)
18. Tan, Y.F., Ting, C.M., Noman, F., Phan, R.C.W., Ombao, H.: Graph-regularized manifold-aware conditional wasserstein gan for brain functional connectivity generation. *arXiv preprint arXiv:2212.05316* (2022)
19. Tan, Y.F., Ting, C.M., Noman, F., Phan, R.C.W., Ombao, H.: A unified framework for static and dynamic functional connectivity augmentation for multi-domain brain disorder classification. In: Proceedings of the 2023 IEEE International Conference on Image Processing (ICIP). pp. 635–639 (2023)
20. Tew, H.H., Ding, F., Li, G., Loo, J.Y., Ting, C.M., Ding, Z.Y., Tan, C.P.: St-hcss: Deep spatio-temporal hypergraph convolutional neural network for soft sensing. *arXiv preprint arXiv:2501.02016* (2025)
21. Tew, H.H., Li, G., Ding, F., Luo, X., Loo, J.Y., Ting, C.M., Ding, Z.Y., Tan, C.P.: Kans: Knowledge discovery graph attention network for soft sensing in multivariate industrial processes. In: 2024 IEEE International Conference on Systems, Man, and Cybernetics (SMC). pp. 4377–4383. IEEE (2024)
22. Ting, C.M., Skipper, J.I., Noman, F., Small, S.L., Ombao, H.: Separating stimulus-induced and background components of dynamic functional connectivity in naturalistic fmri. *IEEE Transactions on Medical Imaging* **41**(6), 1431–1442 (2022)
23. Xu, T., Wenliang, L.K., Munn, M., Acciaio, B.: Cot-gan: Generating sequential data via causal optimal transport. *Advances in neural information processing systems* **33**, 8798–8809 (2020)
24. Yan, C.G., Chen, X., Li, L., Castellanos, F.X., Bai, T.J., Bo, Q.J., Cao, J., Chen, G.M., Chen, N.X., Chen, W., et al.: Reduced default mode network functional connectivity in patients with recurrent major depressive disorder. *Proceedings of the National Academy of Sciences* **116**(18), 9078–9083 (2019)

25. Yan, C., Zang, Y.: Dparsf: a matlab toolbox for " pipeline" data analysis of resting-state fmri. *Frontiers in systems neuroscience* **4**, 1377 (2010)
26. Yap, S.Y., Loo, J.Y., Ting, C.M., Noman, F., Phan, R.C.W., Razi, A., Dowe, D.L.: A deep probabilistic spatiotemporal framework for dynamic graph representation learning with application to brain disorder identification. In: *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*. pp. 5353–5361 (2024)
27. Yoon, J., Jarrett, D., Van der Schaar, M.: Time-series generative adversarial networks. *Advances in neural information processing systems* **32** (2019)
28. Yuan, X., Qiao, Y.: Diffusion-ts: Interpretable diffusion for general time series generation. *arXiv preprint arXiv:2403.01742* (2024)
29. Zhao, H., Lou, H., Yao, L., Peng, W., Adeli, E., Pohl, K.M., Zhang, Y.: Diffusion models for computational neuroimaging: A survey. *arXiv preprint arXiv:2502.06552* (2025)