Proceedings of Machine Learning Research 1–16

Improving brain disorder diagnosis with advanced brain function representation and Kolmogorov-Arnold Networks

Tyler Ward TBWA233@UKY.EDU
Abdullah-Al-Zubaer Imran AIMRAN@UKY.EDU
University of Kentucky, Lexington, KY, USA

Abstract

Quantifying functional connectivity (FC), a vital metric for the diagnosis of various brain disorders traditionally relies on the use of a pre-defined brain atlas. However, using such atlases can lead to issues regarding selection bias and lack of regard for specificity. Addressing this, we propose a *novel* transformer-based classification network (AFBR-KAN) with effective brain function representation, to aid in diagnosing autism spectrum disorder (ASD). AFBR-KAN leverages Kolmogorov-Arnold Network (KAN) blocks replacing traditional multi-layer perceptron (MLP) components. Thorough experimentation reveals the effectiveness of AFBR-KAN in improving the diagnosis of ASD under various configurations of the model architecture.

Keywords: Brain disorder diagnosis, brain function representation, classification, deep learning, Kolmogorov-Arnold Network, Transformer

1. Introduction

Diagnoses of autism spectrum disorder (ASD) are becoming increasingly prevalent across the world (Ge et al., 2024). As such, research into effective methods to improve the diagnosis of this brain disorder is vital. Traditional methods of diagnosing ASD have relied on analysis of functional connectivity (FC) in the brain, quantified from blood-oxygen-level-dependent (BOLD) signals obtained during resting-state functional magnetic resonance imaging (rs-fMRI), but this approach has several flaws.

FC analysis performed in this matter typically relies on regions-of-interest (ROIs) produced by registering a subject's brain with a pre-defined atlas. This approach can lead to subjective selection bias, disregard for individual specificity, and a lack of interaction between brain regions and FC analysis (Liu et al., 2024a). Despite research into various methods of addressing these issues, such as data-driven (Jensen et al., 2024), individualized (Li et al., 2022), and multi-atlas (Xu et al., 2024) setups, a definitive resolution to all of the challenges associated with atlas-based parcellation techniques has not yet emerged.

Given that one of the largest drawbacks of traditional FC analysis in the high dimensionality and complexity of the functional representations, solutions that address this particular issue are desired. Recently, Kolmogorov-Arnold Networks (KANs) (Liu et al., 2024b) have emerged as an alternative to traditional multi-layer perceptrons (MLPs), leveraging learnable activation functions on edges rather than fixed activation functions on nodes. Inspired by the Kolmogorov-Arnold representation theorem, KANs replace conventional weight matrices with univariate functions parameterized as splines, offering improved expressiveness and flexibility in function approximation. This design enables KANs to model complex

transformations more efficiently while maintaining better interpretability and scaling properties compared to MLPs. We hypothesize that replacing MLPs with KANs in brain disorder diagnosis modes can better capture intricate relationships in FC patterns, leading to more robust and individualized ASD diagnoses.

In this paper, we propose *ABFR-KAN*, a novel workflow for brain disorder diagnosis. Building upon state-of-the art methods, we propose novel sampling and function representation strategies and investigate the impact of KANs under various configurations in transformer networks. Our specific contributions are summarized as:

- Randomized anchor patch selection, which helps avoid structural bias, boosts individual-specific representations and increases robustness and generalizability by reducing dependence on atlas-based parcellation.
- Iterative sampling of patches from a subject's brain, aimed to create multiple function representations for the same subject, introducing variance while preserving meaningful FC information.
- Extensive experimentation demonstrating the effectiveness of replacing traditional MLP components in two transformer networks (ViT and DeiT).

2. Related Work

There generally exist three different setups for brain disorder analysis using an atlas: single-atlas, multi-atlas, and individual-specific atlas. An example of a model constructed from a single-atlas approach is BrainGNN (Li et al., 2021), a graph neural network (GNN) based on the Desikan-Killiany (Desikan et al., 2006) atlas that is capable of analyzing fMRI images and discovering neurological biomarkers. Another group employed multiple atlases (Kennedy et al., 1998; Craddock et al., 2012; Rolls et al., 2020) to build a spectral GNN that enabled the identification of potential disease-related patterns associated with major depressive disorder (Lee et al., 2024). PFC-DBGNN- STAA (Cui et al., 2023) was proposed as a method for identifying mild cognitive impairment (MCI) based on individual-specific FC features.

As an alternative to using pre-defined atlases for ROI parcellation, several data-driven approaches have been proposed. For example, attention-guided hybrid deep learning networks have been used to automatically localize discriminative brain regions for Alzheimer's disease and MCI diagnosis (Lian et al., 2020). RandomFR (Liu et al., 2024a) is an innovative approach for brain function representation, and operates via randomized selection of brain patches, and well as the use of novel function and position description methods. RandomFR serves as the main inspiration for the research presented in this paper.

Given the early stage of research into KANs, there exist few studies on the use of KANs for similar tasks that we propose in this paper. One study explores the use of KANs as deep feature extractors for MRI reconstruction, finding that incorporating Chebyshev polynomials into KANs (SS et al., 2024) led to both improved convergence and MRI reconstruction quality based on total variation and peak signal-to-noise ratio (Penkin and Krylov, 2024). Another study demonstrated the usefulness of KANs for chemical exchange saturation transfer (CEST) MRI analysis of the human brain (Wang et al., 2024), and

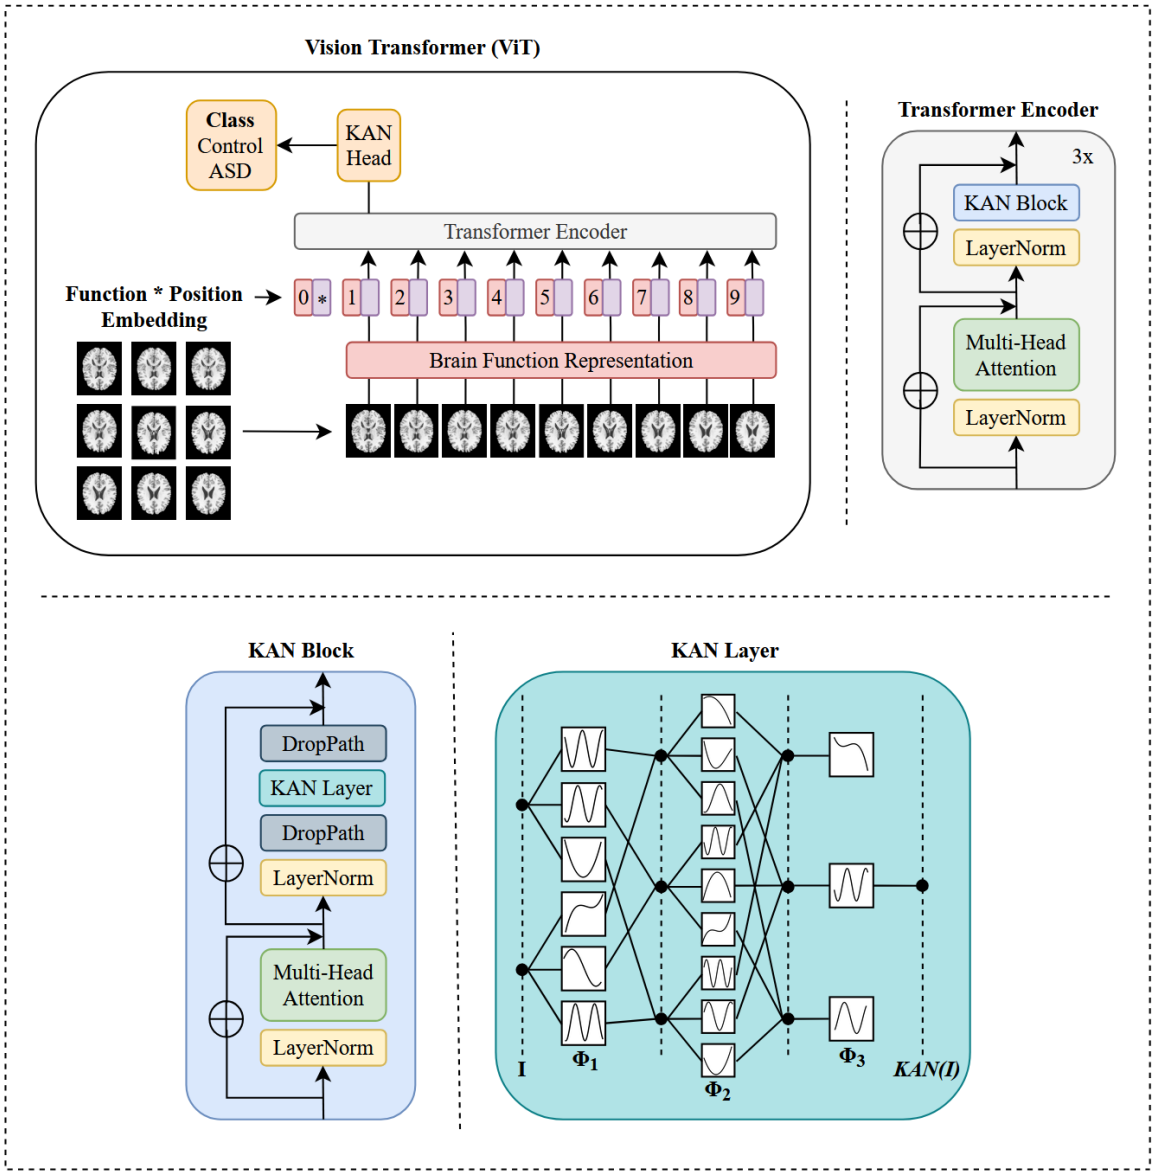

Figure 1: Proposed AFBR-KAN model: The transformer network is fed fMRI-derived patches that are embedded with spatial position information and passed through the encoder. The binary classification prediction (control or ASD) is produced by the KAN head. The encoder is in the ViT style, with a KAN block replacing the MLP block. The KAN block is similar to an MLP block, but with DropPath regularization and KAN layers to handle nonlinear transformations. In the KAN layer, input I passes through multiple learnable nonlinear functions (ϕ_n) that are combined in a structured manner to form the final transformations.

(a) Gray matter mask (b) Grid-based anchor selection (c) Randomized anchor selection

Figure 2: (a) The gray matter mask from which our anchor patches are selected. (b) The baseline grid-based anchor selection process. Note how certain patches fall outside of the gray matter region entirely. (c) Our randomized anchor selection process, which captures the full scope of the gray matter region, reduces structural bias and enhances individual specificity.

another found that a model integrating the learnable spline activation functions of KAN into convolution layers, ConvKAN, outperformed traditional convolutional neural network (CNN) and graph convolution network (GCN) approaches at classifying Parkinson's disease (Patel et al., 2024). To our knowledge, we are the first to investigate the efficacy of KANs for FC analysis and ASD diagnosis.

3. Methods

In this work, we follow the workflow structure described by (Liu et al., 2024a) for brain function representation, which is divided into three stages: **sampling**, where anchor patches are selected from gray matter along with position coordinates, **function representation**, where sampled patches are combined with function and position descriptions for characterization of brain function, and **transformer network**, where embeddings based on the fusion the function and position descriptions are passed to a transformer network for classification. In this section, we detail our methodology based on these three phases.

3.1. Random Anchor Selection

Previous works (Liu et al., 2024a), have explored selecting anchor patches using a grid-based method, where a grid of coordinates to sample anchor patches from is constructed from ROIs in a unified parcellation of gray matter, along with stride and offset values. This approach, which has proven effective, does limit models in terms of flexibility and adaptability because the same grid is used for every subject in a dataset, imposing a structural bias and a disregard for individual specificity, as a subject may have functionally distinct regions that do not align well with predefined anchor patches.

To address this flaw, we propose an alternate approach, a randomized one, for anchor patch selection. Our method works as follows. First, bounding boxes encompassing ROIs in the gray matter are calculated:

$$ROI = [x_{min}, y_{min}, z_{min}, x_{max}, y_{max}, z_{max}]. \tag{1}$$

Next, the starting coordinates are randomly sampled using:

$$x_{start} \sim \text{Uniform}(x_{min}, x_{max} - p_s), \tag{2}$$

where p_s is the patch size. Similar equations are used to calculate y_{start} and z_{start}. Once the anchor patches are sampled, they are validated to ensure sufficient overlap with the gray matter mask, g_m:

$$\sum(p_m \cdot g_m) \geq \tau, \tag{3}$$

where p_m is the patch mask and τ is a threshold. If the condition is not met, resample until a valid sample is found. This process repeats until the desired number of anchor patches is sampled. A visualization of both the grid-based anchor selection and our random anchor selection can be seen in Figure 2.

3.2. Iterative Patch Sampling

As opposed to sampling the patches used to represent the brain function of each subject only once (Liu et al., 2024a), we propose an iterative modification of this process to serve as a form of data augmentation. Our proposed method can be expressed as:

$$\text{Samples} = \bigcup_{i=1}^{m} \{(x, y, z, p_i) \mid p(x, y, z, p_i) \cap g_m \neq \emptyset\}, \tag{4}$$

where p is the patch. In our method, patches are samples over multiple iterations, with varying patch sizes. Each subject is sampled during every iteration.

3.3. Transformer Network

In this work, we explore KAN integration in two popular transformer networks: vision transformer (ViT) (Dosovitskiy, 2020) and data-efficient image transformer (Touvron et al., 2021). A visual depiction of our implementation of the KAN-based ViT is shown in Figure 1. Both networks have two main locations that traditionally are constructed with MLP components: in the encoder, and in the classification head. We experiment with three different configurations of KAN integration: KAN-KAN, where the MLPs in both the encoder and classification head are replaced with KANs, KAN-MLP, where only the MLP in the encoder is replaced, and MLP-KAN, where the opposite is true.

4. Experiments and Results

4.1. Data

We evaluated our proposed ABFR-KAN using pre-processed neuroimaging data from the Autism Brain Imaging Data Exchange (ABIDE) (Craddock et al., 2013; Di Martino et al.,

Table 1: Classification performance of ABFR-KAN with grid-based anchor selection and random patch sampling. The best and second best results are **bolded** and underlined, respectively.

Backbone	Model	ACC	AUC	F1	P	R	SPE
ViT	MLP-MLP	0.731±0.064	0.695±0.114	0.778±0.054	0.745±0.057	0.899±0.062	0.718±0.094
	KAN-KAN	0.708±0.092	0.706±0.080	0.768±0.039	0.762±0.145	0.867±0.121	0.692±0.124
	KAN-MLP	0.714±0.049	**0.718±0.089**	0.771±0.057	**0.765±0.108**	0.866±0.093	0.704±0.068
	MLP-KAN	**0.737±0.050**	0.686±0.096	**0.785±0.050**	0.741±0.070	**0.910±0.058**	**0.719±0.049**
DeiT	MLP-MLP	0.720±0.050	0.696±0.097	0.762±0.050	0.724±0.080	0.856±0.011	0.693±0.057
	KAN-KAN	**0.725±0.057**	**0.697±0.049**	**0.773±0.043**	**0.750±0.131**	0.846±0.060	**0.711±0.068**
	KAN-MLP	0.696±0.053	0.663±0.097	0.756±0.048	0.742±0.121	0.826±0.123	0.672±0.052
	MLP-KAN	0.679±0.058	0.662±0.025	0.767±0.017	0.733±0.080	**0.939±0.049**	0.663±0.066

2014). The preprocessed ABIDE repository contains data collected from a total of 1,112 patients at various sites, preprocessed using a variety of methods. We selected data from 171 patients that were collected from the NYU Langone Medical Center site that had been processed using the Data Processing Assistant for Resting-State fMRI (DPARSF) (Yan and Zang, 2010) method. Our selected subset of data contained information for 73 patients diagnosed with ASD, and 98 from a control group.

4.2. Implementation Details

The ABFR-KAN model was implemented with PyTorch and trained on a single GPU (NVIDIA GeForce RTX 4070, 16GB memory). A train-test split of 80-20 was used, meaning 136 fMRI scans were used for training and 35 for testing. A 5-fold cross-validation strategy was used to assess the model's performance. For classification, we minimize the cross-entropy loss. The model was trained for 100 epochs, using the Adam optimizer with a learning rate of 0.0009. The model's performance is gauged using traditional metrics for classification tasks, namely accuracy (ACC), area under curve (AUC), F1 score (F1), precision (P), recall (R), and specificity (SPE).

4.3. Results

4.3.1. Grid-based anchor selection, random patch sampling

In the first experiment, the main goal was to evaluate the performance of KAN integration, with no changes to the anchor selection or patch sampling processes. Our primary findings for this experiment are reported in Table 1. For the ViT backbone, we show that the final configuration, MLP-KAN achieves the best performance, outperforming the other models across four of the six evaluated metrics (KAN-MLP second best).

For the DeiT backbone, a different trend is observed. This time, KAN-KAN outperforms the other methods in five of the six metrics (MLP-KAN second best). Overall, the performance of the DeiT-based models is poorer than the ViT-based ones.

Table 2: Classification performance of ABFR-KAN with random anchor selection and random patch sampling. The best and second best results are **bolded** and underlined, respectively.

Backbone	Model	ACC	AUC	F1	P	R	SPE
ViT	MLP-MLP	0.702±0.033	0.667±0.081	0.737±0.057	0.723±0.061	0.776±0.125	0.704±0.039
	KAN-KAN	**0.743±0.068**	**0.727±0.112**	**0.783±0.061**	**0.764±0.084**	**0.890±0.132**	**0.729±0.072**
	KAN-MLP	0.720±0.045	0.704±0.041	0.768±0.055	0.734±0.082	0.825±0.116	0.699±0.037
	MLP-KAN	0.708±0.059	0.693±0.071	0.763±0.069	0.732±0.146	0.876±0.143	0.688±0.057
DeiT	MLP-MLP	0.696±0.072	0.671±0.092	0.755±0.065	0.725±0.080	0.865±0.139	0.682±0.085
	KAN-KAN	0.708±0.058	0.676±0.066	0.775±0.014	0.741±0.054	0.900±0.127	0.682±0.059
	KAN-MLP	0.684±0.051	0.622±0.020	0.767±0.032	0.722±0.117	**0.959±0.050**	0.665±0.065
	MLP-KAN	**0.713±0.058**	**0.689±0.060**	**0.779±0.038**	**0.762±0.114**	0.950±0.045	**0.687±0.082**

4.3.2. Random anchor selection, random patch sampling

In the second experiment, the randomized patch sampling process remained the same, but the anchor selection process was switched to a randomized selection strategy as defined in subsection 3.1. Our findings for this experiment are reported in Table 2. Comparing these results to the previous experiment, we observe inverse trends. This time, for the ViT-based models, KAN-KAN outperforms all of the models in all of the metrics, and outperforms the baseline at higher rates compared to the KAN-based models in the previous experiment.

For the DeiT backbone, this time MLP-KAN proves best, outperforming the other models in five of the six evaluated metrics (KAN-MLP second best). As with the previous experiment, the DeiT-based models generally performed poorer compared to the ViT-based models.

4.3.3. Grid-based anchor selection, iterative patch sampling

For the third experiment, we return to the grid-based anchor selection process but modify the patch sampling process to use an iterative approach as defined in subsection 3.2. This time, we observe identical trends between the ViT and DeiT backbones. For the ViT-based models, KAN-MLP performs the best, outperforming the other models in four of the six metrics (MLP-KAN second best).

KAN-MLP again performs the best for the DeiT models, achieving the best performance in five of the six metrics (MLP-KAN second best). Notably, in this experiment, the DeiT-based models are much more competitive with the ViT-based ones compared with the previous experiments, outperforming them in four of the six metrics.

4.4. Discussion

From our experiments, we make several observations. In each of the three experiments, our ABFR-KAN framework achieves better performance compared to the baseline RandomFR, indicating both the efficacy of our proposed method, as well as KANs in general. Interestingly though, taking into account the results of all the experiments, there is not a clear winner in terms of the configuration of ABFR-KAN.

Table 3: Classification performance of ABFR-KAN with grid-based anchor selection and iterative patch sampling. The best and second best results are **bolded** and underlined, respectively.

Backbone	Model	ACC	AUC	F1	P	R	SPE
ViT	MLP-MLP	0.643±0.047	0.612±0.040	0.760±0.049	0.725±0.043	0.839±0.124	0.652±0.040
	KAN-KAN	0.673±0.043	0.621±0.043	0.757±0.033	0.702±0.108	0.898±0.128	0.638±0.037
	KAN-MLP	**0.690±0.048**	**0.674±0.090**	**0.772±0.035**	0.683±0.045	**0.918±0.068**	0.651±0.051
	MLP-KAN	0.684±0.029	0.623±0.056	0.725±0.043	**0.733±0.064**	0.819±0.153	**0.671±0.048**
DeiT	MLP-MLP	0.667±0.047	0.608±0.089	0.733±0.029	0.685±0.063	0.844±0.156	0.647±0.055
	KAN-KAN	0.673±0.020	0.588±0.095	0.755±0.018	0.688±0.051	0.888±0.081	0.646±0.023
	KAN-MLP	**0.707±0.054**	**0.667±0.101**	**0.779±0.030**	0.706±0.075	**0.898±0.085**	**0.674±0.064**
	MLP-KAN	0.673±0.060	0.661±0.074	0.722±0.039	**0.769±0.096**	0.805±0.070	0.660±0.061

Considering the evaluations of the ViT and DeiT backbones as separate experiments within the three main experiments, we are left with a total of six scenarios. The three ABFR-KAN variants, KAN-KAN, KAN-MLP, and MLP-KAN have two scenarios each where they outperform the others, with a ViT-based model and a DeiT-based model present in the best cases for each of the three variants, making it difficult to evaluate whether one particular configuration definitively performs the best. However, given that four of the six scenarios are won by a configuration containing a hybrid of MLP and KAN blocks, it appears that taking a hybrid approach to KAN integration works better than having either all KANs or all MLPs.

Moving away from the discussion on KAN integration, we now discuss the modifications we made to the RandomFR architecture itself. Modifying the anchor selection process to be randomized as opposed to grid-based had a positive effect, with the best-performing model in the second experiment outperforming the best-performing model of the first experiment in three metrics for both the ViT-based and DeiT-based methods. The iterative patch sampling process on the other hand did not perform as well compared to the original RandomFR architecture. As the focus of this work was largely on exploring the efficacy of KAN integration, proper optimization of the randomized anchor selection process and iterative patch sampling process were not explored in depth. Additionally, we did not explore the use of random anchor patches with iteratively sampled patches, which could be an interesting exploration.

5. Conclusions

In this paper, we have introduced a novel architecture for improving ASD diagnosis, ABFR-KAN. Our results demonstrate the effectiveness of our proposed technique, and specifically highlight the effectiveness of KANs in this domain. Interestingly, strictly using KANs without any MLP components was shown to have decreased components compared to configurations where MLPs were involved in some way. This indicates the need for further research in this area, as there are many different ways that KANs can be incorporated into existing deep learning architectures, and optimal frameworks are not yet clear.

References

Michael Brant-Zawadzki, Gary D Gillan, and Wolfgang R Nitz. Mp rage: a three-dimensional, t1-weighted, gradient-echo sequence–initial experience in the brain. *Radiology*, 182(3):769–775, 1992.

Cameron Craddock, Yassine Benhajali, Carlton Chu, Francois Chouinard, Alan Evans, András Jakab, Budhachandra Singh Khundrakpam, John David Lewis, Qingyang Li, Michael Milham, et al. The Neuro Bureau Preprocessing Initiative: open sharing of preprocessed neuroimaging data and derivatives. *Frontiers in Neuroinformatics*, 7(27):5, 2013.

R Cameron Craddock, G Andrew James, Paul E Holtzheimer III, Xiaoping P Hu, and Helen S Mayberg. A whole brain fmri atlas generated via spatially constrained spectral clustering. *Human brain mapping*, 33(8):1914–1928, 2012.

Weigang Cui, Yulan Ma, Jianxun Ren, Jingyu Liu, Guolin Ma, Hesheng Liu, and Yang Li. Personalized functional connectivity based spatio-temporal aggregated attention network for mci identification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:2257–2267, 2023.

Rahul S Desikan, Florent Ségonne, Bruce Fischl, Brian T Quinn, Bradford C Dickerson, Deborah Blacker, Randy L Buckner, Anders M Dale, R Paul Maguire, Bradley T Hyman, et al. An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. *Neuroimage*, 31(3):968–980, 2006.

Adriana Di Martino, Chao-Gan Yan, Qingyang Li, Erin Denio, Francisco X Castellanos, Kaat Alaerts, Jeffrey S Anderson, Michal Assaf, Susan Y Bookheimer, Mirella Dapretto, et al. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular Psychiatry*, 19(6):659–667, 2014.

Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

Wenrong Ge, Cancan Zhang, Guang Yang, and Bo Zhang. Prevalence and trends of autism spectrum disorder and other developmental disabilities among children and adolescents in the United States from 2019 to 2021. *Frontiers in Psychiatry*, 15, 2024.

Kyle M Jensen, Jessica A Turner, Vince D Calhoun, and Armin Iraji. Addressing inconsistency in functional neuroimaging: A replicable data-driven multi-scale functional atlas for canonical brain networks. *bioRxiv*, 2024.

David N Kennedy, Nicholas Lange, Nikos Makris, Julianna Bates, James Meyer, and Verne S Caviness Jr. Gyri of the human neocortex: an mri-based analysis of volume and variance. *Cerebral Cortex (New York, NY: 1991)*, 8(4):372–384, 1998.

Deok-Joong Lee, Dong-Hee Shin, Young-Han Son, Ji-Wung Han, Ji-Hye Oh, Da-Hyun Kim, Ji-Hoon Jeong, and Tae-Eui Kam. Spectral graph neural network-based multi-atlas brain network fusion for major depressive disorder diagnosis. *IEEE Journal of Biomedical and Health Informatics*, 2024.

Xiaoxiao Li, Yuan Zhou, Nicha Dvornek, Muhan Zhang, Siyuan Gao, Juntang Zhuang, Dustin Scheinost, Lawrence H Staib, Pamela Ventola, and James S Duncan. Braingnn: Interpretable brain graph neural network for fmri analysis. *Medical Image Analysis*, 74, 2021.

Yu Li, Aiping Liu, Xueyang Fu, Martin J Mckeown, Z Jane Wang, and Xun Chen. Atlas-guided parcellation: Individualized functionally-homogenous parcellation in cerebral cortex. *Computers in Biology and Medicine*, 150:106078, 2022.

Chunfeng Lian, Mingxia Liu, Yongsheng Pan, and Dinggang Shen. Attention-guided hybrid network for dementia diagnosis with structural mr images. *IEEE transactions on cybernetics*, 52(4):1992–2003, 2020.

Mengjun Liu, Huifeng Zhang, Mianxin Liu, Dongdong Chen, Zixu Zhuang, Xin Wang, Lichi Zhang, Daihui Peng, and Qian Wang. Randomizing human brain function representation for brain disease diagnosis. *IEEE Transactions on Medical Imaging*, 2024a.

Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks. *arXiv preprint arXiv:2404.19756*, 2024b.

Salil B Patel, Vicky Goh, James F FitzGerald, and Chrystalina A Antoniades. 2d and 3d deep learning models for mri-based parkinson's disease classification: A comparative analysis of convolutional kolmogorov-arnold networks, convolutional neural networks, and graph convolutional networks. *arXiv preprint arXiv:2407.17380*, 2024.

Maksim Penkin and Andrey Krylov. Kolmogorov-arnold networks as deep feature extractors for mri reconstruction. In *Proceedings of the 2024 9th International Conference on Biomedical Imaging, Signal Processing*, pages 92–97, 2024.

Edmund T Rolls, Chu-Chung Huang, Ching-Po Lin, Jianfeng Feng, and Marc Joliot. Automated anatomical labelling atlas 3. *Neuroimage*, 206:116189, 2020.

Caio Pinheiro Santana, Emerson Assis de Carvalho, Igor Duarte Rodrigues, Guilherme Sousa Bastos, Adler Diniz de Souza, and Lucelmo Lacerda de Brito. rs-fmri and machine learning for asd diagnosis: a systematic review and meta-analysis. *Scientific reports*, 12(1):6030, 2022.

Sidharth SS, Keerthana AR, Anas KP, et al. Chebyshev polynomial-based kolmogorov-arnold networks: An efficient architecture for nonlinear function approximation. *arXiv preprint arXiv:2405.07200*, 2024.

Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR, 2021.

Jiawen Wang, Pei Cai, Ziyan Wang, Huabin Zhang, and Jianpan Huang. Cest-kan: Kolmogorov-arnold networks for cest mri data analysis. *arXiv preprint arXiv:2406.16026*, 2024.

Jiaxing Xu, Mengcheng Lan, Xia Dong, Kai He, Wei Zhang, Qingtian Bian, and Yiping Ke. Multi-atlas brain network classification through consistency distillation and complementary information fusion. *arXiv preprint arXiv:2410.08228*, 2024.

Chaogan Yan and Yufeng Zang. DPARSF: a MATLAB toolbox for "pipeline" data analysis of resting-state fmri. *Frontiers in Systems Neuroscience*, 4:1377, 2010.

Appendix A. Demographics

Figure 3 provides a detailed breakdown of the demographics among patients in the data subset. In total, 64 male (age range: 7-39) and 9 female (age range: 10-38) patients with ASD diagnoses were selected, along with 72 male (age range: 6-31) and 26 female (age range: 8-29) patients from the control group.

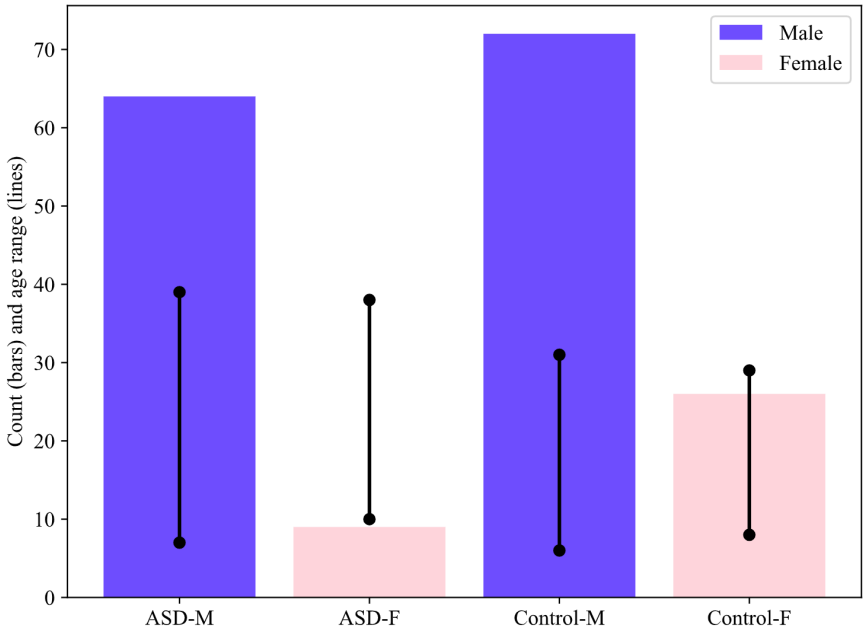

Figure 3: The patient demographics of our selected subset of data from the ABIDE I dataset.

Appendix B. Dataset Samples

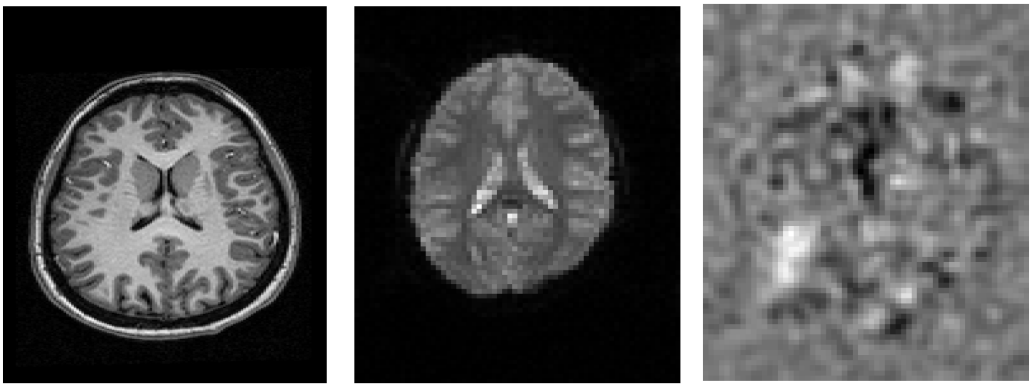

(a) Magnetization-Prepared Rapid Gradient Echo (MP-RAGE)

(b) rs-fMRI

(c) DPARSF pre-processed rs-fMRI patch

Figure 4: Raw and preprocessed data for a single patient in the ABIDE I dataset.

Figure 4 shows different patient-specific scans from the ABIDE I dataset. The first, MP-RAGE, is an MRI pulse sequence optimizing for T1-weighted imaging, allowing easy identification of anatomical features with high gray/white matter contrast (Brant-Zawadzki et al., 1992). The second, rs-fMRI, is a noninvasive technique used to measure and analyze brain activity when a subject is at rest, i.e., not engaged in a specific task. rs-fMRI's are widely used to study FC between brain regions (Santana et al., 2022). The third, and the most relevant to this research, is the DPARSF pre-processed fMRI data. DPARSF is a toolkit enabling easy pre-processing tasks such as slice timing, realignment, normalization, and smoothing data (Yan and Zang, 2010).

Appendix C. Patch Sampling

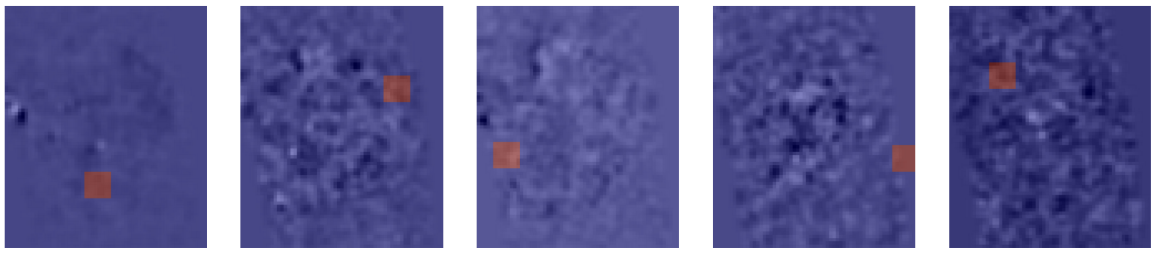

(a) Random patch sampling

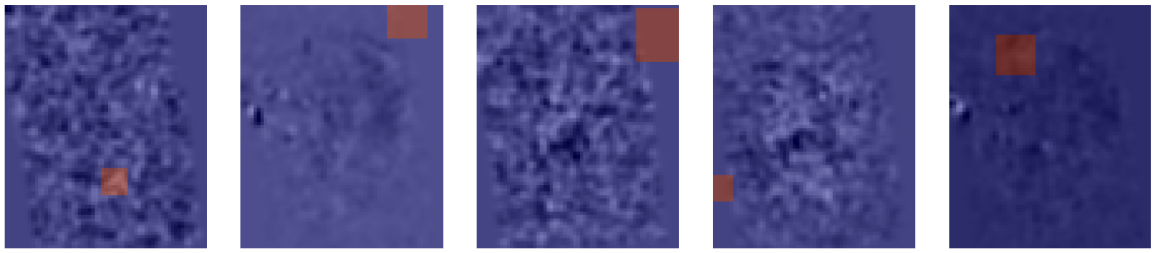

(b) Iterative patch sampling

Figure 5: (a) The random patch sampling process. Observe how the size of the patches is consistent. (b) The iterative patch sampling process, where each subject is processed three times as a form of data augmentation, with patch sizes varying from 8×8, 12×12, and 16×16.

Figure 5 shows the two different patch sampling processes used in our study. The first, random patch sampling, uses consistent patch sizes randomly selected from across the gray matter region, reducing structural bias while maintaining functional specificity. The second, the iterative sampling method, acts as a data augmentation technique with the aim of introducing variability while preserving meaningful FC information.

Appendix D. Receiver Operating Characteristic (ROC) Curves

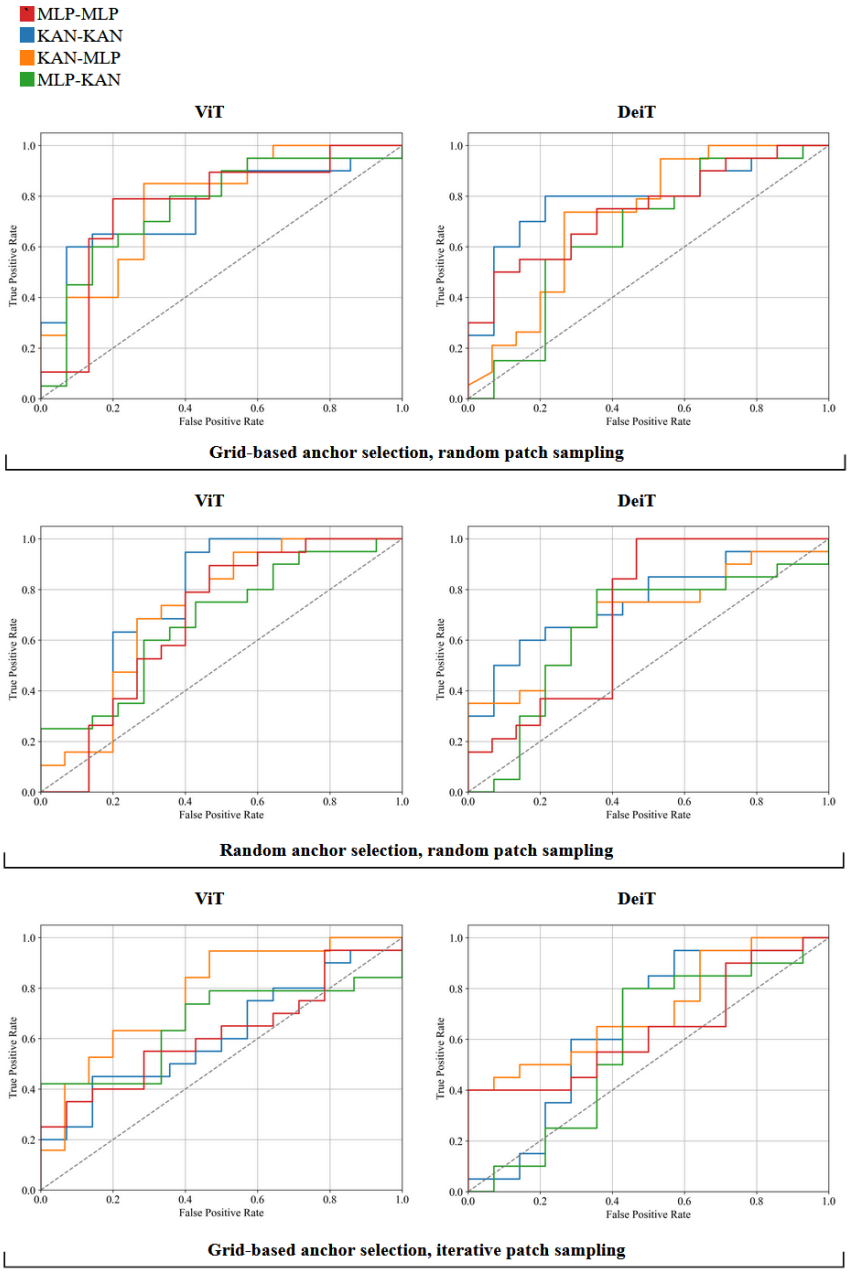

Figure 6: ROC curve comparison of our ABFR-KAN models vs. the baseline, which is reported as MLP-MLP. Note that our ABFR-KAN models generally achieve better curves compared to the MLP-MLP models.

Figure 6 presents a visual performance comparison between our ABFR-KAN models and the baseline. The ROC curves illustrate the trade-offs between true and false positive

rates. Our ABFR-KAN models demonstrate consistent performance improvement over the baseline, as indicated by the higher ROC curves. This suggests that replacing traditional MLP components with KANs enhances the model's ability to distinguish between ASD and control subjects. The observed improvements highlight the effectiveness of our approach in capturing complex FC patterns within the brain.