

Hierarchical Characterization of Brain Dynamics via State Space-based Vector Quantization

Yanwu Yang¹ and Thomas Wolfers^{1,2}(✉)

¹ University of Tübingen, Tübingen, Germany

² German Center for Mental Health, Tübingen, Germany
yangyanwu1111@gmail.com dr.thomas.wolfers@gmail.com

Abstract. Understanding brain dynamics through functional Magnetic Resonance Imaging (fMRI) remains a fundamental challenge in neuroscience, particularly in capturing how the brain transitions between various functional states. Recently, metastability, which refers to temporarily stable brain states, has offered a promising paradigm to quantify complex brain signals into interpretable, discretized representations. In particular, compared to cluster-based machine learning approaches, tokenization approaches leveraging vector quantization have shown promise in representation learning with powerful reconstruction and predictive capabilities. However, most existing methods ignore brain transition dependencies and lack a quantification of brain dynamics into representative and stable embeddings. In this study, we propose a Hierarchical State space-based Tokenization network, termed HST, which quantizes brain states and transitions in a hierarchical structure based on a state space-based model. We introduce a refined clustered Vector-Quantization Variational AutoEncoder (VQ-VAE) that incorporates quantization error feedback and clustering to improve quantization performance while facilitating metastability with representative and stable token representations. We validate our HST on two public fMRI datasets, demonstrating its effectiveness in quantifying the hierarchical dynamics of the brain and its potential in disease diagnosis and reconstruction performance. Our method offers a promising framework for the characterization of brain dynamics, facilitating the analysis of metastability.

Keywords: fMRI · brain dynamics · tokenization · vector quantization.

1 Introduction

In recent years, functional magnetic resonance imaging (fMRI), which uses the blood oxygen level-dependent (BOLD) effect, has emerged as a powerful tool to study temporal brain dynamics in cognitive processes and characterize dysfunction in patients [11, 28, 30]. Despite advances in modeling paradigms [20, 25, 26], understanding these dynamic patterns remains a challenge in neuroscience, as the human brain exhibits dynamic high-dimensional, non-linear, and rapidly evolving dynamics that continuously shift between distinct functional states in

response to varying cognitive demands [4, 6]. Recent studies on brain metastability have provided a new perspective for understanding these complex dynamics [16]. Metastability refers to temporarily stable brain states that persist before transitioning to different configurations, offering a paradigm to quantify continuous brain signals into discretized, and interpretable representations. This approach not only captures the inherent temporal discontinuities of neural activity but also provides a principled way to quantize continuous brain dynamics into finite meaningful states [16].

However, challenges remain in quantifying brain dynamics through the lens of metastability. Traditional approaches, such as LEiDA [5, 13], rely on simple clustering methods to identify patterns of recurrent brain function. However, their linearity and reliance on statistical clustering may oversimplify brain dynamics by ignoring nonlinear interactions, failing to capture rapid state transitions, and reducing high-dimensional data to overly simplistic partitions. Recent advances in deep learning have introduced tokenization or quantization approaches as a promising solution [19, 22]. These methods map inputs into structured representations such as discrete tokens. Vector Quantized Variational Autoencoders (VQ-VAE) [22] have demonstrated remarkable success in various domains, enabling probabilistic modeling and token prediction for pre-training [1, 2, 7]. Unlike traditional VAEs that learn continuous latent spaces, VQ-based methods map features to a finite codebook of embeddings, producing more compact and interpretable representations. Several variants have been proposed to improve the reconstruction and quantization performance, including VQ-VAE-2, Online Cluster VQ-VAE, and RQ-VAE [14, 17, 32].

Despite advances in vector quantization, most existing approaches on brain dynamics potentially ignore brain state transitions in neural dynamics, which underlie the mechanisms of cognitive flexibility and dynamic patterns of neural activity between different cognitive states or processing modes. Studies have shown that abnormal transition patterns are associated with various neurological and psychiatric conditions, making them crucial for understanding brain function [8, 15]. Moreover, another significant challenge lies in the trade-off between quantization precision and stability [31, 32]. Increasing the size of the codebook tends to improve reconstruction performance, but it can lead to unstable and fragmented token representations that are sensitive to noise in neural signals [17, 31]. This instability makes it difficult to characterize the metastability in brain dynamics under different experimental conditions or between individuals.

In this paper, we propose to investigate the characterization of brain dynamics with a hierarchical structure that captures both brain states and transitions simultaneously. Moreover, to improve quantization stability and reconstruction performance, we introduce a refined cluster-based VQ-VAE by optimizing codebook utilization and clustering effectiveness. Our key contributions are:

- (1) We propose a state space-based tokenization method for the quantization of time series fMRI data that enables hierarchical mapping of brain states and transitions.




Fig. 1: The architecture of our HST, including a spatio-temporal encoder, and a state space-based model with a hierarchical VQ-VAE for quantization. For downstream tasks, the quantization embeddings are concatenated and fed into a multi-layer perceptron for classification.

(2) We introduce a refined cluster VQ-VAE that incorporates quantization error feedback and clustering to improve quantization performance while facilitating metastability with stable and representative discretized embeddings.

(3) Evaluations on two public datasets for disease diagnosis and reconstruction performance highlight the superiority of our method in brain dynamics quantization and its potential for analyzing clinical brain states.

2 Method

The framework of our proposed HST is shown in Fig. 1. The inputs are first encoded by a spatio-temporal encoder and then fed into a hierarchical vector quantization based on brain state and transition representations. The quantized embeddings are concatenated and decoded for reconstruction.

2.1 Spatio-temporal Encoder

For a given input BOLD signal $X \in \mathbb{R}^{T \times M}$, the spatio-temporal encoder encodes the sequence from both temporal and spatial perspectives and then fuses them attentively. The temporal and spatial encoding modules are implemented using transformer networks as: $H^t = \text{Enc}(X^{\text{temporal}})$, where $X^{\text{temporal}} = X_{1,:}, X_{2,:}, \dots, X_{T,:}$, and $H^s = \text{Enc}(X^{\text{spatial}})$, where $X^{\text{spatial}} = X_{:,1}, X_{:,2}, \dots, X_{:,M}$. In particular, X^{temporal} and X^{spatial} are obtained by sequencing the BOLD signal along the temporal and spatial dimensions. For each encoder (Enc),

$H = \text{MHSA}(X)$, based on the query, key, and value matrices $Q_t^{l,c}, K_t^{l,c}, V_t^{l,c}$ and $Q_s^{l,c}, K_s^{l,c}, V_s^{l,c}$. The temporal and spatial attention scores are obtained by squeezing and exciting the corresponding representations: $A^t = \sigma(w_2^s \cdot \delta(w_1^s \cdot \frac{1}{M} \sum_{j=1}^M H_{:,j}^t + b_1^s) + b_2^s)$, where $A^t \in \mathbb{R}^{T \times 1}$, and $A^s = \sigma(w_2^t \cdot \delta(w_1^t \cdot \frac{1}{T} \sum_{i=1}^T H_{i,:}^t + b_1^t) + b_2^t)$, where $A^s \in \mathbb{R}^{1 \times M}$. Here, $w_1^t, w_2^t, w_1^s, w_2^s$ and $b_1^t, b_2^t, b_1^s, b_2^s$ are learnable parameters for attention projection. Finally, cross-matrix multiplication is used for fusion of spatial and temporal representations: $H^f = H^t \otimes A^s + H^s \otimes A^t$.

2.2 State Space-based Modeling Encoder

The cross-fusion spatio-temporal representations are split into sequences $H^f = H_1^f, H_2^f, \dots, H_T^f$ along the temporal dimension. These sequences are then fed into the state space-based modeling encoder for encoding brain dynamic states and transitions between states. In general, this is obtained as:

$$h_t = f(h_{t-1}, x_t, \dots; \Theta), \quad o_t = g(h_t, x_t, \dots; \Phi) \quad (1)$$

Here, f and g denote the mapping functions for transition h_t and state o_t updates at the t -th step, with Θ and Φ as their corresponding parameters. In particular, for each update, the t -th transition and state representations depend on previous representations h_{t-1} , or h_{t-2} , varying with the dependency range of different methods. In our study, we examine this using various approaches: RNN updating on the previous step, GRU and LSTM with attentive forgetting and updating, and Mamba with selective state updates. For RNN, the state updates as follow:

$$h_t = \sigma(W_h h_{t-1} + W_x H_t + b_h) \quad (2)$$

For Mamba [10], the states are obtained by:

$$\bar{\mathbf{A}} = \exp(\Delta \mathbf{A}), \quad \bar{\mathbf{B}} = (\Delta \mathbf{A})^{-1}(\exp(\Delta \mathbf{A}) - \mathbf{I}) \cdot \Delta \mathbf{B}, \quad h_t = \bar{\mathbf{A}} h_{t-1} + \bar{\mathbf{B}} x_t \quad (3)$$

where $\bar{\mathbf{A}}$ and $\bar{\mathbf{B}}$ are discrete equivalents of the corresponding continuous evolution matrix $\mathbf{A}$ and projection matrix $\mathbf{B}$. $\Delta \mathbf{A}$ and $\Delta \mathbf{B}$ are learnable parameters representing the discretized versions of the continuous state space-based model, where $\Delta \mathbf{A}$ captures the state transitions over time steps. More details can be found in the previous paper [10]. We formulate the state output representations using a two-layer MLP as: $o_t = W_{o2}(W_{o1}(h_{t-1} \parallel x_t) + b_{o1}) + b_{o2}$.

2.3 Refined Cluster-based VQ-VAE

While deep learning models often benefit from larger codebooks, metastability requires fewer and more meaningful clusters. We address this by proposing a refined cluster VQ-VAE that optimizes codebook efficiency, achieving high reconstruction quality with a compact, interpretable codebook. Our model employs a hierarchical architecture with error feedback and dynamic clustering. The first level implements separate quantization for brain states and transitions:

$$\hat{o}_t = e_k^S(k = \arg\min_i \|o_t - e_i^S\|_2), \quad \hat{h}_t = e_k^T(k = \arg\min_j \|h_t - e_j^T\|_2) \quad (4)$$

where e^S and e^T denote the embeddings of brain state and transition codebooks, and $\hat{o}$ and $\hat{h}$ are the corresponding representations after quantization. Secondly, an auxiliary quantization module computes and feeds back errors to optimize the codebooks:

$$\hat{e}_{o,t} = e_{o,k}^S (k = \operatorname{argmin}_p \|(o_t - \hat{o}_t) - e_p^S\|_2), \quad (5)$$

$$\hat{e}_{h,t} = e_{h,k}^S (k = \operatorname{argmin}_q \|(h_t - \hat{h}_t) - e_q^T\|_2) \quad (6)$$

where e_o^S and e_h^T are the corresponding error feedback codes. Finally, the clustering quantization dynamically updates underutilized code vectors by repositioning them closer to active feature regions:

$$\mathbf{e}_j^{(t+1)} = \mathbf{e}_j^{(t)} \cdot (1 - \alpha_j) + \hat{\mathbf{z}}_j \cdot \alpha_j \quad (7)$$

where $\alpha_j = \exp(-N_j \cdot K \cdot \frac{10}{1-\gamma})$, and $N_j^{(t+1)} = \gamma N_j^{(t)} + (1 - \gamma)n_j^{(t)}$ [32]. Here, $\hat{\mathbf{z}}_j$ represents the feature closest to the j -th codebook vector.

2.4 Decoder and Optimization

After quantifying the spatio-temporal dynamic representations of the brain into state and transition tokens, we leverage a Transformer network to decode the latent embeddings into BOLD signals. The codebook embeddings of brain states and transitions are concatenated and then fed into the decoder. Finally, the loss function is to optimize the reconstruction more closer to the input, as well as a weighted combination of codebook loss with α and β . We use $\operatorname{sg}(\cdot)$ as a gradient stop operation to control the flow of gradients:

$$\begin{aligned} \mathcal{L} = & \alpha \|x - \hat{x}\|_2^2 + \beta \|\operatorname{sg}(o_t) - \hat{o}_t\|_2^2 + \gamma \|\operatorname{sg}(o_t - \hat{o}_t) - e_t\|_2^2 \\ & + \beta \|\operatorname{sg}(h_t) - \hat{h}_t\|_2^2 + \gamma \|\operatorname{sg}(h_t - \hat{h}_t) - e_t\|_2^2 \end{aligned} \quad (8)$$

3 Experiments and results

3.1 Dataset and experimental settings

Datasets. In this study, we use two fMRI datasets on different scales for evaluation. (1) **The ADHD-200 dataset** [21] is a publicly available collection of 1,260 individuals in multiple data centers. After excluding images with scan durations shorter than 100 TRs, our final sample includes 954 subjects: 405 patients with Attention Deficit Hyperactivity Disorder (ADHD) and 549 healthy controls (HC) with sex and age matching. (2) **The SchizoConnect dataset** [23] is a multisite neuroimaging resource providing MRI data in schizophrenia studies, comprising 97 patients with schizophrenia (SCZ) and 91 healthy controls (HC).

Pre-processing. All images were pre-processed using Configurable Pipeline for the Analysis of Connectomes (C-PAC) software, including anterior commissure-posterior commissure (AC-PC) alignment, skull stripping, bias field




Fig. 2: Reconstruction performance on the test set of ADHD-200 dataset in terms of A) correlation (r) and B) Mean Squared Error (MSE), and C) downstream classification accuracy (Acc) with different numbers of embeddings.

correction, and normalization into the standard Montreal Neurological Institute (MNI) space.

Implementations. The transformer networks in the spatio-temporal encoder and decoder are implemented with 2 layers and 4 heads. The spatial state modeling networks are analyzed in the results section, including RNN, LSTM, GRU, and Mamba, where the hidden feature size is set to 256 and the number of layers is 2. The number of tokens in brain state and transition quantization is searched within the range [8, 16, 32, 64, 128]. For hyperparameters, β is set to 0.1, and α is set to 1.0. The learning rate is set to 0.0002. All models are trained using the Adam optimizer. The training process consists of two phases. The tokenizer quantization training is conducted for 4,000 steps, followed by 200 epochs of downstream evaluation. During the second phase, the quantizers are frozen and only the encoder and a 3-layer MLP classifier are trained for classification.

Table 1: Diagnosis performance based on different state-space modeling methods.

	ADHD-200			SchizoConnect		
	Acc	Sen	Spe	Acc	Sen	Spe
RNN	63.81±2.98	61.81±5.42	64.91±1.74	71.57±2.52	68.45±3.57	77.94±3.08
LSTM	65.03±3.01	64.10±3.77	65.78±3.40	73.71±5.05	70.36±4.82	79.18±6.43
GRU	62.59±2.30	59.53±4.65	65.03±3.35	73.57±2.52	69.45±3.57	80.94±3.08
Mamba	66.28±2.54	66.68±3.73	66.18±3.87	75.72±2.15	75.00±3.00	83.75±4.25

3.2 Results

1. The effect of the number of quantization embeddings. We first evaluated the effect of embedding numbers on reconstruction and downstream classification. The state space-based model was implemented using Mamba [10].

We show the results of the ADHD-200 dataset in Figure 2. Figure 2 A) and B) demonstrate that our method achieved better reconstruction performance (higher correlation and lower MSE) with fewer embeddings, compared with other baselines. In particular, HST with 8 embeddings even outperformed both C-VAE and VQ-VAE using 16 embeddings. This demonstrates the efficiency of using a hierarchical quantization structure. For downstream classification, optimal performance was achieved using 4-16 brain state embeddings and transitions, as shown in Figure 2 C). Although increasing the size of the codebook improved reconstruction, the classification performance initially increased and then decreased. This reveals a trade-off between reconstruction fidelity and discriminative power. In general, using eight embeddings for quantifying brain states and transitions yields the best classification performance. Similarly, a grid search on the SchizoConnect dataset showed that 16 embeddings provided the best results.

Table 2: Classification results based on different sequential modeling methods. The best results are shown in bold.

	ADHD-200			SchizoConnect		
	Acc	Sen	Spe	Acc	Sen	Spe
SVM	57.16 $\pm$ 3.14	56.08 $\pm$ 4.64	57.84 $\pm$ 3.89	59.46 $\pm$ 2.96	57.39 $\pm$ 1.60	58.87 $\pm$ 3.89
MLP	59.60 $\pm$ 2.40	55.21 $\pm$ 5.41	59.46 $\pm$ 2.52	64.61 $\pm$ 5.05	68.86 $\pm$ 1.36	68.53 $\pm$ 8.27
ST-GCN [9]	60.73 $\pm$ 2.24	60.17 $\pm$ 1.88	59.63 $\pm$ 2.21	66.25 $\pm$ 1.54	65.87 $\pm$ 2.99	68.90 $\pm$ 4.51
FBNetGen [12]	61.41 $\pm$ 2.54	62.87 $\pm$ 2.00	60.08 $\pm$ 2.67	71.97 $\pm$ 1.54	70.96 $\pm$ 4.64	73.90 $\pm$ 2.63
BolT [3]	63.51 $\pm$ 2.55	59.28 $\pm$ 4.39	63.17 $\pm$ 3.12	71.22 $\pm$ 1.60	74.45 $\pm$ 2.52	70.60 $\pm$ 2.04
BrainWaveNet [11]	64.95 $\pm$ 2.92	65.39 $\pm$ 3.30	65.99 $\pm$ 5.44	73.75 $\pm$ 3.89	73.12 $\pm$ 3.91	75.61 $\pm$ 6.68
HST (Ours)	66.28$\pm$2.54	66.68$\pm$3.73	66.18$\pm$3.87	75.72$\pm$2.15	75.00$\pm$3.00	83.75$\pm$4.25

2. The effect of different SSM methods on diagnostic predictions.

We further evaluated four state space-based modeling methods (RNN, LSTM, GRU, and Mamba) for ADHD/SCZ diagnosis across both datasets. Table 1 presents the performance metrics of accuracy, sensitivity, and specificity across folds, reported as the mean and standard deviation, with the best results highlighted in bold. The results show that Mamba consistently outperformed the other approaches, highlighting the effectiveness of its selective long-range modeling capabilities. Although other methods achieved comparable performance in some cases, Mamba’s consistently superior performance across metrics and datasets underscores its robustness in sequential modeling tasks.

3. Comparisons with other baselines on diagnostic predictions.

Furthermore, we compared our HST with other sequential modeling methods for disease diagnosis. For downstream classification tasks, we extracted the hidden embeddings from the layer preceding the decoder and passed them through an MLP classifier. During this phase, we kept the codebook parameters fixed while fine-tuning both the encoder and classifier components to optimize classification performance. Table 2 presents the results across folds, reported as the mean and standard deviation. Despite the inherent reconstruction loss associated with VQ-VAE approaches, our method achieved performance comparable to or




Fig. 3: Averaged brain activation maps of brain function states and brain transitions with corresponding occurrence probability in significant differences (p -value < 0.05) between ADHD and HC groups on the ADHD-200 dataset. The differences remain significant after FDR correction, except for transition state 3.

superior to established baselines. These representations, with appropriate fine-tuning, demonstrated better generalizability. This enhanced performance can be attributed to the similarity of our method to metric learning, where input features are systematically optimized to align with selected codebook vectors. The success of this approach provides empirical support for the feasibility of using metastability frameworks in brain dynamics analysis.

4. Brain state and transition quantization analysis. We conducted a statistical analysis of brain state and transition occurrence probabilities between groups on the ADHD-200 dataset. Figure 3 presents the results of comparisons between the groups, highlighting significant differences ($p < 0.05$). Our analysis revealed that patients exhibited a higher probability of occurrence in brain state 3 and transition 3, while showing a lower probability in brain state 4 and transition 4. In particular, brain state 4 is associated with reduced activation of brain function. This suggests that patients with ADHD are more likely to experience activated brain function states, aligning with the hyperactivity symptoms commonly observed in ADHD [24, 27]. Furthermore, significant differences were observed in other brain states and transitions, particularly in the prefrontal and parietal regions, consistent with previous studies on ADHD [18, 29]. These findings underscore the feasibility and superiority of our method in dynamic quantization, offering valuable insights into metastability and disease analysis.

4 Discussion and Conclusion

In this study, we proposed the HST for brain dynamics quantization, enabling hierarchical quantization of brain states and transitions. We enhanced the metastability quantization of the model by implementing a refined cluster VQ-VAE to improve codebook efficiency and reconstruction performance. Our evaluations

show that the HST, with hierarchical quantization, facilitates superior reconstruction and downstream diagnostic performance. Additionally, the HST has the potential to contribute to the development of a brain dynamics foundation model with self-supervised learning, particularly through next-token prediction.

Future work will focus on benchmarking our approach against a broader range of brain state quantization methods and scaling the framework toward large-scale brain dynamics modeling across diverse populations.

Acknowledgments. This work was supported by the German Research Foundation (DFG) Emmy Noether Program (513851350, TW) and the BMBF/DLR project FEDORA (01EQ2403G, TW). We acknowledge the affiliation with the International Max Planck Research School for Intelligent Systems (IMPRS-IS) and the computational resources provided by the de.NBI Cloud, part of the German Network for Bioinformatics Infrastructure (de.NBI)

Disclosure of Interests. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. Baevski, A., Zhou, Y., Mohamed, A., Auli, M.: wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems* **33**, 12449–12460 (2020)
2. Bao, H., Dong, L., Piao, S., Wei, F.: Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254* (2021)
3. 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)
4. Beim Graben, P., Jimenez-Marin, A., Diez, I., Cortes, J.M., Desroches, M., Rodriguez, S.: Metastable resting state brain dynamics. *Frontiers in Computational Neuroscience* **13**, 62 (2019)
5. Cabral, J., Vidaurre, D., Marques, P., Magalhães, R., Silva Moreira, P., Miguel Soares, J., Deco, G., Sousa, N., Kringelbach, M.L.: Cognitive performance in healthy older adults relates to spontaneous switching between states of functional connectivity during rest. *Scientific reports* **7**(1), 5135 (2017)
6. Cavanna, F., Vilas, M.G., Palmucci, M., Tagliazucchi, E.: Dynamic functional connectivity and brain metastability during altered states of consciousness. *Neuroimage* **180**, 383–395 (2018)
7. Devlin, J.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
8. Du, M., Zhang, L., Li, L., Ji, E., Han, X., Huang, G., Liang, Z., Shi, L., Yang, H., Zhang, Z.: Abnormal transitions of dynamic functional connectivity states in bipolar disorder: A whole-brain resting-state fmri study. *Journal of affective disorders* **289**, 7–15 (2021)
9. Gadgil, S., Zhao, Q., Pfefferbaum, A., Sullivan, E.V., Adeli, E., Pohl, K.M.: Spatio-temporal graph convolution for resting-state fmri analysis. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VII* 23. pp. 528–538. Springer (2020)

10. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752 (2023)
11. Jeong, A.Y., Heo, D.W., Kang, E., Suk, H.I.: Brainwavenet: Wavelet-based transformer for autism spectrum disorder diagnosis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 56–66. Springer (2024)
12. Kan, X., Cui, H., Lukemire, J., Guo, Y., Yang, C.: Fbnetgen: Task-aware gnn-based fmri analysis via functional brain network generation. In: International Conference on Medical Imaging with Deep Learning. pp. 618–637. PMLR (2022)
13. Kurtin, D.L., Scott, G., Hebron, H., Skeldon, A.C., Violante, I.R.: Task-based differences in brain state dynamics and their relation to cognitive ability. *NeuroImage* **271**, 119945 (2023)
14. Lee, D., Kim, C., Kim, S., Cho, M., Han, W.S.: Autoregressive image generation using residual quantization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11523–11532 (2022)
15. Liu, C., Xue, J., Cheng, X., Zhan, W., Xiong, X., Wang, B.: Tracking the brain state transition process of dynamic function connectivity based on resting state fmri. *Computational Intelligence and Neuroscience* **2019**(1), 9027803 (2019)
16. Rabinovich, M.I., Huerta, R., Varona, P., Afraimovich, V.S.: Transient cognitive dynamics, metastability, and decision making. *PLoS computational biology* **4**(5), e1000072 (2008)
17. Razavi, A., Van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems* **32** (2019)
18. Schneider, M.F., Krick, C.M., Retz, W., Henges, G., Retz-Junginger, P., Reith, W., Rösler, M.: Impairment of fronto-striatal and parietal cerebral networks correlates with attention deficit hyperactivity disorder (adhd) psychopathology in adults—a functional magnetic resonance imaging (fmri) study. *Psychiatry Research: Neuroimaging* **183**(1), 75–84 (2010)
19. Song, C., Hui, C., Lin, Q., Zhang, W., Li, S., Zhang, S., Zhu, H., Li, Z., Liu, S., Jiang, F., Li, X.: Lvpnet: A latent-variable-based prediction-driven end-to-end framework for lossless compression of medical images (2025), <https://arxiv.org/abs/2506.17983>
20. Song, C., Hui, C., Zhang, W., Zhu, H., Liu, S., Huang, H., Jiang, F.: Bpclip: A bottom-up image quality assessment from distortion to semantics based on clip (2025), <https://arxiv.org/abs/2506.17969>
21. ADHD-200 CONSORTIUM: The adhd-200 consortium: a model to advance the translational potential of neuroimaging in clinical neuroscience. *Frontiers in systems neuroscience* **6**, 62 (2012)
22. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. *Advances in neural information processing systems* **30** (2017)
23. Wang, L., Alpert, K.I., Calhoun, V.D., Cobia, D.J., Keator, D.B., King, M.D., Kogan, A., Landis, D., Tallis, M., Turner, M.D., et al.: Schizconnect: Mediating neuroimaging databases on schizophrenia and related disorders for large-scale integration. *Neuroimage* **124**, 1155–1167 (2016)
24. Weyandt, L., Swentosky, A., Gudmundsdottir, B.G.: Neuroimaging and adhd: fmri, pet, dti findings, and methodological limitations. *Developmental neuropsychology* **38**(4), 211–225 (2013)
25. Yang, Y., Chen, H., Hu, J., Guo, X., Ma, T.: Advancing brain imaging analysis step-by-step via progressive self-paced learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 58–68. Springer (2024)

26. Yang, Y., Guo, X., Ye, C., Xiang, Y., Ma, T.: Creg-kd: Model refinement via confidence regularized knowledge distillation for brain imaging. *Medical Image Analysis* **89**, 102916 (2023)
27. Yang, Y., Ye, C., Cai, G., Song, K., Zhang, J., Xiang, Y., Ma, T.: Hypercomplex graph neural network: Towards deep intersection of multi-modal brain networks. *IEEE Journal of Biomedical and Health Informatics* (2024)
28. Yang, Y., Ye, C., Guo, X., Wu, T., Xiang, Y., Ma, T.: Mapping multi-modal brain connectome for brain disorder diagnosis via cross-modal mutual learning <https://ieeexplore.ieee.org/abstract/document/10182318/>, publisher: IEEE
29. Yang, Y., Ye, C., Ma, T.: A deep connectome learning network using graph convolution for connectome-disease association study. *Neural Networks* **164**, 91–104 (2023)
30. Yang, Y., Ye, C., Su, G., Zhang, Z., Chang, Z., Chen, H., Chan, P., Yu, Y., Ma, T.: Brainmass: Advancing brain network analysis for diagnosis with large-scale self-supervised learning. *IEEE Transactions on Medical Imaging* (2024)
31. Zhao, W., Zou, Q., Shah, R., Liu, D.: Representation collapsing problems in vector quantization. arXiv preprint arXiv:2411.16550 (2024)
32. Zheng, C., Vedaldi, A.: Online clustered codebook. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 22798–22807 (2023)