# GG-AE: Genetically-Guided Autoencoder for Disentangling Structural Networks of Brain Atrophy

Marilena De Pian[1], Junhao Wen[2], and Christos Davatzikos[1]

[1] Artificial Intelligence in Biomedical Imaging Lab, University of Pennsylvania, Philadelphia PA 19104, USA
{depianm,Christos.Davatzikos}@pennmedicine.upenn.edu
[2] Laboratory of AI & Biomedical Science, Columbia University, New York City NY 10032, USA
jw4745@cumc.columbia.edu

**Abstract.** Neurodegenerative diseases exhibit substantial phenotypic heterogeneity arising from the interplay of genetic risk and non-genetic factors, which is often obscured in standard neuroimaging analyses. Although structural MRI captures brain atrophy, it reflects a final overlapping pathway of diverse etiological processes, making it difficult to disentangle genetically driven degeneration from environmental or lifestyle-related effects. We propose a genetically guided autoencoder (GG-AE), a self-supervised framework that leverages disease-associated genetic variation as a biological anchor to constrain the representation learning of brain atrophy. The framework first learns a compact genetic latent representation using a variational autoencoder, and subsequently decomposes voxel-wise imaging data into a genetically aligned subspace and a residual subspace using a dual-latent imaging autoencoder with non-negative, part-based decoding. On semi-synthetic data with known ground truth, GG-AE achieves improved recovery of genetically driven atrophy patterns compared to unguided and correlation-based baselines. When applied to real ADNI data, the model identifies spatially interpretable atrophy patterns and genetic associations that are consistent with established Alzheimer's disease findings, while isolating non-genetic variation in a residual component. These results demonstrate the potential of genetically guided representation learning for disentangling disease heterogeneity and enabling mechanistically grounded imaging–genetics analysis. Our code is available at https://github.com/marilenadp/GG-AE.

**Keywords:** Imaging-Genetics · Disentangled Representation Learning · Multimodal Fusion · Interpretable Deep Learning · Neurodegeneration

## 1 Introduction

Neurodegenerative diseases arise from a complex interplay of genetic risk and environmental exposures, leading to marked heterogeneity in clinical presentation and disease progression [24, 22]. Structural magnetic resonance imaging (MRI)

is the primary tool for quantifying neurodegeneration, yet brain atrophy represents a final overlapping pathway of diverse pathological processes [27, 17]. As a result, patients may exhibit similar spatial atrophy patterns driven by fundamentally different etiological mechanisms, such as genetic susceptibility versus environmental or lifestyle-related factors [8, 12]. Current neuroimaging analyses largely treat atrophy as a single phenotype, failing to distinguish its underlying sources and thereby limiting interpretability, biomarker discovery, and precision medicine efforts.

Most pattern recognition approaches in neuroimaging, including multivariate analyses and autoencoder-based models, operate in a purely unsupervised manner and rely solely on imaging covariance [18, 23]. While these methods can uncover dominant structural patterns, they lack biological constraints and inevitably entangle distinct sources of variation into mixed representations. Even recent deep generative learning models, despite improved representational capacity, often learn features that are statistically efficient but biologically ambiguous, hindering mechanistic interpretation [26].

Imaging–genetics approaches attempt to bridge this gap by linking genetic variation to brain structure [21, 24]. However, existing methods typically fall into two extremes. Voxel-wise or region-based association analyses offer statistical rigor but ignore the network-level organization of neurodegeneration [19, 28]. Conversely, many deep multimodal models symmetrically fuse genetic and imaging data, via concatenation, attention mechanisms, or correlation maximization, optimizing predictive performance at the expense of interpretability [7, 5, 16, 25, 9]. These approaches often obscure how genetic risk manifests spatially in the brain and implicitly assume bidirectional information flow between genetics and imaging, contrary to the biological pathway from genotype to phenotype.

To address these limitations, we propose a genetically guided self-supervised framework, GG-AE, that explicitly disentangles genetically driven brain atrophy from non-genetic variation. Rather than symmetrically fusing modalities or optimizing for disease classification, our approach uses disease-associated genetic variation as a biological anchor to constrain imaging representation learning. We first learn a compact representation of genetic risk using a variational autoencoder (VAE). We then train an imaging autoencoder (AE) with a dual latent space that separates a genetically aligned subspace from a residual subspace capturing environmental, lifestyle, or unexplained variability. By incorporating an Non-negative Matrix Factorization (NMF)-inspired decoder with non-negativity and soft orthogonality constraints, the model yields additive, spatially coherent, and biologically interpretable structural brain networks [18]. Our main contributions are the following:

1. An NMF-inspired AE that enables interpretable, part-based decomposition of voxel-wise MRI data while remaining scalable to high-dimensional imaging inputs.
2. An explicit disentanglement of imaging-derived NMF atrophy patterns into distinct biological sources rather than treating imaging variation as a single homogeneous phenotype.

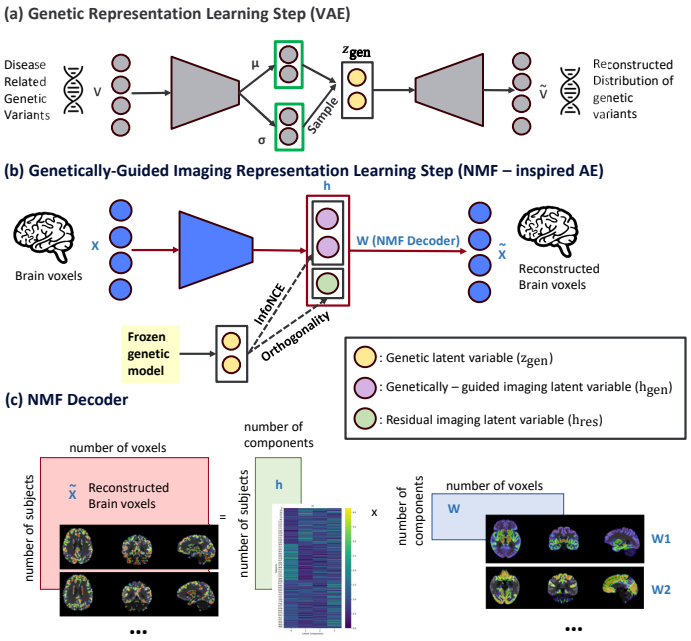


**Fig. 1.** Overview of the proposed GG-AE framework. (a) VAE for Genetic Representation Learning, (b) AE for Genetically Guided Imaging Representation Learning, and (c) NMF Decoder of the AE for non-negative, part-based imaging decoding.

3. A modular, genetically guided two-stage framework that mirrors the biological pathway from genetic risk to brain structure, supports genetic pretraining, and allows extension to additional biological risk factors.

## 2 Methodology

An overview of the proposed framework is shown in Fig. 1. Our objective is to learn a disentangled latent representation that separates genetically driven structural variation from residual non-genetic effects. Let $\mathcal{X}_{img} \in \mathbb{R}^{N \times V}$ and $\mathcal{X}_{snp} \in \mathbb{R}^{N \times G}$ denote subject-level imaging and genetic inputs, respectively. We formulate this problem as a genetically guided deep non-negative matrix factorization framework implemented using neural autoencoders.

### 2.1 Genetic Representation Learning

Genotype data $\mathcal{X}_{snp}$ are modeled using a VAE to obtain a compact representation of subject-level genetic variation. For each subject, the encoder maps $\mathbf{x}_{snp} \in \mathbb{R}^G$ to a latent distribution $q(\mathbf{z}_{gen} \mid \mathbf{x}_{snp})$ parameterized by a mean and

diagonal covariance. Latent variables $\mathbf{z}_{gen} \in \mathbb{R}^K$ are sampled, and the decoder reconstructs SNP counts with a binomial likelihood. The VAE is trained by minimizing the negative evidence lower bound (ELBO):

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{q(\mathbf{z}_{gen}|\mathbf{x}_{snp})}\left[-\log p(\mathbf{x}_{snp} \mid \mathbf{z}_{gen})\right] + \beta\,\text{KL}(q(\mathbf{z}_{gen} \mid \mathbf{x}_{snp}) \,\|\, \mathcal{N}(\mathbf{0}, \mathbf{I}))\,, \quad (1)$$

where $\beta$ controls Kullback-Leibler (KL) regularization and is gradually increased during training. Missing genotypes are handled by masking their contribution to the reconstruction loss.

### 2.2 Genetically Guided Imaging Representation Learning

After training, the genetic encoder $E_{gen}$ is fixed and used to obtain genetic embeddings $\mathbf{z}_{gen} = E_{gen}(\mathbf{x}_{snp}) \in \mathbb{R}^K$. An AE is then trained on voxel-wise imaging data $\mathbf{x}_{img} \in \mathbb{R}^V$. The imaging encoder maps each input to a latent representation explicitly partitioned into a genetically aligned subspace and a residual subspace: $[\mathbf{h}_{gen}, \mathbf{h}_{res}] = E_{img}(\mathbf{x}_{img})$, where $\mathbf{h}_{gen} \in \mathbb{R}^K$ is encouraged to align with $\mathbf{z}_{gen}$ and $\mathbf{h}_{res} \in \mathbb{R}^L$ captures remaining non-genetic variation.

To promote interpretability and part-based structure, non-negativity is enforced on latent activations and decoder weights. Reconstruction is given by: $\hat{\mathbf{x}}_{img} = \text{ReLU}(W\,[\,\sigma(\mathbf{h}_{gen}), \sigma(\mathbf{h}_{res})\,])$, where $\sigma(\cdot)$ denotes the sigmoid function and $W \in \mathbb{R}^{V\times(K+L)}$ is a linear decoder, yielding additive spatial components analogous to NMF.

### 2.3 Objective Function

To encourage disentanglement and interpretability, we apply reconstruction, alignment, orthogonality, and magnitude regularization losses, balancing them via scalar weights.

Imaging reconstruction is enforced using mean squared error, while a soft orthonormality constraint is applied to decoder weights to reduce redundancy while allowing biologically realistic overlap [18]:

$$\mathcal{L}_{\text{img}} = \frac{1}{N}\sum_{i=1}^{N}\left\|\mathbf{x}_{img}^{(i)} - \hat{\mathbf{x}}_{img}^{(i)}\right\|_2^2 \;+\; \lambda_W \left\|W^\top W - I\right\|_F^2. \quad (2)$$

Genetic and imaging representations are coupled using an InfoNCE contrastive loss [13] that aligns $\mathbf{h}_{gen}$ with $\mathbf{z}_{gen}$ at the subject level:

$$\mathcal{L}_{align} = -\frac{1}{B}\sum_{i=1}^{B}\log\frac{\exp\Big(\text{sim}(\mathbf{h}_{gen}^{(i)}, \mathbf{z}_{gen}^{(i)})/\tau\Big)}{\sum_{j=1}^{B}\exp\Big(\text{sim}(\mathbf{h}_{gen}^{(i)}, \mathbf{z}_{gen}^{(j)})/\tau\Big)}, \quad (3)$$

where $\text{sim}(\cdot,\cdot)$ denotes cosine similarity and $\tau$ is a temperature parameter. Disentanglement between genetic and residual subspaces is further encouraged via

latent orthogonality:

$$\mathcal{L}_{h\text{-ortho}} = \frac{1}{N}\sum_{i=1}^{N}\left\|\mathbf{h}_{gen}^{(i)\top}\mathbf{h}_{res}^{(i)}\right\|_2^2. \tag{4}$$

To prevent domination of either latent subspace and encourage balanced utilization, we apply $\ell_2$ regularization to both genetic ($\mathcal{L}_{l2gen}$) and residual ($\mathcal{L}_{l2res}$) latent activations. The full objective is:

$$\mathcal{L}_{total} = \mathcal{L}_{\text{img}} + \lambda_{align}\mathcal{L}_{align} + \lambda_h\mathcal{L}_{h\text{-ortho}} + \lambda_{l2}(\mathcal{L}_{l2gen} + \mathcal{L}_{l2res}). \tag{5}$$

## 3 Experiments

### 3.1 Imaging Data Preprocessing

For semi-synthetic and real-data experiments, T1-weighted MRI scans from UKBB [11] and ADNI [15] were used. Images were preprocessed using a standard pipeline: gray matter (GM) maps were extracted via multi-atlas segmentation [2, 3], skull-stripped, registered to MNI space using deformable registration [14], smoothed with an 8 mm FWHM Gaussian kernel, and downsampled to $2 \times 2 \times 2$ mm.

Voxel-wise imaging inputs were constructed as atrophy maps by $z$-score normalization relative to a cognitively normal (CN) reference cohort. Voxel-wise CN means and standard deviations were used to normalize both CN and diseased subjects. Low-variance voxels were masked to suppress background noise, and negative values were inverted such that higher intensities correspond to greater atrophy.

### 3.2 Semi-Synthetic Experiments

**Data Construction** Semi-synthetic datasets were generated by imposing simulated disease-related patterns onto real CN imaging data from UKBB [11], preserving realistic anatomical variability ($N = 2354$). Ground-truth atrophy patterns were defined using atlas-based ROI masks. Two genetic patterns, $A_{gen1}$ and $A_{gen2}$, were constructed from temporal+subcortical and frontal+occipital ROIs, respectively, while an environmental confound $A_{env}$ combined temporal+parietal ROIs (Fig. 2a). Pattern-specific severity coefficients were sampled per subject, with genetically assigned subjects exhibiting dominant coefficients for the corresponding pattern and low values for others. Within each mask, voxel intensities were randomly reduced by 0–30% to simulate variable atrophy. Subjects assigned to $A_{env}$ shared imaging effects without shared genetic risk, while control subjects exhibited no imposed atrophy. Imaging inputs consisted of voxel-wise atrophy maps derived via CN-based $z$-scoring, as described in the previous section.

Genetic data were fully simulated as minor allele counts (0/1/2) across $G = 100$ SNPs. Baseline minor allele frequencies (MAFs) were sampled from distribution Beta$(0.8, 3.0)$ and scaled to $[0, 0.5]$. For each genetically driven subgroup, five SNPs were designated as risk-associated by increasing their MAFs relative to baseline, while remaining SNPs followed the baseline distribution.

**Evaluation and Experimental Procedure** Unless otherwise specified, experiments used $K = 2$ genetically aligned patterns and a residual dimensionality of $L = 1$. The genetic VAE was trained on all available genetic data ($N = 2354$) to reflect realistic large-scale genetic cohorts, while the imaging AE was trained on a balanced subset ($N = 1000$) to match the scale of typical imaging–genetics overlap. Performance was evaluated using complementary subject- and spatial-level metrics: subject-level disentanglement was assessed via Pearson correlations between latent activations and ground-truth genetic ($s_{gen1}, s_{gen2}$) and environmental ($s_{env}$) severity variables, while spatial alignment was quantified using cosine similarity between learned decoder patterns and ground-truth atrophy maps. To verify that observed alignment was not driven by architectural bias, a negative-control experiment with permuted genetic correspondence was performed (Perm GG-AE). Robustness was further assessed by increasing residual dimensionality while monitoring stability of genetic spatial alignment, and a Deep Canonical Correlation Analysis (CCA) variant replacing the InfoNCE loss was included to isolate the effect of the alignment objective. Qualitative visualization of spatial maps was used to assess correspondence and leakage between genetic and residual components.

### 3.3 Real-Data Experiments

The framework was applied to ADNI [15]. MRI data is pre-processed as described in previous section. Whole-genome sequencing data were processed following standard quality control [21, 20], and 54 Alzheimer's disease–related SNPs were selected via manual review of the NHGRI–EBI GWAS Catalog [1]. Only AD/MCI subjects with available genetic data were retained ($N = 975$). Voxel-wise atrophy maps were constructed for all AD/MCI subjects using CN-based $z$-score normalization. In the absence of ground truth, the number of genetic patterns was selected based on stability of the genetic latent space: for $K = 1$–$5$, the genetic VAE was trained five times per setting, and the largest $K$ yielding consistent subject stratification across runs was chosen ($K = 3$). The full framework was then applied with $K = 3$ genetically aligned patterns and one residual component ($L = 1$). Genetic specificity was evaluated via SNP–latent association testing with false discovery rate correction, and spatial maps were qualitatively inspected to assess biological interpretability.

## 4 Results & Discussion

### 4.1 Results on Semi-Synthetic Data

On semi-synthetic data with known ground truth, the proposed model consistently recovered genetically driven imaging components aligned with the imposed disease patterns. As shown in Table 1, our method achieved the highest spatial cosine similarity of the genetically-aligned learned patterns with the ground-truth genetic patterns ($A_{gen1}$ and $A_{gen2}$), while simultaneously exhibiting the

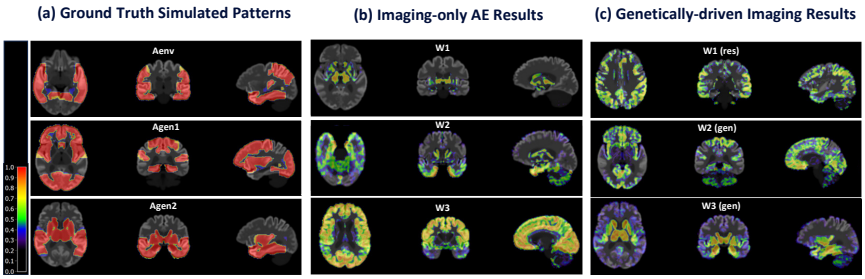


**Fig. 2.** Qualitative comparison of spatial components on semi-synthetic data: (a) ground-truth simulated genetic atrophy patterns, (b) proposed genetically guided model, (c) imaging-only autoencoder

**Table 1.** Quantitative evaluation on semi-synthetic data. For each model and ground-truth genetic pattern ($A_{gen1}$, $A_{gen2}$), spatial cosine similarity between learned components ($W_1$–$W_3$) and ground-truth masks, and Pearson correlations with simulated genetic ($s_{gen1}$, $s_{gen2}$) and environmental ($s_{env}$) severities is reported. Bold denotes highest spatial alignment for each model; underlined values indicate strongest genetic association. "Gen" and "Res" label genetically aligned and residual components.

| **Model** | **GT Pattern** | **Spatial Cosine Similarity ↑** | | | **Pearson Corr ↑** | | |
|---|---|---|---|---|---|---|---|
| | | W1 | W2 | W3 | sgen1 | sgen2 | senv |
| NMF | Agen1 | 0.69 | **0.72** | 0.60 | 0.34 | -0.27 | 0.10 |
| | Agen2 | **0.84** | 0.57 | 0.60 | -0.55 | 0.59 | -0.06 |
| Imaging AE | Agen1 | **0.69** | 0.53 | 0.31 | 0.34 | -0.18 | 0.04 |
| | Agen2 | 0.20 | 0.21 | **0.85** | 0.18 | -0.15 | 0.06 |
| | | W1 (res) | W2 (gen) | W3 (gen) | | | |
| Deep CCA | Agen1 | 0.54 | 0.19 | **0.56** | -0.09 | -0.13 | 0.10 |
| | Agen2 | **0.71** | 0.59 | 0.29 | 0.56 | -0.50 | -0.02 |
| Perm GG-AE | Agen1 | **0.74** | 0.26 | 0.27 | -0.04 | 0.12 | -0.00 |
| | Agen2 | 0.34 | **0.78** | 0.77 | -0.07 | 0.11 | -0.03 |
| GG-AE | Agen1 | 0.48 | 0.22 | **0.69** | <u>0.60</u> | -0.46 | 0.02 |
| | Agen2 | 0.57 | **0.76** | 0.38 | -0.47 | <u>0.61</u> | -0.05 |

strongest subject-level correlations with the corresponding genetic severity variables ($s_{gen1}$ and $s_{gen2}$). Imaging-only models and standard NMF recovered partially overlapping spatial patterns but showed weak or inconsistent associations with the severity coefficients, indicating limited disentanglement without genetic information. Deep CCA captured some genetic correlations but failed to consistently localize the correct spatial patterns. Negative control experiments (Perm GG-AE) yielded near-zero genetic correlations, confirming that alignment arises from meaningful genetic-imaging coupling rather than architectural bias. Increasing residual dimensionality did not alter spatial localization or genetic alignment, indicating that additional residual capacity absorbs non-genetic variability without contaminating the genetic subspace.

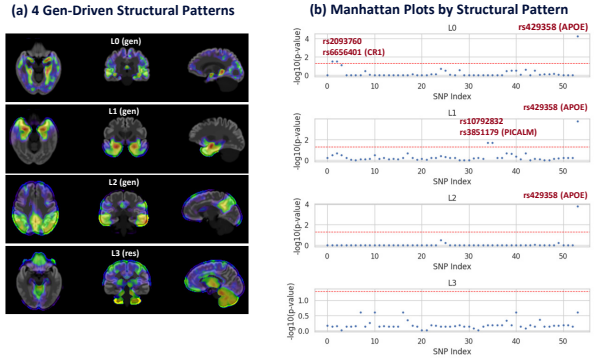


**Fig. 3.** (a) Spatial patterns learned on ADNI data and (b) their genetic associations.

Qualitative comparisons in Fig. 2 further support these findings. While the imaging-only autoencoder recovers broad spatial structure with substantial mixing between genetic and non-genetic regions, the proposed model produces spatial maps that closely resemble the ground-truth genetic patterns, capturing core affected regions and assigning genetically driven effects to the corresponding latent variables.

### 4.2 Results on Real Data

On the ADNI dataset, the proposed framework identified three genetically informed imaging patterns and one residual component (Fig. 3). The first genetic pattern shows widespread atrophy across frontal, parietal, occipital, and temporal regions with prominent medial temporal involvement and is associated with *APOE* and *CR1* variants, implicating lipid metabolism and immune pathways [10, 4]. The second pattern captures focal medial temporal atrophy centered on the hippocampus, amygdala, and entorhinal cortex and is preferentially associated with *APOE* and *PICALM*, genes involved in amyloid processing and synaptic vesicle trafficking [10, 24]. The third pattern highlights the precuneus and posterior cingulate cortex and is primarily associated by the *APOE* $\epsilon$4-defining variant, reflecting a global genetic risk component affecting default-mode network regions [10, 6]. In contrast, the residual component exhibits diffuse cerebellar, subcortical, and frontal involvement with no significant genetic association, indicating effective separation of non-genetic variability.

## 5 Conclusion

We proposed a genetically guided framework, GG-AE, to disentangle genetically driven brain atrophy from non-genetic variation in neuroimaging data. On semi-synthetic data, the method demonstrates improved disentanglement compared to imaging-only and correlation-based baselines, while on real ADNI data it yields

spatial patterns and genetic associations that are consistent with established Alzheimer's disease findings. This framework enables mechanistically grounded imaging-genetics analysis without relying on supervised prediction.

**Acknowledgments.** This work was partially supported by NIH grants RF1MH134236, R01AG054409, and U01AG068057.

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

## References

1. Buniello, A., Macarthur, J.A., Cerezo, M., Harris, L.W., Hayhurst, J., Malangone, C., et al.: The nhgri-ebi gwas catalog of published genome-wide association studies, targeted arrays and summary statistics 2019. Nucleic acids research **47**, D1005–D1012 (1 2019). https://doi.org/10.1093/nar/gky1120
2. Davatzikos, C., Genc, A., Xu, D., Resnick, S.M.: Voxel-based morphometry using the ravens maps: methods and validation using simulated longitudinal atrophy. NeuroImage **14**, 1361–1369 (2001). https://doi.org/10.1006/nimg.2001.0937
3. Doshi, J., Erus, G., Ou, Y., Resnick, S.M., Gur, R.C., Gur, R.E., et al.: Muse: Multi-atlas region segmentation utilizing ensembles of registration algorithms and parameters, and locally optimal atlas selection. NeuroImage **127**, 186–195 (2 2016). https://doi.org/10.1016/j.neuroimage.2015.11.073
4. Karch, C.M., Goate, A.M.: Alzheimer's disease risk genes and mechanisms of disease pathogenesis. Biological Psychiatry **77**, 43–51 (1 2015). https://doi.org/10.1016/j.biopsych.2014.05.006
5. Ko, W., Jung, W., Jeon, E., Suk, H.I.: A deep generative-discriminative learning for multimodal representation in imaging genetics. IEEE transactions on medical imaging **41**, 2348–2359 (9 2022). https://doi.org/10.1109/TMI.2022.3162870
6. Kuang, L., Jia, J., Zhao, D., Xiong, F., Han, X., Wang, Y.: Default mode network analysis of apoe genotype in cognitively unimpaired subjects based on persistent homology. Frontiers in Aging Neuroscience **12**, 545755 (6 2020). https://doi.org/10.3389/fnagi.2020.00188
7. Lee, S., Lee, J., Lee, M., Choi, J., Kang, K., Kim, Y.: Integrating genetic information for early alzheimer's diagnosis through mri interpretation. BHI 2023 - IEEE-EMBS International Conference on Biomedical and Health Informatics, Proceedings (2023). https://doi.org/10.1109/BHI58575.2023.10313442
8. Lenroot, R.K., Giedd, J.N.: The changing impact of genes and environment on brain development during childhood and adolescence: Initial findings from a neuroimaging study of pediatric twins. Development and psychopathology **20**, 1161 (2008). https://doi.org/10.1017/S0954579408000552
9. Li, Y., Niu, D., Qi, K., Liang, D., Long, X.: An imaging and genetic-based deep learning network for alzheimer's disease diagnosis. Frontiers in Aging Neuroscience **17**, 1532470 (2025). https://doi.org/10.3389/fnagi.2025.1532470
10. Liu, C.C., Kanekiyo, T., Xu, H., Bu, G.: Apolipoprotein e and alzheimer disease: risk, mechanisms and therapy. Nature reviews. Neurology **9**, 106–118 (2 2013). https://doi.org/10.1038/nrneurol.2012.263

11. Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub, E., Xu, J., et al.: Multimodal population brain imaging in the uk biobank prospective epidemiological study. Nature Neuroscience **19**, 1523–1536 (10 2016). https://doi.org/10.1038/nn.4393
12. Murray, M.E., Graff-Radford, N.R., Ross, O.A., Petersen, R.C., Duara, R., Dickson, D.W.: Neuropathologically defined subtypes of alzheimer's disease with distinct clinical characteristics: A retrospective study. Lancet neurology **10**, 785 (9 2011). https://doi.org/10.1016/S1474-4422(11)70156-9
13. van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding (1 2019)
14. Ou, Y., Sotiras, A., Paragios, N., Davatzikos, C.: Dramms: Deformable registration via attribute matching and mutual-saliency weighting. Medical Image Analysis **15**, 622–639 (8 2011). https://doi.org/10.1016/j.media.2010.07.002
15. Petersen, R.C., Aisen, P.S., Beckett, L.A., Donohue, M.C., Gamst, A.C., Harvey, D.J., et al.: Alzheimer's disease neuroimaging initiative (adni): clinical characterization. Neurology **74**, 201–209 (2010). https://doi.org/10.1212/WNL.0b013e3181cb3e25
16. Qiang, Y.R., Zhang, S.W., Li, J.N., Li, Y., Zhou, Q.Y.: Diagnosis of alzheimer's disease by joining dual attention cnn and mlp based on structural mris, clinical and genetic data. Artificial Intelligence in Medicine **145**, 102678 (11 2023). https://doi.org/10.1016/j.artmed.2023.102678
17. Skampardoni, I., Nasrallah, I.M., Abdulkadir, A., Wen, J., Melhem, R., Mamourian, E., et al.: Genetic and clinical correlates of ai-based brain aging patterns in cognitively unimpaired individuals. JAMA Psychiatry **81**, 456–467 (5 2024). https://doi.org/10.1001/jamapsychiatry.2023.5599
18. Sotiras, A., Resnick, S.M., Davatzikos, C.: Finding imaging patterns of structural covariance via non-negative matrix factorization. NeuroImage **108**, 1–16 (3 2015). https://doi.org/10.1016/j.neuroimage.2014.11.045
19. Stein, J.L., Hua, X., Lee, S., Ho, A.J., Leow, A.D., Toga, A.W., et al.: Voxelwise genome-wide association study (vgwas). NeuroImage **53**, 1160–1174 (11 2010). https://doi.org/10.1016/j.neuroimage.2010.02.032
20. Wen, J., Fu, C.H., Tosun, D., Veturi, Y., Yang, Z., Abdulkadir, A., et al.: Characterizing heterogeneity in neuroimaging, cognition, clinical symptoms, and genetics among patients with late-life depression. JAMA psychiatry **79**, 464–474 (5 2022). https://doi.org/10.1001/jamapsychiatry.2022.0020
21. Wen, J., Nasrallah, I.M., Abdulkadir, A., Satterthwaite, T.D., Yang, Z., Erus, G., et al.: Genomic loci influence patterns of structural covariance in the human brain. Proceedings of the National Academy of Sciences of the United States of America **120** (2023). https://doi.org/10.1073/pnas.2300842120
22. Wen, J., Skampardoni, I., Tian, Y.E., Yang, Z., Cui, Y., Erus, G., et al.: Neuroimaging endophenotypes reveal underlying mechanisms and genetic factors contributing to progression and development of four brain disorders. Nature Biomedical Engineering 2025 9:11 **9**, 1920–1937 (6 2025). https://doi.org/10.1038/s41551-025-01412-w
23. Wen, J., Varol, E., Sotiras, A., Yang, Z., Chand, G.B., Erus, G., et al.: Multi-scale semi-supervised clustering of brain images: deriving disease subtypes. Med. Image Anal. **75** (1 2021). https://doi.org/10.1016/j.media.2021.102304
24. Wen, J., Yang, Z., Nasrallah, I.M., Cui, Y., Erus, G., Srinivasan, D., et al.: Genetic and clinical correlates of two neuroanatomical ai dimensions in the alzheimer's disease continuum. Translational Psychiatry 2024 14:1 **14**, 420– (10 2024). https://doi.org/10.1038/s41398-024-03121-5

25. Wu, S., Venkataraman, A., Ghosal, S.: Girus-net: A multimodal deep learning model identifying imaging and genetic biomarkers linked to alzheimer's disease severity. Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference **2023** (2023). https://doi.org/10.1109/EMBC40787.2023.10341000
26. Yang, Z., Wen, J., Davatzikos, C.: Surreal-gan:semi-supervised representation learning via gan for uncovering heterogeneous disease-related imaging patterns. ICLR 2022 - 10th International Conference on Learning Representations (5 2022)
27. Yang, Z., Wen, J., Erus, G., Govindarajan, S.T., Melhem, R., Mamourian, E., et al.: Brain aging patterns in a large and diverse cohort of 49,482 individuals. Nature medicine **30**, 3015–3026 (10 2024). https://doi.org/10.1038/s41591-024-03144-x
28. Zhang, Y., Zhan, L., Thompson, P.M., Huang, H.: Biological knowledge guided deep neural network for brain genotype-phenotype association study. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) **11846 LNCS**, 84–92 (2019). https://doi.org/10.1007/978-3-030-33226-6_10