

Learning Explainable Imaging-Genetics Associations Related to a Neurological Disorder

Jueqi Wang¹, Zachary Jacokes², John Darrell Van Horn^{2,3}, Michael C. Schatz⁴, Kevin A. Pelphrey⁵, and Archana Venkataraman¹

¹ Department of Electrical and Computer Engineering, Boston University
{jueqiw, archanav}@bu.edu

² School of Data Science, University of Virginia
zj6nw@virginia.edu

³ Department of Psychology, University of Virginia
jdv7g@virginia.edu

⁴ Departments of Computer Science and Biology, Johns Hopkins University
mschatz@cs.jhu.edu

⁵ Department of Neurology, University of Virginia
kevin.pelphrey@virginia.edu

Abstract. While imaging-genetics holds great promise for unraveling the complex interplay between brain structure and genetic variation in neurological disorders, traditional methods are limited to simplistic linear models or to black-box techniques that lack interpretability. In this paper, we present NeuroPathX, an explainable deep learning framework that uses an early fusion strategy powered by cross-attention mechanisms to capture meaningful interactions between structural variations in the brain derived from MRI and established biological pathways derived from genetics data. To enhance interpretability and robustness, we introduce two loss functions over the attention matrix – a sparsity loss that focuses on the most salient interactions and a pathway similarity loss that enforces consistent representations across the cohort. We validate NeuroPathX on both autism spectrum disorder and Alzheimer’s disease. Our results demonstrate that NeuroPathX outperforms competing baseline approaches and reveals biologically plausible associations linked to the disorder. These findings underscore the potential of NeuroPathX to advance our understanding of complex brain disorders.

Keywords: Imaging-genetics · Neurological disorder · Multimodal fusion · Interpretability · Biologically-inspired deep learning

1 Introduction

Neurological disorders are marked by considerable heterogeneity, which is in part driven by inherited genetic factors [16,3]. These genetic influences can not only modulate disease susceptibility but can also shape the trajectory of disease progression. In parallel, brain imaging offers a unique window into how genetic variations may affect brain structure and function [23,30]. Thus, integrating

neuroimaging and genetics data promises a more comprehensive framework to unravel the complex interplay between neural and biological pathways, which in turn, will advance our understanding of multifaceted and heritable disorders.

A promising avenue for exploring the imaging-genetic interplay is brain imaging transcriptomics, which measures gene expression in brain tissue samples [14,7]. However, these studies rely on postmortem data from a small number of donors. Thus, they cannot capture dynamic processes in the brain both in time and across neurodevelopment [1]. In contrast, traditional *in vivo* imaging-genetics studies have largely modeled linear associations between single-nucleotide polymorphisms (SNPs) and neuroimaging phenotypes [26,33]. This simple framework may overlook the complex and nonlinear interactions that are critical to fully understanding the biological mechanisms underlying the disorder [6,27].

Machine learning (ML) has emerged as a powerful tool for identifying biomarkers in neurological diseases. Most often, a classifier is trained to distinguish between patients and neurotypical controls, and a post-hoc evaluation of the model is used to reveal novel biomarkers [4,9,17,27,34,11]. While ML models excel in capturing nonlinear interactions, they are often applied to a small subset of SNP values [34,11,19]. This approach is limiting, as the heritability of neurological traits is thought to result from the additive effects of hundreds of SNPs spread throughout the genome and genes functioning as part of intricate pathways [23].

A few works have explored deep learning for imaging-genetics by a *latent fusion* method, i.e., by combining the low-dimensional encoding for each modality [34,11,19]. The interpretability of these models are derived from a learned feature importance score. However, the latent fusion approach may not capture direct interactions between specific brain regions and SNPs. More recent studies have explored the *early-fusion* method, which models cross-modal interactions between the input features themselves [17,4,27]. This is done by using the cross-attention layer inside a transformer architecture to learn the interaction between the imaging and genetics data. However, these models are unable to control or to enforce consistency in the learned interactions across subjects. In addition, the proposed transformer architecture [17] uses a genetic encoder that contains a large number of parameters. Finally, the work of [12] takes the unique approach of constructing a graph neural network based on an established ontology of biological pathways [2]. However, this model operates at the gene level, which requires a large training sample and is prone to overfitting.

We propose NeuroPathX, an explainable AI framework that uses cross-attention to capture the intricate interplay between structural brain variation and genetic influences. To better align with the complexity of neurological disorders, NeuroPathX aggregates SNPs into well-established biological pathways [18], with each representing a network of genes that describes a specific cellular and/or metabolic function. Hence, NeuroPathX can deliver interpretable insights into the underlying neurobiology of a particular disorder and can learn effectively on small multimodal datasets. We promote interpretability via a pairwise loss term that enforces pathway-level coherence across subjects. We demonstrate NeuroPathX on two distinct datasets, one of autism spectrum disorder (ASD)




Fig. 1. NeuroPathX overview. Our framework integrates genetic variations (SNPs) and brain MRI features using a pathway-guided attention layer (yellow panel) to capture interactions. We use a sparsity and a pathway similarity loss to enhance interpretability.

and another of Alzheimer’s disease (AD). Our experiments underscore the robustness, interpretability, and predictive power of our proposed framework.

2 Pathway-driven Multimodal Fusion and Interpretability

Fig. 1 illustrates our proposed NeuroPathX framework. The inputs are a collection of imaging features extracted from brain MRI and a set of genetic features that encapsulate different biological processes. By employing a cross-attention layer for early fusion, the model not only effectively combines the two modalities but also offers interpretability through the analysis of the attention matrix. Formally, let $\mathbf{i}^n \in \mathbb{R}^{d \times N_I}$ be the input imaging feature vector for subject n , where N_I is the total number of ROIs, and d is the number of features extracted per ROI. The input construction from genetics data is described below.

2.1 Biological Pathway Features from SNPs

We first preprocess the raw SNP data using the conventional genome-wide polygenic risk score (PRS) pipeline [5] and select disease-specific SNPs based on a corresponding genome-wide association study (GWAS) from the literature. Let $\mathbf{x}^n \in \mathbb{R}^{N_G}$ denote the SNP vector for subject n , where N_G is the number of SNP positions from the GWAS. The selected SNPs are then mapped to genes based on established biological pathway databases [18]. Specifically, for each gene i , we compute a weighted sum of the corresponding SNP values using effect sizes β_j (where j is the SNP position from the disease-specific GWAS [26,31]). This process yields $\mathbf{g}_i^n = \sum_{j \in \Omega_i} \beta_j \mathbf{x}_j^n$, where Ω_i is the set of SNP indices mapped to gene i . We generate a pathway-level representation by aggregating these gene-level features for each pathway $\mathbf{p}_k^n$ as follows: $\mathbf{p}_k^n = \sum_{l \in \Gamma_k} g_l$, where Γ_k defines the set of genes associated with pathway k [18]. The pathway vector $\mathbf{p}^n \in \mathbb{R}^{1 \times N_p}$ is the input to NeuroPathX, where N_p denotes the total number of pathways.

2.2 Pathway-Guided Attention Layer

We use a series of (small) modality-specific encoding layers at the input. This transform the genetics input $\mathbf{p}^n$ into a representation $\mathbf{P}^n \in \mathbb{R}^{N_P \times d_q}$, where d_q is the pathway encoding dimension. For the imaging data, we construct $\mathbf{I}_k^n \in \mathbb{R}^{N_I \times d_k}$ and $\mathbf{I}_v^n \in \mathbb{R}^{N_I \times d_v}$ with encoding dimensions d_k and d_v , respectively. Unlike a traditional cross-attention mechanism, in which $\mathbf{I}_v^n$ and $\mathbf{I}_k^n$ are identical, we design separate encoding layers to learn more flexible associations. In addition, we preserve the input feature dimensionality to learn the direct interactions between individual genetic pathways and brain ROIs. The encoded representations are then projected via three learnable linear mappings $\mathbf{W}_q \in \mathbb{R}^{d_q \times d_{qk}}$, $\mathbf{W}_k \in \mathbb{R}^{d_k \times d_{qk}}$, and $\mathbf{W}_v \in \mathbb{R}^{d_v \times d}$ for a subsequent cross-attention layer.

Motivated by the works of [4,17], we derive a Pathway-guided Cross Attention (PathAttn) layer to learn the interactions between imaging and genetics. This layer uses the pathway guidance signal $\mathbf{P}^n$ to direct the aggregation of imaging features into a clustered set of Imaging-Guided pathway Concepts (IGC), denoted as $\mathbf{X}_{\text{IGC}}^n \in \mathbb{R}^{N_I \times d}$. Mathematically, this operation is defined as

$$\begin{aligned} \text{PathAttn}(\mathbf{I}_k^n, \mathbf{I}_v^n, \mathbf{P}^n) &= \text{SoftSign}\left(\text{ReLU}\left(\frac{\mathbf{P}^n \mathbf{W}_q \mathbf{W}_k^\top (\mathbf{I}_k^n)^\top}{C}\right)\right) \mathbf{W}_v \mathbf{I}_v^n \\ &= \mathbf{A}^n \mathbf{W}_v \mathbf{I}_v^n \triangleq \mathbf{X}_{\text{IGC}}^n \end{aligned} \quad (1)$$

where $\mathbf{A}^n \in \mathbb{R}^{N_I \times N_G}$ is the pathway-guided attention matrix for subject n and $\text{SoftSign}(x) = \frac{x}{0.5+x}$. We use a ReLU activation instead of the traditional softmax function in the cross-attention layer to accommodate the possibility that certain pathways may not be relevant to the disease and should have uniformly lower attention weights. To stabilize the ReLU output, we use a scaling constant $C = \gamma\sqrt{N_G}/2$ introduced by [25]. Finally, the imaging-guided pathway concept $\mathbf{X}_{\text{IGC}}^n$ is forwarded to a feed-forward classifier for disease prediction.

2.3 Enhancing Interpretability via the Loss Functions

We design two complementary loss functions on the attention matrix $\mathbf{A}^n$ to improve the explainability of NeuroPathX. These are described below.

Sparsity Loss: Our first loss encourages the attention matrix to focus on the most salient interactions between genetic pathways and brain regions. This is done via a sparsity penalty on the entries $\mathbf{A}_{ij}^n$ of the attention matrix

$$\mathcal{L}_{\text{sp}} = \sum_i \sum_j \text{KL}\left(\text{Ber}(q) \parallel \text{Ber}(\mathbf{A}_{ij})\right) \quad (2)$$

where $\text{Ber}(q)$ is a Bernoulli distribution with q determining the sparsity level [11]. Since Eq. (2) is computationally intensive, we approximate $\mathcal{L}_{\text{sp}}$ by applying the loss to two randomly selected subjects from each batch during training.

Pathway Similarity Loss: We hypothesize that there is a shared set of biological pathways associated with a particular disorder, and that these pathways have different influences across the brain ROIs between patients (PAT) and neurotypical controls (NC). We integrate this assumption into NeuroPathX by designing a novel paired similarity loss that encourages the learned pathway-level representations to be consistent between PAT and NC subjects.

Mathematically, we first sum the attention matrix $\mathbf{A}^n$ across its columns, i.e., over brain regions, yielding the vector $\mathbf{s}^n = \mathbf{A}^n \mathbf{1}_{N_{\mathcal{I}}}$, where $\mathbf{1}_{N_{\mathcal{I}}}$ is the all-ones vector of dimension $N_{\mathcal{I}}$ and $\mathbf{s}^n \in \mathbb{R}^{1 \times N_{\mathcal{G}}}$ for each subject n . During each training epoch, we randomly sample M pairs of PAT and NC subjects. For each sampled pair m , we construct the vectors $\mathbf{s}_{\text{PAT}}^{(m)}$ and $\mathbf{s}_{\text{NC}}^{(m)}$ from the corresponding attention matrices and define the similarity loss for the batch as:

$$\mathcal{L}_{\text{path}} = \frac{1}{M} \sum_{m=1}^M \left\| \mathbf{s}_{\text{NC}}^{(m)} - \mathbf{s}_{\text{PAT}}^{(m)} \right\|_1, \quad (3)$$

where the ℓ_1 norm both promotes sparsity and encourages consistent pathway-level representations across the patient and control cohorts.

Finally, we implement a weighted cross-entropy loss on the end classification network. The overall loss for our model is the sum of these terms:

$$\mathcal{L} = \lambda_{\text{sp}} \mathcal{L}_{\text{sp}} + \lambda_{\text{path}} \mathcal{L}_{\text{path}} + \sum_{n=1}^N \left[\delta \cdot y^n \log(\hat{y}^n) + (1 - \delta)(1 - y^n) \log(1 - \hat{y}^n) \right] \quad (4)$$

where y^n and $\hat{y}^n$ are the true and predicted classes of subject n , and δ controls the class weighting in the cross entropy function. The hyperparameter δ and the hyperparameters $\{\lambda_{\text{sp}}, \lambda_{\text{path}}\}$ are fixed based on the training data.

The multifaceted loss function NeuroPathX not only drives sparse and interpretable attention matrices, but it also ensures consistent pathway selection across the cohort while maintaining robust predictive performance.

2.4 Implementation Details

We maintain consistent hyperparameters across both datasets in Section 3. We set the learning rate to 5×10^{-5} and the Bernoulli probability to $q = 1 \times 10^{-2}$. The hyperparameters in Eq. (4) were fixed at $\lambda_{\text{path}} = 1 \times 10^{-3}$ and $\lambda_{\text{sp}} = 1 \times 10^{-6}$. In each epoch, we sample $M = 10$ (PAT/NC) pairs and employed a batch size of 128. The dimensions for the query, key, and value representations were set to $d_q = 32$, $d_k = 4$, $d_v = 8$, and the projected dimension $d_{qk} = 32$ for both datasets. NeuroPathX was trained for 1300 epochs on ACE and 100 epochs on ADNI.

NeuroPathX⁶ was optimized using AdamW with a weight decay of 1×10^{-5} and momentum parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$ [21]. Batch normalization was applied during training [15]. We use an MLP with dropout as the final classifier.

⁶ Code is available at <https://github.com/jueqiw/NeuroPathX>

Baseline Comparisons: We evaluate NeuroPathX against two ablated variants of our model: (1) omitting the sparsity loss $\mathcal{L}_{\text{sp}}$ and (2) omitting the pathway similarity loss $\mathcal{L}_{\text{path}}$, and to state-of-the-art methods in the literature. These baselines include G-MIND [11], which uses a decoder for regularization and learns a shared latent space between imaging and genetic data, and SurvPATH [17], which captures intermodal and intramodal interactions through a unified cross attention matrix. To mitigate overfitting, we did not implement the pathway embedding in SurvPATH. For consistency, we used the same pathway inputs and imaging feature representations as in NeuroPathX. Finally, we compared NeuroPathX with Canonical Correlation Analysis (CCA) as a classical baseline. For CCA, we project the z-scored imaging and genetic pathway features into a 10-D canonical space using linear CCA. The resulting imaging and genetics embeddings are concatenated and input to a logistic regression classifier.

3 Data and Preprocessing

We evaluate NeuroPathX on a publicly available dataset of autism spectrum disorder from an Autism Center of Excellence (ACE)⁷ and on the public Alzheimer’s Disease Neuroimaging Initiative (ADNI) dataset [29]. We remove related samples from both datasets to prevent the inflation of brain-pathway associations and to limit data leakage. This results in $N=165$ (98 ASD, 67 NC) subjects for ACE and $N=438$ subjects for ADNI (168 AD, 270 NC).

Genetics: Genotyping for ACE was done using the HumanOmni 2.5M BeadChip and for ADNI was done using the Human610-Quad BeadChip [24]. For both datasets, we include only autosomal SNPs with a genotyping rate above 99%, a Hardy–Weinberg equilibrium p-value larger than 1×10^{-6} , a minor allele frequency above 0.01, and a hard call threshold of at most 0.1.

We compute linkage disequilibrium (LD) matrices from the European population of the 1000 Genomes Project. For LD clumping, we set an R^2 threshold of 0.1 [28] and a physical distance threshold of 50 kilo-base pairs (KB) for both the ACE and ADNI datasets. After this process, 44,102 SNPs in ACE and 45,074 SNPs in ADNI remained available for SNP-to-gene mapping.

For SNP-to-gene mapping, we assign SNPs located within 50 KB to each gene. These genes are further aggregated according to the KEGG canonical pathway gene sets [18] to generate the input feature vector $\mathbf{p}^n$ (see Section 2.1). To better focus on neurological relevance, we omit KEGG pathways associated with cancer. This filtering resulted in $N_{\mathcal{P}} = 177$ pathways, encompassing 3,907 genes for ACE and 3,943 genes for ADNI, each with at least one mapped SNP.

Neuroimaging: We use T1-weighted MRI scans to extract anatomical features from the human brain, with preprocessing performed using FreeSurfer [10]. We parcellate the brain according to the Brainnetome atlas [20] ($N_{\mathcal{I}} = 210$ ROIs) for ACE and the Desikan-Killiany atlas [8] ($N_{\mathcal{I}} = 68$ ROIs) for ADNI. For each

⁷ Data available from https://nda.nih.gov/edit_collection.html?id=2021

Table 1. Quantitative evaluation of disease status prediction by NeuroPathX and the baseline and ablation on the ACE and ADNI datasets. The best performance for each metric is highlighted in bold. * denotes statistically worse performance than best model. ($p < 0.05$, paired two-tailed t-test)

	Methods	Accuracy $\uparrow$	Sensitivity $\uparrow$	Specificity $\uparrow$	AUC $\uparrow$
ACE	CCA	0.557 $\pm$ 0.07*	0.900$\pm$0.10	0.060 $\pm$ 0.10*	0.478 $\pm$ 0.06*
	G-MIND	0.569 $\pm$ 0.09*	0.593 $\pm$ 0.12*	0.567 $\pm$ 0.19*	0.618 $\pm$ 0.16*
	SurvPATH	0.538 $\pm$ 0.06*	0.469 $\pm$ 0.14*	0.638 $\pm$ 0.21*	0.534 $\pm$ 0.08*
	NeuroPathX w/o $\mathcal{L}_{path}$	0.586 $\pm$ 0.09*	0.512 $\pm$ 0.16*	0.690 $\pm$ 0.30*	0.652 $\pm$ 0.11*
	NeuroPathX w/o $\mathcal{L}_{sp}$	0.620 $\pm$ 0.10	0.743 $\pm$ 0.27*	0.743 $\pm$ 0.27	0.662 $\pm$ 0.12
	NeuroPathX	0.644$\pm$0.08	0.562 $\pm$ 0.13*	0.762$\pm$0.19	0.680$\pm$0.11
ADNI	CCA	0.592 $\pm$ 0.04*	0.242 $\pm$ 0.11*	0.810 $\pm$ 0.05*	0.526 $\pm$ 0.05*
	G-MIND	0.800 $\pm$ 0.06*	0.838 $\pm$ 0.09*	0.776 $\pm$ 0.09*	0.897 $\pm$ 0.05*
	SurvPATH	0.818 $\pm$ 0.06*	0.867$\pm$0.09	0.787 $\pm$ 0.08*	0.908$\pm$0.04
	NeuroPathX w/o $\mathcal{L}_{path}$	0.746 $\pm$ 0.06*	0.781 $\pm$ 0.09*	0.724 $\pm$ 0.07*	0.856 $\pm$ 0.05*
	NeuroPathX w/o $\mathcal{L}_{sp}$	0.758 $\pm$ 0.05*	0.800 $\pm$ 0.09*	0.731 $\pm$ 0.06*	0.858 $\pm$ 0.05*
	NeuroPathX	0.827$\pm$0.06	0.733 $\pm$ 0.08*	0.884$\pm$0.06	0.880 $\pm$ 0.06*

ROI, we extract $d = 4$ structural features: volume size, surface area, average cortical thickness, and standard deviation of the cortical thickness. The ADNI imaging features were available precomputed as part of the TADPOLE challenge [22]. All training, validation, and test imaging data were normalized using the mean and standard deviation computed from the training set.

4 Experimental Results

Quantitative Disease Prediction: Table 1 reports the patient vs. control classification performance of the testing folds within a 10-fold cross validation setup. As seen, NeuroPathX attains the highest accuracy and specificity for both the ACE and ADNI datasets. NeuroPathX also achieves the highest AUC on the ACE dataset and comparable performance to other deep learning models on the ADNI dataset. We note a drop in sensitivity for NeuroPathX, which may be attributed to the small sizes of the encoding and classifier networks. We note that removing either the pathway similarity or sparsity loss in the ablation models diminishes performance, reinforcing the critical role of these terms in capturing robust interactions. While SurvPATH and G-MIND achieve competitive sensitivity and AUC on ADNI, they under-perform in three of the four metrics on ACE, thus underscoring the generalizability of NeuroPathX.

Learned Imaging-Genetics Associations: To assess the interpretability of our model, we compute the mean attention matrices $\bar{\mathbf{A}}_{PAT}$ and $\bar{\mathbf{A}}_{ND}$ for the PAT and NC groups, respectively, across all testing folds in our 10-fold cross-validation. We then sum these matrices across columns to obtain pathway-level influences $\bar{\mathbf{s}}_{PAT}$, $\bar{\mathbf{s}}_{ND}$ for both the ACE and ADNI datasets. We extract the




Fig. 2. Visualization of the top brain-pathway associations learned by the mean attention matrices $\bar{\mathbf{A}}_{\text{PAT}}$ and $\bar{\mathbf{A}}_{\text{ND}}$ for ACE (left) and ADNI (right). The upper portion of each circle corresponds to brain ROIs, and the top portion corresponds to biological pathways. Purple connections are from $\bar{\mathbf{A}}_{\text{PAT}}$ and yellow connections are from $\bar{\mathbf{A}}_{\text{ND}}$.

intersection of the top seven pathways for each group (PAT and NC) for interpretation. Correspondingly, we extract four imaging ROIs that exhibit the highest values in the mean attention matrix for both PAT and NC groups.

Fig 2 right illustrates the predominant pathway–imaging feature interactions for the ACE and ADNI datasets. We observe that each pathway in the ACE dataset interacts with diverse brain ROIs, whereas in ADNI, each pathway is linked to more consistent ROIs. The patient and control groups in the ACE dataset also tend to select overlapping brain regions, suggesting a broader variety of neural alterations in autism. In contrast, for ADNI, the ROIs chosen by patients and controls diverge more noticeably, implying a greater separability between cohorts. These observations align with established clinical findings.

For autism (ACE), NeuroPathX highlights pathways involved in perception, neuro-inflammation, and cellular communication. These pathways interact with brain regions crucial for sensory processing, perception, emotion, and cognition. Such findings have been previously implicated in the ASD literature and support the biological relevance of our learned associations [13,32]. For Alzheimer’s disease (ADNI), our model identifies pathways related to neuroinflammation and neurodegeneration, interacting with brain regions responsible for memory, language, and social functions. These associations are consistent with well-documented patterns of cortical atrophy and functional decline in AD [8].

5 Conclusion

We have presented NeuroPathX, an explainable deep learning framework that integrates genetics data with MRI scans to elucidate the complex interactions between biological pathways and brain structure in neurological disorders. Our approach leverages an early fusion strategy with a cross-attention layer to effectively combine multimodal information. To enhance interpretability, we introduced two specialized loss functions: a sparsity loss to focus on the most critical interactions within the attention matrix, and (2) a pathway similarity loss to ensure consistent pathway identification across cohorts. We validated NeuroPathX on datasets for ASD and AD. Our results demonstrate that the proposed framework not only achieves high prediction accuracy but also uncovers biologically meaningful interactions between genetic pathways and brain regions.

Acknowledgments. This work was supported by the National Institutes of Health awards 1R01HD108790 (PI Venkataraman) and 1R01EB029977 (PI Caffo) and the National Science Foundation CAREER award 1845430 (PI Venkataraman).

Disclosure of Interests. All authors disclosed no relevant relationships.

References

1. Arnatkeviciūtė, A., Fulcher, B.D., Fornito, A.: A practical guide to linking brain-wide gene expression and neuroimaging data. *Neuroimage* **189**, 353–367 (2019)
2. Ashburner, M.o.: Gene ontology: tool for the unification of biology. *Nature genetics* **25**(1), 25–29 (2000)
3. Buch, A.M., et al.: Molecular and network-level mechanisms explaining individual differences in autism spectrum disorder. *Nature neuroscience* **26**(4), 650–663 (2023)
4. Chen, R.J., et al.: Multimodal co-attention transformer for survival prediction in gigapixel whole slide images. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4015–4025 (2021)
5. Choi, S.W., et al.: Tutorial: a guide to performing polygenic risk score analyses. *Nature protocols* **15**(9), 2759–2772 (2020)
6. Choi, S.W., et al.: Prset: Pathway-based polygenic risk score analyses and software. *PLoS genetics* **19**(2), e1010624 (2023)
7. Consortium, G.: The gtex consortium atlas of genetic regulatory effects across human tissues. *Science* **369**(6509), 1318–1330 (2020)
8. Desikan, R.S., 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)
9. Dvornek, N.C., et al.: Copy number variation informs fmri-based prediction of autism spectrum disorder. In: *International Workshop on Machine Learning in Clinical Neuroimaging*. pp. 133–142. Springer (2023)
10. Fischl, B.: Freesurfer. *Neuroimage* **62**(2), 774–781 (2012)
11. Ghosal, S., et al.: G-mind: an end-to-end multimodal imaging-genetics framework for biomarker identification and disease classification. In: *Medical Imaging 2021: Image Processing*. vol. 11596, pp. 63–72. SPIE (2021)

12. Ghosal, S., et al.: A biologically interpretable graph convolutional network to link genetic risk pathways and imaging phenotypes of disease. In: International conference on learning representations (2022)
13. Hartig, R., et al.: Genetic influences of autism candidate genes on circuit wiring and olfactory decoding. *Cell and tissue research* **383**(1), 581–595 (2021)
14. Hawrylycz, M.J., et al.: An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature* **489**(7416), 391–399 (2012)
15. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International conference on machine learning. pp. 448–456. pmlr (2015)
16. Jack, A., et al.: A neurogenetic analysis of female autism. *Brain* **144**(6), 1911–1926 (2021)
17. Jaume, G., et al.: Modeling dense multimodal interactions between biological pathways and histology for survival prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11579–11590 (2024)
18. Kanehisa, M., Goto, S.: Kegg: kyoto encyclopedia of genes and genomes. *Nucleic acids research* **28**(1), 27–30 (2000)
19. Lee, S., et al.: Integrating genetic information for early alzheimer’s diagnosis through mri interpretation. In: 2023 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI). pp. 1–6 (2023)
20. Li, W., et al.: Brainnetome atlas of preadolescent children based on anatomical connectivity profiles. *Cerebral Cortex* **33**(9), 5264–5275 (2023)
21. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
22. Marinescu, R.V., et al.: Tadpole challenge: Accurate alzheimer’s disease prediction through crowdsourced forecasting of future data. In: Predictive Intelligence in Medicine: Second International Workshop, PRIME 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 13, 2019, Proceedings 2. pp. 1–10. Springer (2019)
23. Pretzsch, C.M., Ecker, C.: Structural neuroimaging phenotypes and associated molecular and genomic underpinnings in autism: a review. *Frontiers in Neuroscience* **17**, 1172779 (2023)
24. Saykin, A.J., et al.: Genetic studies of quantitative mci and ad phenotypes in adni: Progress, opportunities, and plans. *Alzheimer’s & Dementia* **11**(7), 792–814 (2015)
25. Shen, K., et al.: A study on relu and softmax in transformer. arXiv preprint arXiv:2302.06461 (2023)
26. Shen, L., et al.: Whole genome association study of brain-wide imaging phenotypes for identifying quantitative trait loci in mci and ad: A study of the adni cohort. *Neuroimage* **53**(3), 1051–1063 (2010)
27. Smedley, N.F., El-Saden, S., Hsu, W.: Discovering and interpreting transcriptomic drivers of imaging traits using neural networks. *Bioinformatics* **36**(11), 3537–3548 (2020)
28. Weiner, D.J., et al.: Polygenic transmission disequilibrium confirms that common and rare variation act additively to create risk for autism spectrum disorders. *Nature genetics* **49**(7), 978–985 (2017)
29. Weiner, M.W., et al.: Recent publications from the alzheimer’s disease neuroimaging initiative: Reviewing progress toward improved ad clinical trials. *Alzheimer’s & Dementia* **13**(4), e1–e85 (2017)
30. Wen, Y., Alshikho, M.J., Herbert, M.R.: Pathway network analyses for autism reveal multisystem involvement, major overlaps with other diseases and convergence upon mapk and calcium signaling. *PloS one* **11**(4), e0153329 (2016)

31. Wightman, D.P., et al.: A genome-wide association study with 1,126,563 individuals identifies new risk loci for alzheimer’s disease. *Nature genetics* **53**(9), 1276–1282 (2021)
32. Xu, N., Li, X., Zhong, Y.: Inflammatory cytokines: potential biomarkers of immunologic dysfunction in autism spectrum disorders. *Mediators of inflammation* **2015**(1), 531518 (2015)
33. Xu, Z., et al.: Imaging-wide association study: integrating imaging endophenotypes in gwas. *Neuroimage* **159**, 159–169 (2017)
34. Zhou, R., et al.: Integrating multimodal contrastive learning and cross-modal attention for alzheimer’s disease prediction in brain imaging genetics. In: 2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). pp. 1806–1811. IEEE (2023)