Task-aligned fMRI Generation Model for Brain Disorder Diagnosis

Yifan Li[1*], Xiaotong Wu[1*], Xiaocai Zhang[2], Haiteng Jiang[3], Weiwen Wu[1], Dinggang Shen[4,5,6], and Jianjia Zhang[1(✉)]

[1] School of Biomeical Engineering, Shenzhen Campus of Sun Yat-sen University, Shenzhen 518107, China, {zhangjj225, wuweiw7}@mail.sysu.edu.cn; {liyf279, wuxt39}@mail2.sysu.edu.cn

[2] Faculty of Engineering and Information Technology, The University of Melbourne, Melbourne VIC3010, Australia, xiaocai.zhang@unimelb.edu.au

[3] Advanced Brain Cognition and Disease Laboratory, Zhejiang University, Hangzhou 310058, China, h.jiang@zju.edu.cn

[4] School of Biomedical Engineering & State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai 201210, China

[5] Shanghai United Imaging Intelligence Co., Ltd., Shanghai 200232, China

[6] Shanghai Clinical Research and Trial Center, Shanghai 201210, China

Dinggang.Shen@gmail.com

Abstract. Functional magnetic resonance imaging (fMRI) is essential for understanding and diagnosing brain disorders. However, the challenge of small sample sizes, due to high acquisition costs and low annotation efficiency, hinders deeper exploration of the mechanisms underlying brain diseases. Recently, generative diffusion models have shown great potential for time series data generation, but directly using them for fMRI generation still has some issues. Firstly, most of them are designed for single time series, ignoring the significant dependency information between multiple time series when applied to fMRI. Since fMRI time series from different brain regions exhibit correlations, it is necessary to consider this characteristic when generating fMRI. Secondly, the generation process often lacks the involvement of label information, which limits their applicability in facilitating classification tasks. Thirdly, the alignment between the generated data and the target tasks is often insufficient, limiting its effectiveness for brain disorder diagnosis. To address these issues, we propose a novel task-aligned fMRI generation method based on the diffusion model. Specifically, a functional brain network (FBN) is incorporated into the diffusion model as prior knowledge to guide and constrain the data generation process, ensuring that the generated fMRI respects the functional connectivity characteristics observed in actual fMRI. To effectively and flexibly generate class-specific fMRI, a representative class-wise FBN is utilized as the prior FBN. Meanwhile, the proposed method ensures that the generated fMRI is well aligned with target brain disorder classification tasks. Extensive experiments are

* These authors contributed equally to this work.

conducted on three datasets, consistently demonstrating the superior performance of the proposed method.

Keywords: Diffusion Model · Functional Brain Network · fMRI · Brain Disorder Diagnosis.

1 Introduction

Functional magnetic resonance imaging (fMRI), with its ability to detect abnormal brain activity, has played a crucial role in diagnosing various brain disorders, such as mild cognitive impairment (MCI) of early Alzheimer's disease (AD) and vascular cognitive impairment (VCI) [3, 4, 18]. Meanwhile, computer-aided diagnosis (CAD) technology has provided powerful capabilities to uncover latent patterns within fMRI data, further advancing research in brain disorders. However, the superior performance of CAD systems often relies on large-scale, high-quality datasets, while the fMRI-based brain disorder diagnosis is typically limited by the small sample size due to the high cost and difficulty of collecting medical samples. Data generation technology, capable of extending datasets, is increasingly recognized as a promising solution to promote diagnosis accuracy and reliability of CAD systems [11, 20, 25].

Typical generative approaches mainly include generative adversarial network (GAN) [8, 17, 19], variational autoencoder (VAE) [12, 15, 16] and diffusion model (DM) [10, 29]. These models have found wide applications in fMRI generation. For example, the deep recurrent VAE (DRVAE) [21] and VAE-GAN [22] models are proposed to model and generate fMRI. In contrast to GAN and VAE, DM is more stable in generating high-quality data [2, 5]. While BrainNetDiffusion [32] leverages the strengths of DM to generate FBNs and focuses on the important functional connectivity of fMRI, it overlooks that the generated FBNs would lose the dynamic temporal information present in the fMRI time series and lacks the flexibility to align with downstream models that dynamically construct and analyze FBNs from raw time series. Therefore, it is crucial to develop a diffusion model-based fMRI time series generation model that retains the temporal dynamics and can be well aligned with downstream tasks.

Despite diffusion models showing advantages in generating various types of time series data, direct application of them in fMRI generation still has some drawbacks. Firstly, most diffusion models for time series generation are designed to generate single sequences in fields such as audio and finance [9, 14, 24], which often fail to fully consider the significant correlation information between different sequences when applied to multiple time series, i.e., fMRI. Note that, the functional connectivity (FC) information between brain regions, which presents the inherent functional interactions of the brain regions, is the crucial characteristic of fMRI signals [6, 27]. Therefore, integrating these complex and non-negligible inter-regional interactions into the generation method can help generate more accurate fMRI. Although the latest studies [26, 31] have proposed models for generating multiple time series, they still have some limitations. These

models typically generate data without label information, reducing their applicability in generating multi-class datasets for classification tasks. Additionally, most generation methods are designed only for data generation, which often fail to align with the downstream tasks, e.g., brain disorder classification.

To address the above limitations, we propose a novel task-aligned fMRI time series generation model, named TA-fMRI-GM, based on the diffusion model and incorporating class-wise FBN as prior knowledge. Concerning the first issue, the prior FBN, which represents the functional dependency patterns between brain regions, is introduced to guide each step of the generation process, ensuring that the FBN of the generated fMRI aligns with the prior FBN. This enables the generation model to generate fMRI data that more closely resemble the actual fMRI samples. To address the second issue, we propose using the representative FBN of each class as the prior FBN during the generation process to flexibly generate fMRI of the corresponding category, extending the applicability of the generated data for classification tasks. Regarding the third issue, to ensure alignment between the generated fMRI and downstream tasks, the generated fMRI is repeatedly evaluated on brain disorder classification tasks. The superior performance across multiple datasets consistently demonstrates the effectiveness of the proposed fMRI generation model.

Major contributions of this paper can be summarized as follows:

1. We propose a novel method for fMRI generation, which introduces the prior FBN into the diffusion model to guide and constrain the generation process, allowing the generated fMRI to fully respect the functional connectivity characteristic and more closely resemble the actual fMRI.
2. The proposed method can flexibly generate fMRI data across different categories by setting the class-wise FBN with label information as the prior FBN during the generation process.
3. The fMRI data generated from the proposed generation method is well aligned with the brain disease classification tasks.
4. Comprehensive evaluations are conducted to validate the effectiveness of the proposed method, and experimental results consistently demonstrate that the generated fMRI could benefit the performance of brain disease diagnosis.

2 Proposed Method

2.1 Overall architecture

The proposed model, designed to generate the fMRI time series that fully considers the rich and complex functional dependencies among different brain regions, is based on the diffusion model and is induced with prior FBN during the data generation process. Meanwhile, this method flexibly generates fMRI of different classes by introducing the class-wise FBN that contains label information. As illustrated in Fig. 1, the proposed method consists of two main processes: the training process, which is utilized to learn the latent features and data distribution, and the generation process, which generates fMRI data from the random noise and is guided with the class-wise prior FBN.

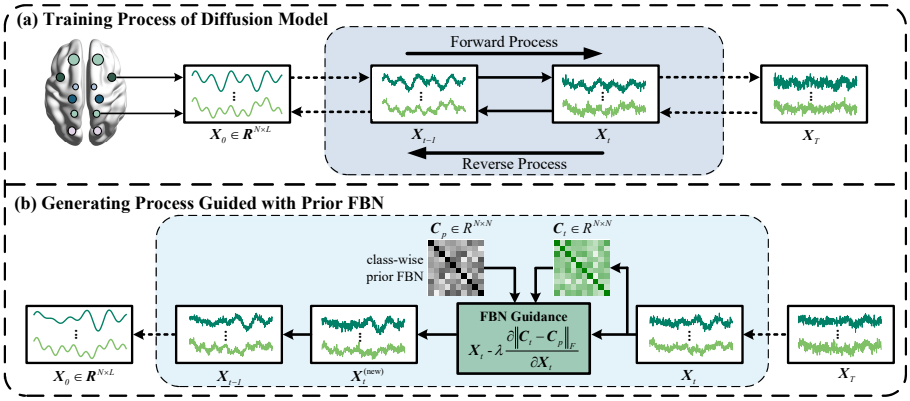

Fig. 1. The overview of the proposed model. (a) is the training process of the diffusion model. (b) represents the generating process guided with class-wise prior FBN.

2.2 Training Process of Diffusion Model

As shown in Fig. 1(a), the diffusion model is a probabilistic model that involves a forward process and a reverse process during the training phase.

Forward Diffusion Process: The forward process of the diffusion model progressively adds noise to the input data. Specifically, for the input time series denoted $\boldsymbol{X}_0$, the Gaussian noise is successively added over T steps, obtaining $\boldsymbol{X}_1$, $\boldsymbol{X}_2$, ..., $\boldsymbol{X}_T$, where $\boldsymbol{X}_t$ denotes the noisy data at the t-th step, and the final $\boldsymbol{X}_T$ is the pure noisy time series. Mathematically, the process from $\boldsymbol{X}_{t-1}$ to $\boldsymbol{X}_t$ satisfies the following transition distribution:

$$q\left(\boldsymbol{X}_t \mid \boldsymbol{X}_{t-1}\right)=\mathcal{N}\left(\boldsymbol{X}_t ; \sqrt{1-\beta_t} \boldsymbol{X}_{t-1}, \beta_t \mathbf{I}\right), \tag{1}$$

where β_t denotes the noise variance that controls the amount of noise introduced at each step and $\mathbf{I}$ is the identity matrix. This indicates that $\boldsymbol{X}_t$ follows a normal distribution with mean $\sqrt{1-\beta_t}\boldsymbol{X}_{t-1}$ and variance β_t $\mathbf{I}$. To generate $\boldsymbol{X}_t$ at a specific step t, the noise-adding process can be reparameterized as:

$$\boldsymbol{X}_t=\sqrt{1-\bar{\beta}_t}\boldsymbol{X}_0+\sqrt{\bar{\beta}_t}\epsilon_t, \epsilon_t \sim \mathcal{N}(0, \mathbf{I}), \tag{2}$$

where ϵ_t is sampled from the standard normal distribution $\mathcal{N}(0, \mathbf{I})$, and $\bar{\beta}_t$ represents the cumulative noise variance.

Reverse Diffusion Process: The reverse process of the diffusion model aims to denoise the final noisy data $\boldsymbol{X}_T$ and gradually recover the original data $\boldsymbol{X}_0$ from $\boldsymbol{X}_T$. Specifically, it learns the joint probability distribution from $\boldsymbol{X}_T$ to $\boldsymbol{X}_0$, which can be formulated as:

$$p_\theta(\boldsymbol{X}_{0:T})=p(\boldsymbol{X}_T)\prod_{t=1}^{T} p_\theta(\boldsymbol{X}_{t-1}|\boldsymbol{X}_t), \tag{3}$$

where the reverse conditional probability distribution from $\boldsymbol{X}_t$ to $\boldsymbol{X}_{t-1}$ can be expressed as:

$$p_\theta(\boldsymbol{X}_{t-1}|\boldsymbol{X}_t) = \mathcal{N}(\boldsymbol{X}_{t-1}; \mu_\theta(\boldsymbol{X}_t, t), \Sigma_\theta(\boldsymbol{X}_t, t)), \tag{4}$$

where the mean $\mu_\theta(\boldsymbol{X}_t, t)$ and covariance $\Sigma_\theta(\boldsymbol{X}_t, t)$ can be predicted from a noise prediction network, e.g., the Attention Residual U-Net [23], and θ represents the parameters of the network. In this case, to make the predicted noise $\epsilon_\theta(\boldsymbol{X}_t, t)$ more consistent with the added noise ϵ_t of the forward diffusion process, the loss function L can be defined as:

$$L = \mathbb{E}_{\boldsymbol{X}_0, \epsilon_t \sim \mathcal{N}(0,\mathbf{I}), t}[\|\epsilon_t - \epsilon_\theta(\boldsymbol{X}_t, t)\|_2^2]. \tag{5}$$

2.3 Generating process guided with prior FBN

Through the above training process, the model can capture the underlying distribution of fMRI data and learn how to progressively recover the original data from pure noise. However, the traditional generating process fails to account for the inherent functional dependencies among the time series of different brain regions, which are crucial for generating accurate fMRI signals. To this end, we propose incorporating a prior FBN to guide the data generation process, which contains rich functional connectivity information between brain regions. Specifically, as shown in Fig. 1(b), the generated data $\boldsymbol{X}_t$ at a specific step is guided to exhibit similar functional connectivity information as in the prior FBN via the following mechanism:

$$\boldsymbol{X}_t^{(\text{new})} = \boldsymbol{X}_t - \lambda \frac{\partial \| \boldsymbol{C}_t - \boldsymbol{C}_p \|_F}{\partial \boldsymbol{C}_t} \cdot \frac{\partial \boldsymbol{C}_t}{\partial \boldsymbol{X}_t}, \tag{6}$$

where $\boldsymbol{C}_t$ denotes the FBN calculated from $\boldsymbol{X}_t$ using the Pearson correlation coefficient. $\boldsymbol{C}_p$ represents a prior FBN obtained by averaging the FBNs of actual fMRI samples, $\| \cdot \|_F$ indicates the Frobenius norm used to quantify the discrepancy between $\boldsymbol{C}_t$ and $\boldsymbol{C}_p$, λ is the update coefficient, and this FBN guidance is conducted every Δt step. Furthermore, to flexibly generate fMRI data from different classes, e.g., normal control (NC) and AD patients, the proposed method introduces the label information into the generation process. Specifically, by defining the prior FBN as the representative class-wise FBN derived from actual samples of a specific class, it enables to adaptively generate the fMRI data for the corresponding class via Eq. (6). By doing this, the generated fMRI signals are more consistent with the actual fMRI signals. Meanwhile, the generated fMRI is aligned with the downstream disease classification tasks.

3 Experiment

3.1 Data and Preprocessing

The proposed method is evaluated on various challenging brain disorder classification tasks on three fMRI datasets, including the public Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset (https://adni.loni.usc.edu/) and

two private datasets from local hospitals named Huashan-MCI and Zhongshan-VCI for differentiating MCI and VCI from NC subjects, respectively. There are 407 subjects (183 MCI vs. 224 NC) on ADNI dataset, 513 subjects (204 MCI vs. 309 NC) on Huashan-MCI dataset and 297 subjects (151 VCI vs. 146 NC) on Zhongshan-VCI dataset. The fMRI data from ADNI, Huashan-MCI and Zhongshan-VCI were acquired with TR/TE set as 3000/30 ms, 800/37 ms and 2000/24 ms, respectively. All data are preprocessed with the DPARSFA toolbox [30] following [7]. The mean time series of 90 brain regions in gray matter are extracted from the preprocessed data using the automated anatomical labeling (AAL) atlas [28].

3.2 Competing Methods

To evaluate the effectiveness of the proposed method, various state-of-the-art competing methods with typical architectures are included for comparison, including one VAE-based method (VAE [12]), one GAN-based method (LSTM-GAN [33]), and four DM-based methods (DDPM [10], LDM [23], Diffusion-TS [31], BrainNetDiffusion [32]). Besides, BolT [1] and GCN [18] are used as classifiers to comprehensively validate the performance of the generated fMRI.

3.3 Experimental Settings

For input $\boldsymbol{X}_0 \in \mathbb{R}^{N \times L}$, the number of brain regions N is set as 90 and the length of time series L is set as 128, 480, 208 for ADNI, Huashan-MCI, Zhongshan-VCI datasets, respectively. In the training process, the batch size is set as 24, the epoch is set as 1000, the learning rate is set as 0.0001, and the Adam optimizer [13] is used. The number of noise-adding steps T is set as 200 and β_t is linearly sampled from [0.0015, 0.0195]. During the generating process, the update coefficient λ is set as 1 and the FBN guidance interval Δt is also set as 1, which are selected by cross-validation in the training dataset. We generate 100 samples for ADNI and Huashan-MCI datasets (50 MCI vs. 50 NC) and 50 for Zhongshan-VCI (25 VCI vs. 25 NC). To ensure a fair comparison, we perform 5-fold cross-validation and use the same data partitions across all competing methods. The performance is evaluated using four metrics, including accuracy (ACC), area under the receiver operating characteristic curve (AUC), sensitivity (SEN) and specificity (SPE).

3.4 Experimental Results

The performance of the proposed method and the competing methods on three datasets is summarized in Table 1. On the ADNI dataset, the proposed method achieves accuracies of 71.7% and 74.5% with the BoIT and GCN classifiers, respectively, outperforming the second-best method by 2.2% and 3.2%. For Huashan-MCI, the accuracies are 70.2% and 70.8% with BoIT and GCN, showing improvements of 2.8% for both classifiers compared to the second-best method.

Table 1. The performance of different generation methods on ADNI, Huashan-MCI and Zhongshan-VCI.

Dataset	Method	Classifier	ACC	AUC	SEN	SPE
ADNI	Original data	BoIT	66.8±3.8	66.2±2.5	63.4±2.9	69.6±4.9
	VAE[12]		69.5±1.9	70.9±2.3	61.9±4.9	76.0±6.1
	LSTM-GAN[33]		67.8±1.5	66.8±3.9	60.7±2.3	73.7±3.6
	Diffusion-TS[31]		67.1±1.8	66.4±1.9	62.4±10.6	71.0±8.2
	LDM[23]		67.6±2.3	67.0±3.7	63.8±6.2	70.4±8.0
	DDPM[10]		67.3±2.9	66.7±3.1	57.7±9.0	75.5±5.1
	TA-fMRI-GM(proposed)		**71.7**±1.5	**72.6**±2.6	64.1±7.8	78.2±6.4
	Original data	GCN	69.0±4.7	65.2±7.9	53.6±14.5	81.1±9.4
	VAE[12]		69.6±4.0	66.2±5.9	44.8±8.7	89.7±3.5
	LSTM-GAN[33]		70.5±1.2	68.0±2.9	58.2±10.8	80.2±10.1
	Diffusion-TS[31]		69.3±3.1	68.4±2.9	55.4±14.0	80.2±6.3
	BrainNetDiffusion[32]		70.3±6.1	67.9±8.3	49.6±18.0	86.5±9.0
	LDM[23]		69.8±3.4	65.3±6.4	53.2±16.0	82.6±6.9
	DDPM[10]		71.3±3.9	69.5±5.1	59.3±13.0	80.7±5.6
	TA-fMRI-GM(proposed)		**74.5**±2.7	**74.5**±4.2	63.5±12.4	83.5±6.7
Huashan-MCI	Original data	BoIT	65.5±5.5	63.0±7.7	40.9±18.3	81.7±10.0
	VAE[12]		67.4±4.8	64.1±8.3	47.2±5.1	80.8±6.1
	LSTM-GAN[33]		66.7±2.1	62.2±5.8	49.2±13.8	78.2±10.7
	Diffusion-TS[31]		66.9±5.6	59.5±9.9	43.3±17.5	82.4±8.9
	LDM[23]		66.7±2.2	60.2±6.7	43.8±6.5	81.8±6.5
	DDPM[10]		66.7±3.0	62.6±7.6	50.7±15.6	77.2±11.3
	TA-fMRI-GM(proposed)		**70.2**±6.8	**70.7**±9.7	57.0±11.1	78.8±5.5
	Original data	GCN	64.1±2.7	55.6±4.3	32.0±2.9	85.4±3.8
	VAE[12]		66.1±1.3	56.4±4.0	27.5±9.9	91.6±8.3
	LSTM-GAN[33]		67.6±1.9	61.1±5.4	35.3±14.7	89.0±8.5
	Diffusion-TS[31]		65.9±1.8	58.1±1.0	41.4±9.3	82.1±9.2
	BrainNetDiffusion[32]		68.0±2.9	62.0±4.2	45.3±10.5	83.0±4.1
	LDM[23]		66.9±2.7	58.2±8.2	36.9±13.6	86.6±7.6
	DDPM[10]		66.9±1.5	62.3±1.8	40.4±8.3	84.4±6.8
	TA-fMRI-GM(proposed)		**70.8**±3.0	**68.8**±3.8	56.1±8.6	80.4±6.2
Zhongshan-VCI	Original data	BoIT	63.6±9.3	57.3±10.8	66.9±13.4	60.2±11.0
	VAE[12]		63.6±5.6	60.3±9.2	72.9±6.2	54.0±16.8
	LSTM-GAN[33]		61.9±5.8	57.4±7.1	80.2±7.4	43.1±14.1
	Diffusion-TS[31]		61.6±3.6	55.8±2.6	54.3±18.9	69.3±16.4
	LDM[23]		61.0±5.6	58.0±7.7	65.6±7.5	56.2±14.5
	DDPM[10]		64.3±5.4	59.7±9.1	63.6±9.6	65.0±7.3
	TA-fMRI-GM(proposed)		**66.7**±5.4	**64.8**±7.0	69.5±10.0	63.7±12.2
	Original data	GCN	64.0±4.4	62.1±6.3	59.6±13.5	68.4±14.2
	VAE[12]		66.0±3.7	66.2±5.5	64.9±8.9	67.1±6.5
	LSTM-GAN[33]		65.6±4.3	60.8±6.7	67.4±16.7	63.8±19.3
	Diffusion-TS[31]		65.0±3.4	65.8±4.9	62.2±8.9	67.8±12.9
	BrainNetDiffusion[32]		66.3±3.8	63.2±7.6	68.9±12.7	63.7±14.7
	LDM[23]		64.0±4.1	61.6±6.2	60.9±18.8	67.1±17.3
	DDPM[10]		67.4±7.2	66.6±12.9	68.2±7.1	66.4±20.5
	TA-fMRI-GM(proposed)		**70.0**±3.7	**67.6**±5.8	74.2±8.7	65.1±6.1

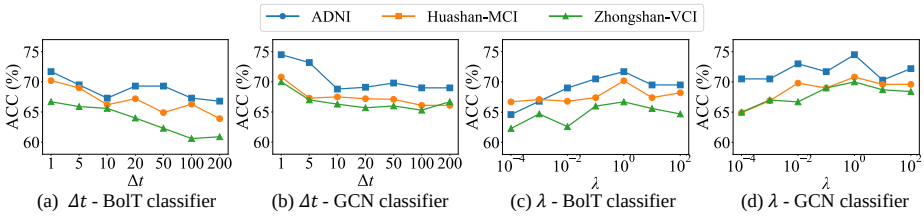

Fig. 2. Performance of the proposed method with different numbers of (a)(b) FBN guidance intervals Δt and (c)(d) update coefficients λ.

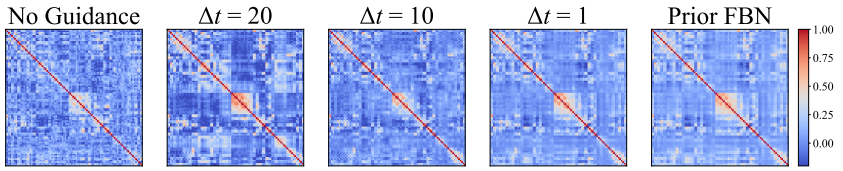

Fig. 3. Illustrations of the FBNs derived from the generated fMRI without prior FBN guidance, as well as those generated under different FBN guidance intervals Δt, and the prior FBN from the actual fMRI on the Huashan-MCI dataset.

Likewise, on Zhongshan-VCI, the proposed method achieves accuracies of 66.7% and 70.0% with BoIT and GCN, outperforming the second-best method by 2.4% and 2.6%. Moreover, the proposed method surpasses the base model DDPM by 3.5% and 3.9% on the Huashan-MCI dataset, and similar improvements are observed on the other two datasets. These results consistently indicate the effectiveness and reliability of the proposed method in generating fMRI data, demonstrating that it can be well aligned with the brain disorder classification tasks.

3.5 Ablation Study

An ablation study is conducted to evaluate the effect of different FBN guidance intervals Δt and update coefficients λ during incorporating prior FBN guidance. As shown in Fig. 2(a) and (b), the performance of the proposed model improves as Δt decreases. This suggests that more frequent integration and interaction between the denoising process and the guidance of prior FBN during the generation process can improve the generated fMRI towards the actual fMRI data. Besides, as seen in Fig. 2(c) and (d), the performance of the proposed model initially increases and then decreases as λ varies from 10^{-4} to 10^{2}, with the best performance when λ is 1. These results imply that adding FBN guidance is beneficial in a wide range of λ. At the same time, excessively low λ may introduce insufficient guidance and excessively high λ may affect the diversity of the generated fMRI, both of which ultimately lead to degraded performance.

4 Visualization and Conclusion

As shown in Fig. 3, the FBNs generated with FBN guidance exhibit connectivity characteristics more similar to the actual FBN compared to those generated without guidance. Besides, as the FBN guidance interval Δt decreases, the resulting FBNs become more closely aligned with the actual FBN. These results further emphasize the importance of incorporating prior FBN guidance to enhance the quality of the generated fMRI data and suggest that more frequent guidance can improve the generated fMRI towards the actual fMRI.

In conclusion, our paper proposes a task-aligned generation method to generate fMRI time series with prior FBN guidance, which is more flexible and efficient to generate fMRI that conforms to the actual fMRI and is well aligned to the brain disorder diagnosis tasks. Experimental results on three datasets demonstrate the effectiveness of the proposed method.

Acknowledgments. This work was supported in part by Shenzhen Medical Research Fund under Grant B2402030; in part by the National Natural Science Foundation of China under Grant 62471501, Grant 62471502; in part by Guangdong Basic and Applied Basic Research Foundation under Grant 2023A1515012278; in part by Shenzhen Science and Technology Program under Grant JCYJ20240813151026034 and JCYJ20220818102414031; in part by the Open Research Fund of the State Key Laboratory of Brain-Machine Intelligence, Zhejiang University, under Grant BMI2400019.

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

References

1. Bedel, H.A., Sivgin, I., Dalmaz, O., Dar, S.U., Çukur, T.: BolT: Fused window transformers for fMRI time series analysis. Medical image analysis **88**, 102841 (2023)
2. Cao, H., Tan, C., Gao, Z., Xu, Y., Chen, G., Heng, P.A., Li, S.Z.: A survey on generative diffusion models. IEEE Transactions on Knowledge and Data Engineering (2024)
3. Chen, D., Zhang, L.: FE-STGNN: Spatio-temporal graph neural network with functional and effective connectivity fusion for MCI diagnosis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 67–76. Springer (2023)
4. Dennis, E.L., Thompson, P.M.: Functional brain connectivity using fMRI in aging and Alzheimer's disease. Neuropsychology review **24**, 49–62 (2014)
5. Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. Advances in neural information processing systems **34**, 8780–8794 (2021)
6. Du, Y., Fu, Z., Calhoun, V.D.: Classification and prediction of brain disorders using functional connectivity: promising but challenging. Frontiers in neuroscience **12**, 525 (2018)
7. Esteban, O., Markiewicz, C.J., Blair, R.W., Moodie, C.A., Isik, A.I., Erramuzpe, A., Kent, J.D., Goncalves, M., DuPre, E., Snyder, M., et al.: fMRIPrep: a robust preprocessing pipeline for functional MRI. Nature methods **16**(1), 111–116 (2019)

8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. Communications of the ACM **63**(11), 139–144 (2020)
9. Guo, Z., Mao, J., Tao, R., Yan, L., Ouchi, K., Liu, H., Wang, X.: Audio generation with multiple conditional diffusion model. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 18153–18161 (2024)
10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
11. Khader, F., Müller-Franzes, G., Tayebi Arasteh, S., Han, T., Haarburger, C., Schulze-Hagen, M., Schad, P., Engelhardt, S., Baeßler, B., Foersch, S., et al.: Denoising diffusion probabilistic models for 3D medical image generation. Scientific Reports **13**(1), 7303 (2023)
12. Kingma, D.P.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
14. Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: Diffwave: A versatile diffusion model for audio synthesis. In: International Conference on Learning Representations (2021)
15. Li, H., Yu, S., Principe, J.: Causal recurrent variational autoencoder for medical time series generation. In: Proceedings of the AAAI conference on artificial intelligence. vol. 37, pp. 8562–8570 (2023)
16. Li, T., Wu, C., Shi, P., Wang, X.: Cumulative difference learning vae for time-series with temporally correlated inflow-outflow. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 13619–13627 (2024)
17. Li, X., Metsis, V., Wang, H., Ngu, A.H.H.: TTS-GAN: A Transformer-based time-series generative adversarial network. In: International conference on artificial intelligence in medicine. pp. 133–143. Springer (2022)
18. Liu, M., Wang, Y., Zhang, H., Yang, Q., Shi, F., Zhou, Y., Shen, D.: Multiscale functional connectome abnormality predicts cognitive outcomes in subcortical ischemic vascular disease. Cerebral Cortex **32**(21), 4641–4656 (2022)
19. Lu, C., Reddy, C.K., Wang, P., Nie, D., Ning, Y.: Multi-label clinical time-series generation via conditional GAN. IEEE Transactions on Knowledge and Data Engineering (2023)
20. Pesteie, M., Abolmaesumi, P., Rohling, R.N.: Adaptive augmentation of medical data using independently conditional variational auto-encoders. IEEE transactions on medical imaging **38**(12), 2807–2820 (2019)
21. Qiang, N., Dong, Q., Liang, H., Ge, B., Zhang, S., Sun, Y., Zhang, C., Zhang, W., Gao, J., Liu, T.: Modeling and augmenting of fMRI data using deep recurrent variational auto-encoder. Journal of neural engineering **18**(4), 0460b6 (2021)
22. Qiang, N., Gao, J., Dong, Q., Yue, H., Liang, H., Liu, L., Yu, J., Hu, J., Zhang, S., Ge, B., et al.: Functional brain network identification and fMRI augmentation using a VAE-GAN framework. Computers in Biology and Medicine **165**, 107395 (2023)
23. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
24. Sattarov, T., Schreyer, M., Borth, D.: FinDiff: Diffusion models for financial tabular data generation. In: Proceedings of the Fourth ACM International Conference on AI in Finance. pp. 64–72 (2023)

25. Singh, N.K., Raza, K.: Medical image generation using generative adversarial networks: A review. Health informatics: A computational perspective in healthcare pp. 77–96 (2021)
26. Suh, N., Yang, Y., Hsieh, D.Y., Luan, Q., Xu, S., Zhu, S., Cheng, G.: TimeAutoDiff: Combining autoencoder and diffusion model for time series tabular data synthesizing. arXiv preprint arXiv:2406.16028 (2024)
27. Supekar, K., Menon, V., Rubin, D., Musen, M., Greicius, M.D.: Network analysis of intrinsic functional brain connectivity in alzheimer's disease. PLoS computational biology **4**(6), e1000100 (2008)
28. Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M.: 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)
29. Vetter, J., Macke, J.H., Gao, R.: Generating realistic neurophysiological time series with denoising diffusion probabilistic models. Patterns **5**(9) (2024)
30. Yan, C., Zang, Y.: DPARSF: a MATLAB toolbox for "pipeline" data analysis of resting-state fMRI. Frontiers in systems neuroscience **4**, 1377 (2010)
31. Yuan, X., Qiao, Y.: Diffusion-TS: Interpretable diffusion for general time series generation. In: International Conference on Learning Representations (2024)
32. Zhao, H., Lou, H., Yao, L., Zhang, Y.: Diffusion transformer-augmented fMRI functional connectivity for enhanced autism spectrum disorder diagnosis. Journal of Neural Engineering (2025)
33. Zhu, G., Zhao, H., Liu, H., Sun, H.: A novel LSTM-GAN algorithm for time series anomaly detection. In: 2019 prognostics and system health management conference (PHM-Qingdao). pp. 1–6. IEEE (2019)