



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

CALM: Interpretable Cross-Modal Alignment for Biomarker Discovery from Unpaired Data

Jueqi Wang¹, Zachary Jacokes², John Darrell Van Horn^{2,3}, Kevin A. Pelphrey⁴, Michael C. Schatz⁵, 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 Department of Neurology, University of Virginia
kevin.pelphrey@virginia.edu

⁵ Departments of Computer Science and Biology, Johns Hopkins University
mschatz@cs.jhu.edu

Abstract. The interaction between brain structure and genetic influences is key to understanding neuropsychiatric disorders. However, most large-scale datasets are unimodal, providing either neuroimaging or genetics data. We propose CALM, a framework that learns interpretable associations between brain ROIs and genetic pathways from *completely disjoint* populations. CALM aligns the two modalities in a shared latent space via linear projections that simultaneously match the class-conditional latent distributions and ensure group separability. These projections provide interpretable pathway–ROI associations. When trained on unimodal imaging and genetics datasets, CALM generalizes to an unseen paired dataset, outperforming several state-of-the-art methods and ablation baselines. We also demonstrate stability of the learned associations against a paired baseline. Our experiments on autism spectrum disorder reveal immune and metabolic pathways linked to specific cortical regions and are consistent with established literature. Thus, CALM opens the door to leveraging large unimodal repositories for studying cross-modal interactions in brain disorders across disparate datasets.

Keywords: Unpaired Multimodal Learning · Imaging-Genetics · ASD

1 Introduction

Disorders such as autism spectrum disorder (ASD) are driven by complex interactions between genetic risk factors and brain morphology [13]. Understanding these relationships is crucial for biomarker identification; however, acquiring both neuroimaging and genetics data from the same individuals is costly and often impractical at scale. As a result, the largest public resources in precision psychiatry are non-overlapping and modality specific. For example, the Autism Brain

Imaging Data Exchange (ABIDE) [4] offers comprehensive neuroimaging data, while the Simons Simplex Collection (SSC) [6] provides extensive genomic data for ASD *for a different set of participants*. Current multimodal analyses require paired samples and cannot directly be applied to these large datasets.

Traditional imaging-genetics is based on linear models between a restricted set of neuroimaging and genetics features [30]. This paradigm struggles to capture the complex dependencies associated with neuropsychiatric disorders [32]. Deep learning offers a promising alternative for imaging-genetics. For instance, latent fusion methods combine low-dimensional encodings from each modality [8, 18]; however, the information compression obscures the link between brain regions and genomic features, thus limiting biological interpretability. Early fusion approaches using cross-attention can explicitly model such interactions [14, 32, 34], but they cannot enforce consistency in learned associations across subjects, and they often require models with a large number of parameters [14, 34]. Graph-based methods have been used to incorporate *a priori* biological knowledge [9] but operate at the gene level and are prone to overfitting on small samples. Critically, these approaches require paired data from the same individuals, whereas public neuroimaging and genetic repositories are largely disjoint.

Recent works have addressed the problem of unpaired multimodal learning, but at the cost of interpretability [11]. For example, spectral embedding methods such as SUE [37] align the modalities through a canonical correlation analysis and distribution matching, but the learned representations lack clear biological meaning. Additive fusion strategies require both modalities to share identical dimensionality, similarly obscuring interpretable structure [8, 38]. In contrast, optimal transport (OT) methods compute a minimum-cost coupling between two distributions [2, 20] and is the most natural when distributions describe comparable entities under a shared geometry. However, brain morphometry and genetic pathways are fundamentally different entities, and no biologically grounded metric exists to quantify distances between them, making the transport cost dependent on modeling assumptions rather than biological structure.

To address these limitations, we propose Class-conditional Alignment with Linear Maps (CALM), a framework that discovers interpretable associations between *completely unpaired* neuroimaging and genetics datasets. CALM pretrains modality-specific encoders that preserve diagnostic groupings while ensuring a direct link between latent dimensions and neurobiological features. Cross-modal alignment is implemented as low-dimensional linear mappings in the latent space that optimize both a class-conditional alignment loss to match the latent imaging and genetics distributions within each diagnostic group, and a contrastive loss to pull same-class samples together while pushing apart different-class samples, regardless of modality. We validate CALM by training on ABIDE (neuroimaging) and SSC (genetics) and testing on the independent and paired ACE dataset. We show that associations recovered from unpaired data closely agree with those obtained from paired data and remain stable across cross-validation folds. CALM lays a crucial foundation to use independent unimodal datasets for cross-modal biomarker discovery, thus capitalizing on existing data resources.




Fig. 1: Overview CALM. Given unpaired MRI and genetics data, pretrained encoders ($\mathcal{E}_I$, $\mathcal{E}_G$) output item-specific representations that are projected into a shared latent space via the learned linear maps $\mathbf{W}_I$ and $\mathbf{W}_G$. The loss function simultaneously matches the within-class distributions across modalities, while separating diagnostic groups. The cross-modal association matrix $\mathbf{A} = \mathbf{W}_I^\top \mathbf{W}_G \in \mathbb{R}^{n_I \times n_G}$ quantifies the association between each brain region and genetic pathway.

2 Interpretable Linear Alignment for Unpaired Data

CALM tackles the problem of discovering biomarker associations between neuroimaging and genetics data, when the modalities are collected from *separate and non-overlapping groups of subjects*. Our framework is outlined in Fig. 1.

The neuroimaging input to CALM consists of region-wise morphological features from structural MRI. Formally, let $\mathbf{S}^{(i)} \in \mathbb{R}^{n_I \times n_F}$ denote the neuroimaging input for subject i , where n_I is the number of brain regions, and n_F is the number of morphological features. The genetics input to CALM is obtained by aggregating genetic risk scores across established biological pathways [16]. Due to the heterogeneity of neuropsychiatric disorders, we include pathway-level features for several phenotypic traits. Details of this process are provided in Section 3. Formally, let $\mathbf{P}^{(j)} \in \mathbb{R}^{n_G \times n_T}$ denote the genetics input for subject j , where n_G and n_T are the number of genetic pathways and phenotypic traits, respectively.

2.1 Data Encoding

CALM uses pretrained autoencoders and classifiers to construct the latent representations for each modality. The loss combines reconstruction quality, classification accuracy, and a contrastive term that encourages latent representations from the same cohort to cluster together. This cohort-level grouping enables CALM to use linear alignment maps to discover stable correspondences between the two modalities. Crucially, the encoders consist of small per-entity layers that preserve the region-level information from the neuroimaging input and pathway-level information from the genetics input. This strategy preserves a one-to-one correspondence between rows and neurobiological entities, thus maintaining the

interpretability of CALM. Let $\mathbf{X}_I^{(i)} \in \mathbb{R}^{n_I \times d}$ be the latent neuroimaging representation for subject i and $\mathbf{X}_G^{(j)} \in \mathbb{R}^{n_G \times d}$ be the latent genetics representation for subject j . The latent dimension d is shared between the modalities.

2.2 Linear Projection

Given the latent encodings $\mathbf{X}_I^{(i)}$ and $\mathbf{X}_G^{(j)}$, we map the modalities into a shared

latent space using learnable linear projections $\mathbf{W}_I \in \mathbb{R}^{d \times n_I}$ and $\mathbf{W}_G \in \mathbb{R}^{d \times n_G}$:

$$\mathbf{Z}_I^{(i)} = \mathbf{W}_I \mathbf{X}_I^{(i)}, \quad \mathbf{Z}_G^{(j)} = \mathbf{W}_G \mathbf{X}_G^{(j)} \quad (1)$$

For convenience, we flatten the matrices in Eq. (1) into the vectors: $\mathbf{z}_I^{(i)} = \text{vec}(\mathbf{Z}_I^{(i)})$ and $\mathbf{z}_G^{(j)} = \text{vec}(\mathbf{Z}_G^{(j)})$ when defining the distributional loss in the next section. The association matrix $\mathbf{A} = \mathbf{W}_I^\top \mathbf{W}_G \in \mathbb{R}^{n_I \times n_G}$ quantifies the similarity between cross-modal projection vectors, i.e., $\mathbf{A}_{r,p} = \mathbf{w}_{I,r}^\top \mathbf{w}_{G,p}$. As seen, $\mathbf{A}_{r,p}$ is large when ROI r and pathway p load along the same latent directions, indicating they co-vary in the aligned space. This formulation parallels that of canonical correlation analysis (CCA) but does not require paired samples.

2.3 Class-Conditional Distribution Alignment

CALM leverages a compound objective function to align the imaging and genetic distributions without requiring paired multi-modal samples:

$$\mathcal{L}_{total} = \lambda_{cmmd} \mathcal{L}_{cmmd} + \lambda_{con} \mathcal{L}_{con} + \lambda_{orth} \mathcal{L}_{orth} \quad (2)$$

The loss terms in Eq. (2) address three challenges of unpaired learning: stabilizing geometric alignment, preventing negative transfer between diagnostic groups, and enforcing semantic discriminability in the shared latent space. The hyperparameters $\lambda_{cmmd}, \lambda_{con}, \lambda_{orth}$ control the contribution of each term.

Distributional Loss ($\mathcal{L}_{cmmd}$): We use Class-Conditional Maximum Mean Discrepancy (CMMD) to align the distributions of latent neuroimaging and genetics encodings *within each diagnostic group*. Unlike the standard MMD [10], which is label agnostic, $\mathcal{L}_{cmmd}$ calculates the discrepancy separately for each class c :

$$\mathcal{L}_{cmmd} = \sum_{c \in \mathcal{C}} \left\| \frac{1}{N_c} \sum_{i \in c} \phi(\mathbf{z}_I^{(i)}) - \frac{1}{M_c} \sum_{j \in c} \phi(\mathbf{z}_G^{(j)}) \right\|_{\mathcal{H}}^2, \quad (3)$$

where $\phi(\cdot)$ maps features to a reproducing kernel Hilbert space (RKHS), and N_c, M_c denote the number of imaging and genetic samples in class c , respectively.

In our experiments, the classes $c \in \mathcal{C}$ are ASD patients and neurotypical controls (NC), but the CMMD loss can be generalized to other scenarios.

Contrastive Loss ($\mathcal{L}_{con}$): While $\mathcal{L}_{cmmd}$ aligns the within-class distributions, it does not enforce a separation *between classes* in the shared latent space. Thus,

we employ a supervised contrastive loss [17] that treats all same-class samples as positives regardless of modality. Let $\mathcal{B}$ denote the set of latent representations from both modalities included within a mini-batch. For an anchor $\mathbf{z}^{(i)}$ with label $y_i \in \mathcal{C}$, the positive set contains all other samples sharing the same label, $\mathcal{P}(i) = \{j \neq i \in \mathcal{B} : y_j = y_i\}$. The contrastive loss is computed

$$\mathcal{L}_{con} = - \sum_{i \in \mathcal{B}} \frac{1}{|\mathcal{P}(i)|} \sum_{p \in \mathcal{P}(i)} \log \frac{\exp(\text{sim}(\mathbf{z}^{(i)}, \mathbf{z}^{(p)})/\tau)}{\sum_{j \in \mathcal{B} \setminus \{i\}} \exp(\text{sim}(\mathbf{z}^{(i)}, \mathbf{z}^{(j)})/\tau)}, \quad (4)$$

where τ is a temperature parameter, and $|\mathcal{P}(i)|$ is the total number of positive samples for anchor $\mathbf{z}^{(i)}$. Eq. (4) simultaneously pulls samples with the same class together across modalities and pushes apart samples in different classes. This procedure encourages a shared discriminative structure in the latent space.

Regularization ($\mathcal{L}_{orth}$): To prevent degenerate solutions for the projection matrices, we use regularization to encourage orthonormality in the rows:

$$\mathcal{L}_{orth} = \|\mathbf{W}_I \mathbf{W}_I^\top - \mathbf{I}\|_F^2 + \|\mathbf{W}_G \mathbf{W}_G^\top - \mathbf{I}\|_F^2 \quad (5)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, and $\mathbf{I}$ is the identity matrix. Each row of $\mathbf{W}_I$ and $\mathbf{W}_G$ defines the projection of a single latent dimension across input features. Encouraging their outer products to be identity ensures that the projections remain orthogonal with unit-norm. This prevents rank collapse and encourages each latent dimension to capture non-redundant information.

2.4 Implementation Details

CALM⁶ follows a two-stage training procedure: modality-specific encoders and classifiers are pretrained for 50 epochs, the classifiers are frozen while the linear projection matrices are trained for cross-modal alignment over 30 epochs. The inputs $\mathbf{S}^{(i)}$ and $\mathbf{P}^{(j)}$ are z-score normalized using per-feature statistics computed solely from the training set. The hyperparameters in Eq. (2) were set to $\lambda_{cmm} = 5$, $\lambda_{con} = 5$, and $\lambda_{orth} = 0.01$, with a contrastive temperature of $\tau = 0.07$. The projector learning rate was set to 1×10^{-3} , with a reduced rate of 1×10^{-5} for fine-tuning the encoders. We used layer normalization, a batch size of 64, and applied Gaussian feature noise ($\sigma = 0.15$) to the imaging modality for regularization. The shared latent dimension was set to $d = 6$. The hyperparameters were tuned via grid search with 5-fold CV on ABIDE and SSC, covering latent dimension d , batch size, learning rate, loss weights, and temperature τ . ACE was never seen during tuning, and Table 1 confirms generalization to it.

Class Prediction on Paired Test Data: The distributional and contrastive losses refine both encoders during the second stage of training, so the updated representations now capture cross-modal structure. As a quantitative validation, we use the updated encoders and previously frozen stage one classifiers to predict

⁶ Code is available at <https://github.com/jueqi/CALM>

the diagnosis of unseen test subjects. As described in Section 3, our test cohort contains paired data. Thus, the final prediction is obtained by averaging the output probabilities for both the neuroimaging and genetics classifiers.

Baseline Comparisons: We evaluate the quantitative performance of CALM in predicting diagnostic labels against several state-of-the-art multimodal baselines. G-MIND [8] and UNSEEN [38] both project modality-specific latent encodings into a shared latent space and merge them through summation. Unlike CALM, these methods blur information across input features, thus limiting our ability to extract feature-wise associations between the modalities. SUE [37] aligns independently learned spectral embeddings across modalities without paired samples, but also lacks interpretable cross-modal associations. For consistency, all baselines use the same pathway inputs and imaging feature representations as CALM. We also compare CALM against an optimal transport ablation that replaces the $\mathcal{L}_{cmmid}$ loss with a global Sinkhorn divergence [2]. Finally, we include two other ablations of CALM, omitting the distributional loss $\mathcal{L}_{cmmid}$ and omitting the contrastive loss $\mathcal{L}_{con}$, to assess their contributions.

3 Data and Preprocessing

We evaluate CALM using four publicly available datasets for ASD. For training, we use neuroimaging data from the ABIDE I and II datasets [4] ($N = 1,898$; 889 ASD) and genetics data from SSC [6] ($N = 2,452$). For SSC simplex families, we use proband genotypes as ASD samples and construct pseudo-controls from untransmitted parental alleles [5]. For testing, we use the *independent* ACE⁷ dataset ($N = 198$; 105 ASD), which contains paired MRI and genetic data.

Neuroimaging: T1-weighted MRI from ABIDE and ACE were preprocessed using FreeSurfer [7] and parcellated according to the Brainnetome atlas [19] ($n_I = 246$ ROIs). For each ROI, we extract $n_F = 4$ structural features: volume size, surface area, average cortical thickness, and standard deviation of the cortical thickness. We use ComBat [15] to remove ABIDE site effects. The input features $\mathbf{S}^{(i)}$ were normalized based on statistics from the training set.

Genetics: SSC and ACE participants were genotyped on different Illumina chips. Quality control was performed separately per batch and cohort using PLINK [28]. Analyses were restricted to a European subpopulation identified via projection onto the first three principal components of the 1,000 Genomes reference panel [1]. We used the genetics preprocessing steps outlined in [22], which after LD clumping, yielded 44,102 SNPs. The SNPs were grouped into 3,907 genes defined by the $n_P = 177$ non-cancer pathways in the KEGG database [16].

Following [34], each entry $\mathbf{P}_{gt}^{(n)}$ in the input $\mathbf{P}^{(n)} \in \mathbb{R}^{n_G \times n_T}$ corresponds to the genetic risk for a single pathway g and phenotypic trait t . This value is obtained by first selecting the SNPs associated with the pathway genes, weighting them by the effect sizes obtained from a published GWAS [35], and computing a

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

Table 1: Classification performance on ACE. Models are trained on ABIDE (neuroimaging) and SSC (genetics). The best performance is highlighted in bold. * denotes statistically worse performance than the best model ($p < 0.05$).

Methods	Accuracy $\uparrow$	Sensitivity $\uparrow$	Specificity $\uparrow$	AUC $\uparrow$
G-MIND [8]	0.518 $\pm$ 0.01*	0.533 $\pm$ 0.04*	0.490 $\pm$ 0.04*	0.540 $\pm$ 0.01*
UNSEEN [38]	0.499 $\pm$ 0.02*	0.490 $\pm$ 0.25	0.510 $\pm$ 0.23	0.522 $\pm$ 0.02*
SUE [37]	0.486 $\pm$ 0.05*	0.430 $\pm$ 0.11*	0.549 $\pm$ 0.09	0.478 $\pm$ 0.04*
Ablation $\mathcal{L}_{OT}$ [2]	0.567 $\pm$ 0.03	0.573 $\pm$ 0.08	0.568$\pm$0.06	0.599 $\pm$ 0.04
Ablation $-\mathcal{L}_{cmm}$	0.570 $\pm$ 0.03	0.568 $\pm$ 0.08	0.548 $\pm$ 0.08	0.585 $\pm$ 0.04
Ablation $-\mathcal{L}_{con}$	0.568 $\pm$ 0.03	0.575 $\pm$ 0.09	0.557 $\pm$ 0.09	0.598 $\pm$ 0.04
CALM	0.572$\pm$0.03	0.583$\pm$0.08	0.562 $\pm$ 0.06	0.606$\pm$0.03




Fig. 2: Validation of learned pathway-ROI associations. (a) Correlation between elements of the mean association matrix $\bar{\mathbf{A}}$ learned by CALM using unpaired data and when learned using a paired baseline. (b) Cross-fold stability measured by the survival function of signal-to-noise ratio (SNR) across five folds.

weighted sum across SNPs. We capture genetic risk across multiple phenotypes by leveraging the effect sizes from $m = 6$ independent GWAS spanning disorder-related traits: ASD [26, 39], ADHD [3], educational attainment (EA) [24], intelligence [29], major depression disorder (MDD) [36], and epilepsy [33].

4 Experimental Results

Quantitative Disease Prediction: Table 1 reports the ASD vs. NC classification performance on the ACE dataset when the models are trained on ABIDE and SSC within a 5-fold CV setup. As seen, CALM achieves the best performance across most metrics, with statistically significant improvement over G-MIND, UNSEEN, and SUE. While below the significance threshold, we note the consistent improvement of CALM over its model ablations. This improvement can be attributed to the complementary roles of the two loss terms.




Fig. 3: (a) Bipartite graph showing the top pathway-ROI associations; edge thickness encodes association strength. (b) Gradient-based attribution of each GWAS disorder channel to the shared latent space, normalized to sum to 100%.

Association Validation: To validate that CALM learns pathway-ROI associations consistent with those obtained from true paired data, we replace the class-conditional distribution matching loss ($\mathcal{L}_{\text{CMMD}}$) with a paired contrastive loss (InfoNCE [25]) that directly maps modalities from the same subject to similar latent representations. Fig. 2(a) correlates the entries of the mean association matrix $\bar{\mathbf{A}}$ between the unpaired CALM and the paired baseline. The strong Pearson correlation ($r = 0.87$) indicates that CALM recovers consistent pathway-ROI associations without requiring paired data. We also assess cross-fold stability within the unpaired setting by computing the signal-to-noise ratio ($\text{SNR} = \mu/\sigma$) of each association across five folds. Fig. 2(b) shows that 72% of associations achieve $|\text{SNR}| > 1$ (stable) and 47% achieve $|\text{SNR}| > 2$ (reproducible), confirming that most associations are robust across data splits.

Learned Imaging-Genetics Associations: To determine the most salient pathway-brain ROI influences, we first rank pathways by their ℓ_2 column norm of the association matrix $\mathbf{A}$ and select the top four pathways with the strongest brain-wide effects. We then merge bilateral ROIs by summing the association values of their left and right counterparts. For visualization, we display the top four brain ROIs with the highest association strength for each pathway in Fig. 3.

The IgA Immune Network, Wnt Signaling, and Sulfur Metabolism pathways converge on occipital, temporal, and parietal regions, linked to atypical visual and multisensory processing in ASD [23]. These pathways are functionally related through homeostatic regulation of neurodevelopment mechanisms that are well-established in ASD pathology [27]. In contrast, NOD-like Receptor Signaling maps onto different ventral temporal and limbic areas, aligning with socioemotional and face processing deficits in ASD [12].

Channel Attribution Analysis: To quantify the contribution of each phenotypic trait to the shared latent space, we compute gradient-based input attribution [31] through the genetics encoder $\mathcal{E}_G$ and the learned projection matrix

$\mathbf{W}_G$, averaged across the top four pathways and five cross-validation folds. Fig. 3(b) shows that ASD, ADHD, EA, and IQ contribute most strongly, while MDD and epilepsy contribute moderately. This is consistent with the well-documented genetic and phenotypic overlap among these neurodevelopmental traits [21].

5 Conclusion

We have proposed CALM, a framework that learns interpretable cross-modal associations between brain ROI morphology and genetic pathways from *entirely unpaired and disjoint* populations. The backbone of CALM is a class-conditional linear transport model that aligns the group-wise distributions across modalities in the a shared latent space. We validate the learned associations by comparing them against a paired baseline trained on matched subjects, and we demonstrate the stability of the association matrix. CALM generalizes to a completely unseen external dataset and surpasses state-of-the-art baselines. The discovered pathway-ROI associations are supported by established ASD literature.

Acknowledgments. This work was supported by the National Institutes of Health awards 1R01HD108790 (PI Venkataraman).

Disclosure of Interests. All authors disclosed no relevant relationships.

References

1. 1000 Genomes Project Consortium: A global reference for human genetic variation. *Nature* **526**(7571), 68 (2015)
2. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems* **26** (2013)
3. Demontis, D., et al.: Genome-wide analyses of ADHD identify 27 risk loci, refine the genetic architecture and implicate several cognitive domains. *Nature genetics* **55**(2), 198–208 (2023)
4. Di Martino, A., 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)
5. Falk, C., et al.: Haplotype relative risks: an easy reliable way to construct a proper control sample for risk calculations. *Annals of human genetics* **51**(3), 227–233 (1987)
6. Fischbach, G.D., et al.: The simons simplex collection: a resource for identification of autism genetic risk factors. *Neuron* **68**(2), 192–195 (2010)
7. Fischl, B.: FreeSurfer. *Neuroimage* **62**(2), 774–781 (2012)
8. 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)
9. 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)

10. Gretton, A., et al.: A kernel two-sample test. *The journal of machine learning research* **13**(1), 723–773 (2012)
11. Hoshen, Y., et al.: Unsupervised correlation analysis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3319–3328 (2018)
12. Hughes, H.K., et al.: Immune dysfunction and autoimmunity as pathological mechanisms in autism spectrum disorders. *Frontiers in cellular neuroscience* **12**, 405 (2018)
13. Jack, A., et al.: A neurogenetic analysis of female autism. *Brain* **144**(6), 1911–1926 (2021)
14. 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)
15. Johnson, W.E., et al.: Adjusting batch effects in microarray expression data using empirical Bayes methods. *Biostatistics* **8**(1), 118–127 (2007)
16. Kanehisa, M., et al.: KEGG: kyoto encyclopedia of genes and genomes. *Nucleic acids research* **28**(1), 27–30 (2000)
17. Khosla, P., et al.: Supervised contrastive learning. *Advances in neural information processing systems* **33**, 18661–18673 (2020)
18. 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)
19. Li, W., et al.: Brainnetome atlas of preadolescent children based on anatomical connectivity profiles. *Cerebral Cortex* **33**(9), 5264–5275 (2023)
20. Li, Y., et al.: AlignMamba: Enhancing multimodal Mamba with local and global cross-modal alignment. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 24774–24784 (2025)
21. Litman, A., et al.: Decomposition of phenotypic heterogeneity in autism reveals underlying genetic programs. *Nature Genetics* **57**(7), 1611–1619 (2025)
22. Marees, A.T., et al.: A tutorial on conducting genome-wide association studies: Quality control and statistical analysis. *International journal of methods in psychiatric research* **27**(2), e1608 (2018)
23. Nickl-Jockschat, T., et al.: Brain structure anomalies in autism spectrum disorder—a meta-analysis of VBM studies using anatomic likelihood estimation. *Human brain mapping* **33**(6), 1470–1489 (2012)
24. Okbay, A., et al.: Polygenic prediction of educational attainment within and between families from genome-wide association analyses in 3 million individuals. *Nature genetics* **54**(4), 437–449 (2022)
25. Oord, A.v.d., et al.: Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018)
26. Pedersen, C.B., et al.: The iPSYCH2012 case-cohort sample: new directions for unravelling genetic and environmental architectures of severe mental disorders. *Molecular psychiatry* **23**(1), 6–14 (2018)
27. Pelphrey, K.A., et al.: Research review: constraining heterogeneity: the social brain and its development in autism spectrum disorder. *Journal of Child Psychology and Psychiatry* **52**(6), 631–644 (2011)
28. Purcell, S., et al.: PLINK: a tool set for whole-genome association and population-based linkage analyses. *The American journal of human genetics* **81**(3), 559–575 (2007)
29. Savage, J.E., et al.: Genome-wide association meta-analysis in 269,867 individuals identifies new genetic and functional links to intelligence. *Nature genetics* **50**(7), 912–919 (2018)

30. 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)
31. Shrikumar, A., et al.: Not just a black box: Learning important features through propagating activation differences. *arXiv preprint arXiv:1605.01713* (2016)
32. Smedley, N.F., et al.: Discovering and interpreting transcriptomic drivers of imaging traits using neural networks. *Bioinformatics* **36**(11), 3537–3548 (2020)
33. Stevelink, R., et al.: GWAS meta-analysis of over 29,000 people with epilepsy identifies 26 risk loci and subtype-specific genetic architecture. *Nature genetics* **55**(9), 1471–1482 (2023)
34. Wang, J., et al.: Learning explainable imaging-genetics associations related to a neurological disorder. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 347–357. Springer (2025)
35. 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)
36. Wray, N.R., et al.: Genome-wide association analyses identify 44 risk variants and refine the genetic architecture of major depression. *Nature genetics* **50**(5), 668–681 (2018)
37. Yacobi, A., et al.: Learning shared representations from unpaired data. *Advances in Neural Information Processing Systems* (2025)
38. Zhang, Y., et al.: Learning unseen modality interaction. *Advances in Neural Information Processing Systems* **36**, 54716–54726 (2023)
39. Zhou, X., et al.: Integrating de novo and inherited variants in 42,607 autism cases identifies mutations in new moderate-risk genes. *Nature genetics* **54**(9), 1305–1319 (2022)