Integrating morphology and gene expression of neural cells in unpaired single-cell data using GeoAdvAE

Jinqiu Turbo Du[1], Kevin Z Lin[1]

University of Washington, Seattle WA 98107, USA, {turbodu,kzlin}@uw.edu,

Abstract. Cellular morphological transitions are widely observed in many diseases; however, the functional role of these morphologies remains unclear, as most technologies cannot profile both form and function simultaneously. However, computationally linking single-cell morphology and neural cell transcriptomics is challenging due to a lack of feature correspondences. We present GeoAdvAE, a geometry-aware adversarial autoencoder for diagonal (unpaired) integration of single-cell morphology and single-cell RNA sequencing. GeoAdvAE combines modality-specific variational autoencoders with a Gromov-Wasserstein regularizer and an adversarial discriminator to embed unpaired morphologies and transcriptomes into a shared latent space, preserving both reconstruction fidelity and cross-modal geometry. To validate the correctness of integration, we leverage Patch-seq neurons with joint morphology-RNA measurements. Using these ground-truth pairings, GeoAdvAE achieves the best cross-modal cell-type matching accuracy among diagonal integration methods, outperforming optimal transport, latent alignment, and adversarial baselines. We then apply GeoAdvAE to microglia from the 5xFAD mouse model, a model system of Alzheimer's disease. We integrate 98 CAJAL-quantified morphologies, spanning amoeboid and ramified forms, with 31,948 single-cell RNA-seq profiles across homeostatic, proliferating, and disease-associated states to recover a one-dimensional axis that aligns the two modalities. We uncover novel biology using integrated gradient attribution, highlighting transcriptomic shifts (DNA repair in ramified; cell killing in amoeboid) and nominating gene markers (*Ms4a6b*; *Ftl1*/*Fth1*) corresponding to morphological changes. Our integration also enables us to identify DAM signatures that do not correspond to morphological changes. GeoAdvAE provides a scalable and interpretable approach to connecting cellular "form" and "function" when joint profiling of morphology and transcriptomics is impractical. Our method is publicly available at https://github.com/turbodu222/GeoAdVAE.

Keywords: Alzheimer's disease, Cross-modal alignment, Deep learning, Diagonal integration, Patch-seq neurons, Microglia

1 Introduction

"Form follows function" is a powerful organizing principle in biology: a cell's shape is often a visible readout of what it is doing. In the brain, for example, neurons in different layers exhibit distinct arborization patterns that enable specialized circuit roles [47,17,5]. As another example, during the progression of Alzheimer's disease (AD), microglia remodel their processes and shift their morphology as they surveil tissue, respond to damage, and interact with pathology. Yet morphology is not a perfect proxy for cellular function [56,9]. Similar shapes can conceal divergent molecular states, and conversely, functional reprogramming does not always manifest as a clear change in shape. This gap complicates efforts to understand how microglia respond to disease burden and limits our ability to translate visible phenotypes into mechanistic insight [45].

Molecular profiling provides a complementary view. High-throughput measurements of RNA, proteins, and chromatin capture internal programs and environmental responses at single-cell resolution. However, for many questions, such as the nervous system, where cell morphology is intricate and tightly linked to function, simultaneous measurement of detailed morphology and transcriptomes remains rare. As a result, most available datasets are *unpaired*: large imaging collections with reconstructed shapes and, separately, large single-cell RNA-seq atlases. Unlocking the biology requires principled methods that can connect these modalities without joint measurements.

Integrating morphology with gene expression poses challenges distinct from those of conventional multi-omics alignment. Only a small fraction of genes directly influence morphological features, so information is intrinsically imbalanced across modalities [58,33]. Unlike RNA–ATAC or RNA–protein, there is no straightforward feature-to-feature correspondence to anchor alignment. Moreover, morphology must first be converted into quantitative descriptors that respect geometric relationships, and naïve embeddings can distort them. Together, these issues make diagonal (unpaired) integration between morphology and transcriptomics both necessary and non-trivial.

We introduce **GeoAdvAE**, a geometry-aware adversarial autoencoder for diagonal integration of single-cell morphology and gene expression. GeoAdvAE learns a shared latent space from unpaired datasets using four complementary ingredients: modality-specific autoencoders, an adversarial objective with a Gromov–Wasserstein regularizer for alignment, and a biological prior for orientation. The result is a joint representation that mixes cells across modalities while maintaining biologically meaningful neighborhoods. We validate GeoAdvAE using Patch-seq neurons, where matched morphology–RNA measurements provide ground truth for cross-modal cell-type matching, and then deploy the framework to microglia in the 5xFAD mouse model of AD. By probing the fitted model, we highlight gene programs associated with morphological transitions and identify transcriptomic signatures that are decoupled from visible shape changes.

1.1 Biological relevance

Alzheimer's disease (AD) lacks a cure [66], motivating strategies that enhance resilience – the capacity of cells and circuits to preserve function despite pathology [11]. Microglia are central to this idea: they adopt distinct morphologies that can encircle amyloid plaques and form barrier-like structures associated with tissue protection, even as dystrophic microglia accumulate with aging and may signal impaired responses [56,25]. Much of this knowledge comes from microscopy, but direct links between these visible phenotypes and underlying molecular programs remain limited. To close that gap, we apply GeoAdvAE to the 5xFAD mouse, an amyloid-only transgenic model with reduced between-sample variability. This setting lets us cleanly interrogate RNA–morphology coupling: even if mouse models do not fully reproduce human resilience or dystrophic signatures, they robustly recapitulate a homeostatic-to-inflammatory microglial continuum (historically framed as M2–M1) [23]. By aligning morphology with transcriptomic state in this controlled context, GeoAdvAE provides a roadmap for dissecting putative resilience mechanisms and prioritizing hypotheses for testing in human tissue.

1.2 Related work and challenges in integrating morphology and gene expression

Our work addresses the key obstacle that integrating cellular morphology with gene expression is substantially different and arguably more difficult than conventional single-cell integration of two different types, such as RNA-ATAC or RNA-protein integration. We highlight two main reasons for this.

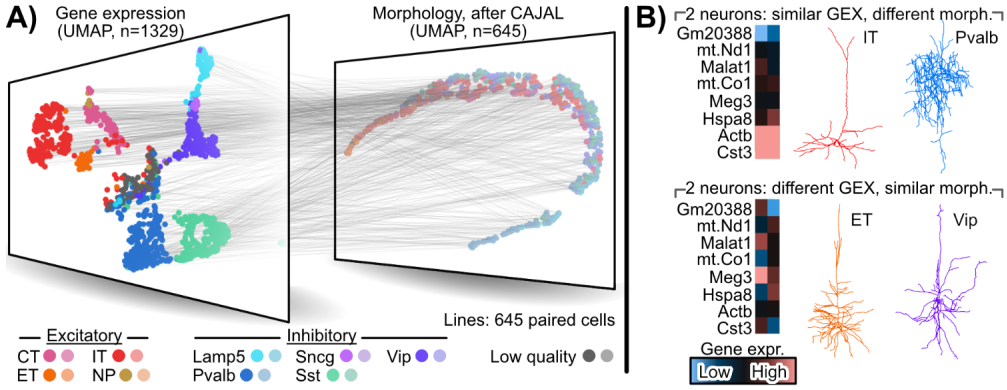

Fig. 1. Patch-seq data of cell GEX and morphology. A) Patch-seq data of 1329 neurons [47], of which 645 cells are measured in both gene expression (GEX) and morphology (quantified by CAJAL [17]) among different subtypes of excitatory and inhibitory neurons. The gray lines reveal the complex relation between the two modalities. B) Exemplary genes and morphologies to demonstrate that despite neurons having similar gene expression, they could have vastly different morphologies (top) or vice-versa (bottom).

- **Imbalanced information and no clear correspondence**: Only a small subset of genes influences cell shape (e.g., cytoskeleton, membrane dynamics), so morphology–GEX integration is intrinsically asymmetric and has a low signal-to-noise ratio. This is exemplified in Figure 1, which illustrates Patch-seq neurons whose GEX and morphology are measured simultaneously. Neurons with very similar GEX could have vastly different morphologies, and vice versa. By contrast, GEX–ATAC and GEX–protein have natural anchors (i.e., regulatory links and protein readouts) supporting feature-level correspondence [7]. Methods built for such settings (e.g., SCOT [12], cross-modal autoencoders [64]) tend to assume symmetric information and therefore underperform on morphology–GEX. Morphology emerges from many pathways rather than single genes, precluding one-to-one anchors and weakening co-occurrence/paired-feature strategies [7]. We note that quantifying morphology itself is also nontrivial. Tools like CAJAL [18] and MorphOMICs [9] yield embeddings, but links to transcriptional programs remain limited.
- **Lack of datasets that measure both genes and morphology of neural (brain) cells**: Despite progress in multi-modal integration, linking single-cell morphology with gene expression remains difficult because paired datasets are rare. We note that there is a plethora of work that studies the relationship between gene expression and morphology using spatial transcriptomics [6,27]. However, these spatial transcriptomics platforms lack the resolution to capture the intricate morphology of neural cells [44]. For example, while the latest versions of 10x Visium HD can capture spots of up to 2 μm, this resolution is still not apt for capturing the intricate morphologies of microglia processes that are typically only 1 μm thick [61]. Patch-seq [17] is one of the few wet-bench protocols that enable accurate profiling of both gene expression and brain cell morphology, but its low throughput contrasts with the scale of scRNA-seq and microscopy, which are typically generated separately. Consequently, new frameworks for effective inference from large, unpaired single-cell gene expression and morphology datasets are needed.

2 Method

To integrate unpaired cellular morphology and gene expression data, we propose **GeoAdvAE** (geometry-aware adversarial autoencoder), a diagonal integration framework that learns joint representations of both modalities in a shared latent space (Fig. 2A). GeoAdvAE consists of two modality-specific Variational Autoencoders (VAEs) and one discriminator: the two encoders extract features from morphology and gene expression, respectively, and project them into a common latent distribution. Unlike conventional VAEs that rely solely on reconstruction and KL-divergence losses, GeoAdvAE introduces three additional and complementary penalty terms (adversarial alignment, Gromov–Wasserstein-based structural regularization, and a coarse biological cluster-level prior) that cooperatively guide the training process toward desirable cross-modal alignment. Together, these components enable the model to preserve modality-specific fidelity while

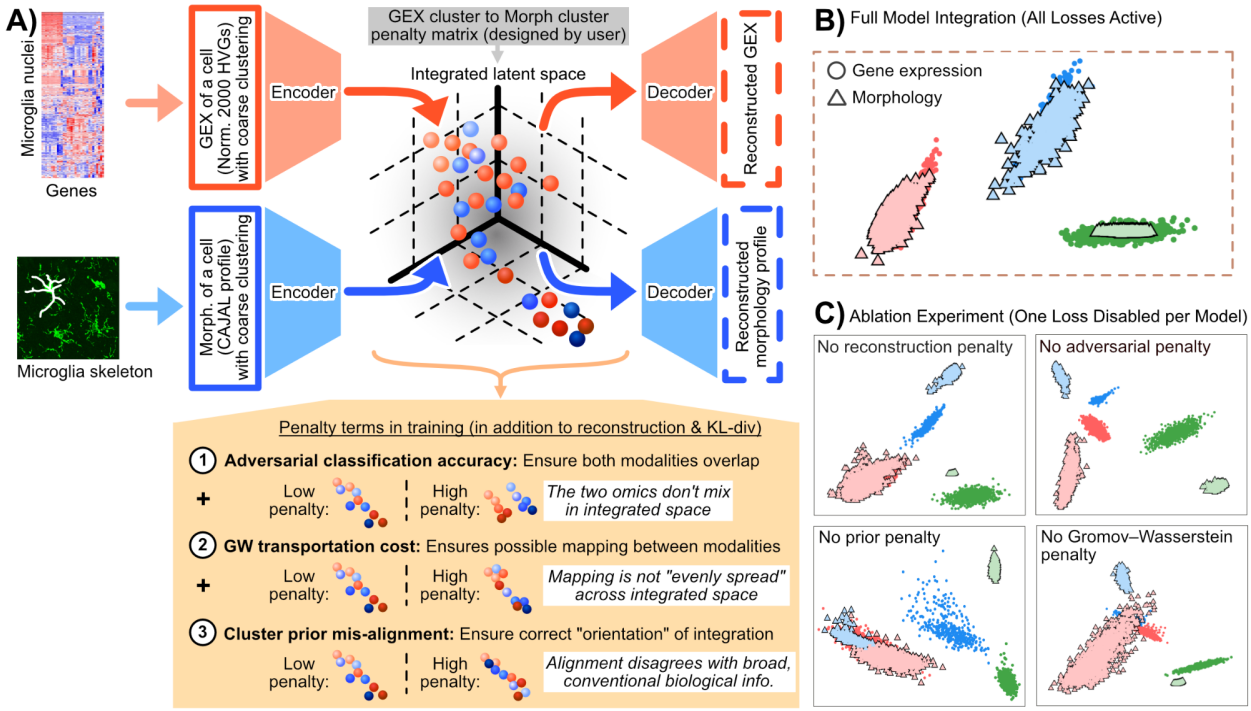

Fig. 2. Architecture and ablation study. A) Overview of the GeoAdvAE architecture integrating unpaired morphology (microglial skeleton; white) and gene expression through two modality-specific encoders and decoders in a shared latent space. Three additional loss terms (adversarial alignment, Gromov–Wasserstein (GW) transport regularization, and cluster-prior guidance) jointly control modality overlap, geometric consistency, and biologically meaningful orientation, respectively. The white boxes denote qualitative statements on how to interpret potential integrations with high penalty terms (i.e., a not desirable integration). B) Integrated latent space from simulation data computed by GeoAdvAE. C) Ablation experiment (one penalty disabled for each fit), where removing any of the penalty term leads to distinct distortions in the integrated latent space. See Appendix S4 for a full Shapley value analysis.

enforcing semantic and geometric consistency across unpaired single-cell modalities, forming the foundation for integrating morphology and transcriptomic information in the subsequent stages of our framework.

For GeoAdvAE to integrate RNA and morphology, we first need to meaningfully represent a cell's morphology as a vector of numerical values. To do this, we use CAJAL, which represents each microglia's morphology as a low-dimensional vector based on the Wasserstein distance to other cells' morphologies. Compared with other methods for quantifying morphology, we have empirically found that CAJAL is the most suitable for our proposed diagonal integration framework (see Fig. S9).

2.1 Model architecture

Encoder and decoder architecture. Each modality uses a VAE. The gene expression profiles ($\mathbf{x}_B \in \mathbb{R}^{2000}$, for the 2000 highly variable genes, log-normalized) are organized into three layers. Both encoders output the mean and log-variance of a $d = 16$ integrated latent space. The morphology profiles ($\mathbf{x}_A \in \mathbb{R}^{30}$, quantified via CAJAL [17] as described in Section 3) pass through two hidden layers. Using the reparameterization trick, let $\mathbf{z}_A$ and $\mathbf{z}_B$ be sampled from the posterior Gaussian distributions based on the gene expression and morphology encoders using $\mathbf{x}_A$ and $\mathbf{x}_B$ as inputs, respectively. Decoders mirror the encoders back to the input size with the same blocks; the final layer is linear to preserve real-valued reconstructions.

Discriminator architecture. To align the latent spaces across modalities, we add a discriminator $D : \mathbb{R}^d \to [0, 1]$ is a three-layer multi-layer perceptron. The encoders are updated adversarially to fool D, yielding a shared, modality-invariant latent.

2.2 Loss Function Formulation for Cross-Modal Alignment

We align gene expression (GEX, modality A) and CAJAL morphology embeddings (modality B) in a shared latent space through a weighted sum of five complementary objectives:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{recon}} \cdot \mathcal{L}_{\text{recon}} + \lambda_{\text{KL}} \cdot \mathcal{L}_{\text{KL}} + \lambda_{\text{GAN}} \cdot \mathcal{L}_{\text{GAN}} + \lambda_{\text{GW}} \cdot \mathcal{L}_{\text{GW}} + \lambda_{\text{prior}} \cdot \mathcal{L}_{\text{prior}}, \tag{1}$$

with hyperparameters $\lambda_{\text{KL}}, \lambda_{\text{recon}}, \lambda_{\text{prior}}, \lambda_{\text{GAN}}, \lambda_{\text{GW}} \geq 0$. Our default choices are $\lambda_{\text{KL}} = 0.1$, $\lambda_{\text{recon}} = 5$, $\lambda_{\text{prior}} = 1$, $\lambda_{\text{GAN}} = 4$, and $\lambda_{\text{GW}} = 1$. Training follows a curriculum: autoencoding ($\mathcal{L}_{\text{KL}}$+$\mathcal{L}_{\text{recon}}$) is active from the start; the adversarial term ($\mathcal{L}_{\text{GAN}}$) is enabled after an initial warmup so that encoders learn stable reconstructions first; the prior-guided semantic alignment ($\mathcal{L}_{\text{prior}}$) ramps up via a schedule; finally, the GW structure term ($\mathcal{L}_{\text{GW}}$) is enabled after the latent geometry stabilizes. As our results in Section 4 will emphasize, integration that focuses solely on reconstruction is insufficient to integrate GEX and morphology. The addition of these other penalty terms encourages GeoAdvAE to prioritize other qualities that enable meaningful biological discovery beyond what existing methods optimize for.

Reconstruction loss and KL divergence. For each modality, the encoder parameterizes a diagonal-Gaussian posterior; we apply the Kullback–Leibler (KL) divergence penalty against a unit-normal prior to regularize the encoder. Each modality's decoder reconstructs its original input; we use an L_1 loss (Manhattan distance) to measure the reconstruction quality for robustness in high dimensions. Both the reconstruction and KL divergence penalties are summed across both modalities.

Adversarial classification based on the discriminator to ensure integration. Leveraging our discriminator, we include an adversarial classification term to encourage the embeddings from both modalities to overlap. This ensures mixing between the two modalities. Specifically, let q_A and q_B denote the posterior distributions produced by E_A and E_B, respectively, and $\mathbf{z}_A$ and $\mathbf{z}_B$ be samples from these distributions. The discriminator $D : \mathbb{R}^d \to [0, 1]$ tries to identify the modality based on the samples $\mathbf{z}_A$ and $\mathbf{z}_B$, while encoders try to produce posterior distributions that fool the discriminator.

To train, we alternate between training the discriminator and the generator. The discriminator minimizes the following classification loss with a fixed hyperparameter $\epsilon = 0.1$ for stability, where a prediction closer to 0 or 1 means the discriminator predicts the cell to originate from modality A or B, respectively.

$$\textbf{Discriminator loss}: \quad \mathcal{L}_D = \mathbb{E}_{\mathbf{z}_A \sim q_A}\left[(D(\mathbf{z}_A) - \epsilon)^2\right] + \mathbb{E}_{\mathbf{z}_B \sim q_B}\left[(D(\mathbf{z}_B) - (1-\epsilon))^2\right]. \tag{2}$$

Training the discriminator D (but keeping the encoders fixed, and hence, the embeddings $\mathbf{z}_A$ and $\mathbf{z}_B$ fixed) to obtain a near-0 would reflect a clear separation between the cell morphology and GEX embeddings, which is not desirable for our integration goal (hence, adversarial).

The generator (i.e., encoders) has a penalty term to encourage better integration between GEX and morphology by updating the encoders (but holding the discriminator fixed),

$$\textbf{Penalty term in GeoAdvAE}: \quad \mathcal{L}_{\text{GAN}} = \mathbb{E}_{\mathbf{z}_A \sim q_A}\left[(D(\mathbf{z}_A) - (1-\epsilon))^2\right] + \mathbb{E}_{\mathbf{z}_B \sim q_B}\left[(D(\mathbf{z}_B) - \epsilon)^2\right]. \tag{3}$$

This particular formulation is inspired by Crossmodal-AE [64], where the updated encoder (and hence, updated embeddings $\mathbf{z}_A$ and $\mathbf{z}_B$) are incentivized to cause the discriminator to mispredict whether the cell's embedding originates from the GEX or the morphology modality.

Gromov–Wasserstein loss to enable uniform alignment. We align intra-modality geometry by minimizing the Gromov–Wasserstein (GW) discrepancy between pairwise distances in the two latent spaces. This penalty term encourages the integration to spread cells uniformly across the integration space so that as much of a one-to-one mapping between cells from each modality as possible can be achieved. To define this, we first compute the intra-modality distances,

$$D_{ij}^{(A)} = \left\|\mathbf{z}_i^{(A)} - \mathbf{z}_j^{(A)}\right\|_2^2, \qquad D_{kl}^{(B)} = \left\|\mathbf{z}_k^{(B)} - \mathbf{z}_l^{(B)}\right\|_2^2, \tag{4}$$

for all cell pairs $\{i, j\}$ and $\{k, l\}$ in modality A or B, respectively. Then, we match the geometry (among the cells in the minibatch) through an optimal transport plan to obtain a mapping T among all valid transportation plans Π, the set of all doubly stochastic matrices (i.e., rows and columns sum to 1).

$$\mathcal{L}_{\text{GW}} = \min_{T \in \Pi} \sum_{i,j,k,l} \left(D_{ij}^{(A)} - D_{kl}^{(B)}\right)^2 T_{ik}\, T_{jl}. \tag{5}$$

Prior-guided cluster alignment. To orient the shared latent space with coarse biology (e.g., excitatory vs. inhibitory in Patch-seq), we impose a prior on broad cluster–cluster correspondences. Let $C_A : \{1, \ldots, N_A\} \to \{1, \ldots, K_A\}$ and $C_B : \{1, \ldots, N_B\} \to \{1, \ldots, K_B\}$ be precomputed broad cluster labels for gene expression (modality A) and morphology (modality B). For instance, in our Patch-seq analysis, our clustering is simply separating excitatory from inhibitory neurons. In this example, there are well-established distinctions between these two categories of neurons from both gene expression and morphological data. Importantly, the clusterings we use here are *not* about more granular types of neurons, since they are more difficult to define from a single modality alone. It ensures that broad, literature-supported cell categories remain coherently aligned across modalities, providing semantic structure that complements the unsupervised adversarial and geometric objectives.

The user also provides a correspondence matrix $\mathbf{P} \in [0,1]^{K_A \times K_B}$ whose entry P_{jk} encodes the expected association strength between GEX cluster j and morphology cluster k. This induces a target similarity

$$S_{ij}^{(\text{target})} = P_{C_A(i),\, C_B(j)}, \quad \text{for } i \in \{1, \ldots, N_A\}, j \in \{1, \ldots, N_B\}.$$

For a minibatch with latent codes $\mathbf{z}_A^{(i)}$ and $\mathbf{z}_B^{(j)}$, we define the predicted cross-modal similarity by cosine similarity of ℓ_2-normalized embeddings,

$$S_{ij}^{(\text{pred})} = \left(\frac{\mathbf{z}_A^{(i)}}{\|\mathbf{z}_A^{(i)}\|_2}\right)^\top \left(\frac{\mathbf{z}_B^{(j)}}{\|\mathbf{z}_B^{(j)}\|_2}\right).$$

The prior loss matches these similarities via a temperature-scaled mean-squared error,

$$\mathcal{L}_{\text{prior}} = \frac{1}{N_A N_B} \cdot \sum_{i=1}^{N_A} \sum_{j=1}^{N_B} \left(\frac{S_{ij}^{(\text{pred})}}{\tau} - \frac{S_{ij}^{(\text{target})}}{\tau} \right)^2,$$

where $\tau > 0$ sharpens contrasts, emphasizing confident matches and down-weighting weak ones.

3 Data processing and evaluation

The model operates on two unpaired modalities: morphology and gene expression. The morphologies are "skeletons" of the cell in `.swc` format, recording the 3-dimensional location and connection between each "point." These morphologies are processed using CAJAL [17], which encodes structural information in a continuous, numerical space suitable for GeoAdvAE's deep-learning framework. The transcriptomics are the 2,000 highly variable genes, which we log-normalize to remove the confounding effect of sequencing depth.

Diagonal integration methods to compare against. We benchmark GeoAdvAE against representative cross-modal single-cell integrators spanning distinct paradigms. For instance, coupled/graph autoencoders or latent embedding methods learn shared spaces for unmatched modalities (ScDART [69]; STACI [68]; DCCA [1], scJoint [38]). Some methods also incorporate adversarial terms to encourage a better integration (CycleGAN [70], sciCAN [63], SCIM [51], Crossmodal-AE [64]). Implicit in such methods is the assumption that the axes of variation that best represent the modality are also the axes shared between the two modalities. As we have seen in Patch-seq data (Fig. 1A), this assumption is not necessarily true for GEX-morphology alignment. In contrast, optimal-transport (i.e., geometry matching) methods align the global structures between the two modalities (SCOT [12], MMD-MA [39], UnionCom [3]). However, these methods implicitly assume that the global structure for each modality has similar structural properties, and they also lack the ability to distinguish the "proper orientation" of the integration. The Patch-seq data (Fig. 1A) demonstrates that both drawbacks are detrimental when aligning GEX-morphology of neural cells.

Evaluation criteria. We evaluate diagonal integration via cross-modal cell-type transfer accuracy. Importantly, these "true" cell-type labels are *not* used during the training of any of the proposed methods. For each query cell in one modality, we take its 1-nearest neighbor (in Manhattan distance) in the other modality within the latent integrated space and count a match if cell-type labels agree. Accuracy is defined as

$$\text{Accuracy} \;=\; \frac{1}{2} \cdot \left(\frac{\#\text{matches}_{A \to B}}{\#\text{queries}_{A \to B}} \;+\; \frac{\#\text{matches}_{B \to A}}{\#\text{queries}_{B \to A}} \right),$$

where higher values indicate a better cross-modal alignment.

Downstream gene investigation. The importance of GeoAdvAE's diagonal integration lies in its ability to refine our biological understanding of which genes are associated with morphological changes. We use integrated gradients [55], a computational framework for assessing how perturbing a gene's expression alters the embedding produced by the GEX encoder $E_A(\cdot)$. By interpreting each gene's impact as its "importance scores," we can identify pathways that concurrently have high importance scores via a gene set enrichment analysis (GSEA, [53]). Importantly, this downstream analysis can also identify pathways known to be differentially expressed across cell populations that do *not* have a morphological component.

4 Results

4.1 Simulated data demonstrates GeoAdvAE's advantage over other method and enables ablation studies

We build a simulator that generates synthetic gene-expression profiles and corresponding neuronal morphologies to evaluate GeoAdvAE under controlled conditions. This enables us to precisely diagnose how each loss component of GeoAdvAE contributes to its superior performance over other integration methods. Three canonical neuron types (pyramidal, multipolar, bipolar) are modeled by sampling GEX from three low-dimensional clusters; morphology is then generated by a process in which GEX controls polarity, branch density, and anisotropy, yielding three separable CAJAL morphology clusters (Supplemental Fig. S1). For diagonal integration, we use a 3×3 correspondence matrix $\mathbf{P}$ (the identity), which maps each GEX cluster to its morphological counterpart (Fig. 3B).

To illustrate the importance of each penalty term in GeoAdvAE, we perform an ablation study, where we remove one of the penalty terms and assess how the resulting integration suffers (Fig. 2C). Removing the adversarial classifier separates modalities in the latent space. Removing the Gromov–Wasserstein penalty breaks local geometric coherence. Removing prior-guided cluster alignment misorients matched clusters. Removing reconstruction loss destabilizes embeddings and blurs cluster boundaries. Notably, although our correspondence matrix $\mathbf{P}$ "reveals" how the two modalities should be integrated, we find that this biological information is *necessary* for GEX-morphology integration. In the presence of no paired samples and no matched features between the two modalities, the information represented by $\mathbf{P}$ provides the necessary orientation signal that the data alone cannot recover.

We next compare GeoAdvAE with other competing methods using both quantitative and qualitative evaluations. GeoAdvAE achieves the highest alignment accuracy (Fig. 3A). Latent embedding methods (scJOINT and STACI) perform reasonably well but tend to under-align fine-grained structures, while optimal-transport methods (UnionComm, MMD-MA) struggle to uncover the correct integration orientation (Fig. 3B, Supplemental Fig. S3). Collectively, these results support combining adversarial alignment with cluster-prior regularization to achieve proper alignment between GEX and morphology. See Appendix S4 for more results, such as a more in-depth ablation analysis via Shapley values and a comparison between using GW and Maximum Mean Discrepancy (MMD) losses to measure alignment.

4.2 Patch-seq neurons provide validation for GeoAdvAE's integration

As mentioned in Section 1, Patch-sequencing is one of the few wet-bench protocols where the GEX and morphology of neurons can be simultaneously measured. This data provides an invaluable resource for validating the performance of any GEX-morphology diagonal integration method. We evaluate our method on a Patch-seq dataset [47], focusing on the 645 neurons spanning a diverse set of excitatory and inhibitory neurons from the mouse motor cortex that both gene expression and morphology profiled. To define the cluster prior matrix in the Patch-seq dataset, we *do not* use the provided cell annotations, as they were defined by the original authors across multiple modalities. Instead, for GEX, we clustered the neurons into 4 clusters and annotated each cluster as either excitatory or inhibitory using canonical marker genes [36]. For morphology, we embedded the neurons into a lower-dimensional space using CAJAL [17] and grouped them into 4 clusters. Based on our visual assessment of 10 randomly selected neurons in each cluster, we annotated whether each neuron was excitatory or inhibitory, and whether it had a pyramidal shape. Based on these manual annotations, we construct the correspondence matrix $\mathbf{P}$ mapping between the GEX and

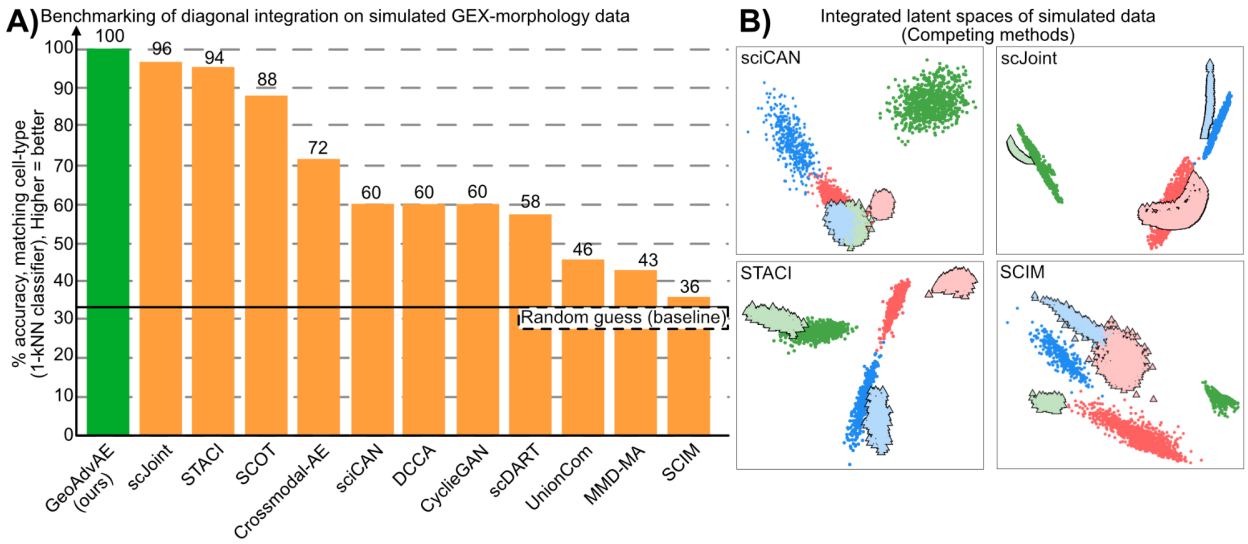

Fig. 3. Simulation study via synthetic neurons. A) Comparison of cell-type matching accuracy on simulation data across GeoAdvAE and competing cross-modal integration methods. The solid horizontal line denotes the accuracy achieved by random guessing. B) Integrated latent spaces of simulated data produced by selected baseline methods (sciCAN [63], STACI [68], scJoint [38], SCIM [51]). The results by these competing methods can be compared against GeoAdvAE's integration shown in Figure 2B.

morphology clusters based on how highly our coarse excitatory-vs.-inhibitory labels align between the two modalities. The details are described in Appendix S3.

We validated GeoAdvAE using Patch-seq neurons with joint GEX–morphology measurements. GeoAdvAE achieved 34% cell-type alignment accuracy, outperforming graph-based baselines such as ScDART (28%) and STACI (27%), as shown in Fig. 4A. Graph-based approaches (GeoAdvAE, ScDART, STACI) perform best, whereas latent-space alignment methods (e.g., scJoint, Crossmodal-AE) and adversarial models (e.g., CycleGAN, sciCAN) degrade on Patch-seq, likely because they struggle to preserve modality consistency in the presence of biological noise and nonlinear morphological variability. Consistent with our simulations (Fig. 3A), optimal transport methods (SCOT, UnionCom) perform poorly, as the high diversity of GEX axes often lacks corresponding morphological features. UMAP visualizations confirm that GeoAdvAE produces a smoother, more biologically consistent latent mixture than Crossmodal-AE (Fig. 4B,D). Additionally, the continuous latent space (Fig. 4B) suggests an alignment of cellular states rather than individual cells (analysis and results shown in Fig. S8). We further benchmarked GeoAdvAE against Supervised PCA [2] to quantify the specific impact of prior information (Fig. S10) and against Tilted-CCA [37] to establish the theoretical "ceiling" for shared signal recovery (Fig. S12).

Model interpretation via integrated gradients on the GEX encoder $E_A(\cdot)$ confirmed that GeoAdvAE captures essential neuronal biology. We highlight two important pathways highlighted by GSEA in Fig. 4C. First, pathways for neuron projection guidance (GO:0097485) and axon guidance play a critical role in axon/neurite guidance and cytoskeletal regulation (example genes: *Slit3*; p-value: 0.016). Additionally, Rho-family small GTPases are master regulators of neuronal morphogenesis, coupling extracellular cues to actin remodeling that governs spine structure, filopodial dynamics, and growth-cone advance or collapse [40]. In diverse contexts, guidance and trophic signaling converge on these GTPases to control growth-cone behavior and neurite extension [60,10]. Consistent with this framework, we find that the relevant GO term is strongly associated with GEX-morphology, via genes such as *Rasgrp1* and *Ngf*. We ensure these findings are supported by analyses that directly leverage the paired nature of Patch-seq data (Fig. S6).

Additional results are in Appendix S4, where we also map the full 1329 GEX cells to the 645 cells with morphology, and also assess how GeoAdvAE performs with classically used morphological features instead of with CAJAL-derived features.

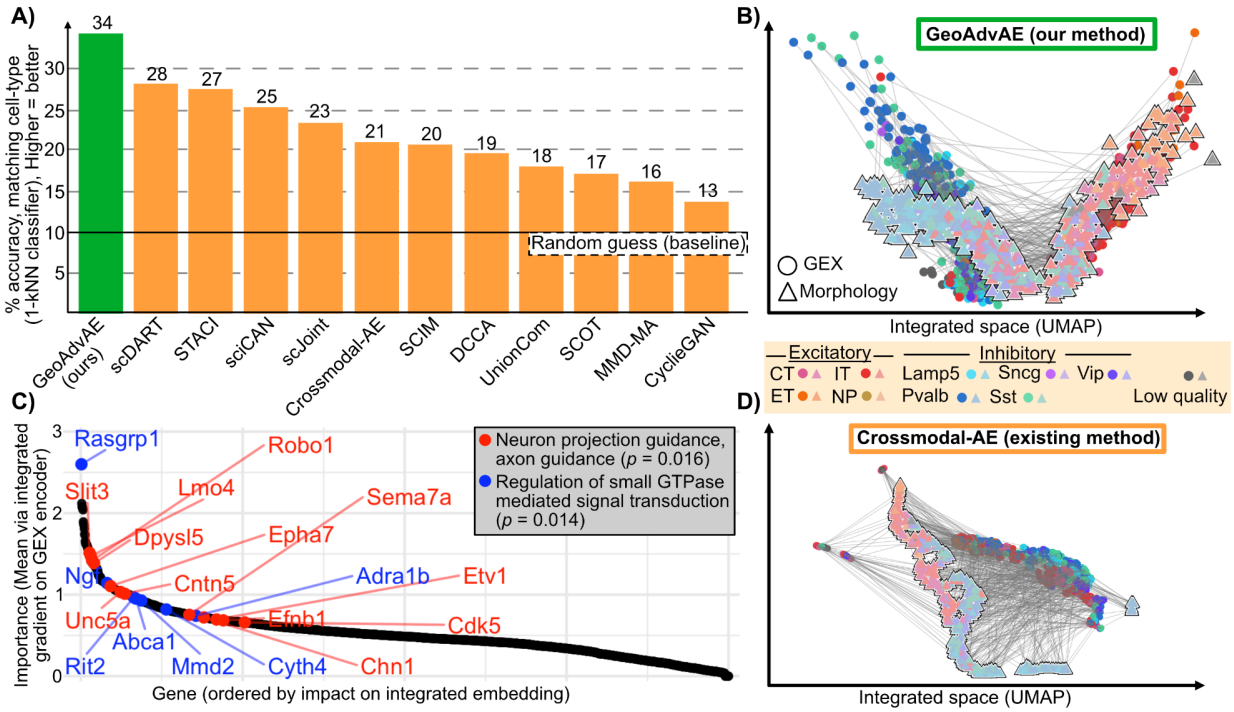

Fig. 4. Comparison and validation of integration methods using Patch-seq neurons. A) Comparison of cell-type matching accuracy on Patch-seq data across GeoAdvAE and competing cross-modal integration methods. GeoAdvAE achieves the highest alignment accuracy. B) Integrated latent space learned by GeoAdvAE on Patch-seq neurons. C) Integrated gradient importance scores of all the 2000 genes based on the GEX encoder. The genes in two critical pathways known to regulate neuron morphologies are highlighted in red and blue, both nominally significant. D) Integrated latent space generated by the existing Crossmodal-AE method [64], showing less continuous cross-modal alignment compared to GeoAdvAE.

4.3 GeoAdvAE uncovers novel biology of 5xFAD microglia

Our primary dataset of interest regards microglia from the 5xFAD mouse model, which recapitulates certain hallmarks of Alzheimer's disease (AD), see Section 1.1. For morphological data, we leverage 98 microglial skeletons from mice at 3 and 6 months of both genders [9]. After applying CAJAL and clustering, we observe 3 primary microglial clusters (Fig. 5A), where one microglial cluster (cyan) is enriched for larger microglia with a more ramified morphology primarily from male, younger mice. In contrast, another microglial cluster (pink) is enriched for smaller microglia with a more amoeboid morphology, primarily from female, older mice. For transcriptomic data, we leverage 31,948 microglial single-nuclei RNA-sequencing [57]. These microglia span many different states that the authors classified using marker genes, including homeostatic, proliferating, interferon-response (IRM), and disease-associated microglia (DAM) (Fig. 5B).

To form the correspondence matrix $\mathbf{P}$ necessary for the prior-guided cluster alignment, we annotate broad correspondences between the two modalities. For GEX, we use the provided 12 GEX microglia clusters and further quantify the proportions of homeostatic, proliferating, IRM, and DAM microglia within each cluster using marker-gene enrichment analysis. This enables us to capture continuous state transitions among microglial states, rather than discretely labeling each cluster as a "pure" state. For morphology, we clustered the CAJAL profiles into 3 groups and labeled them as ramified, intermediate, and amoeboid based on visual inspection of microglia in each group. Based on these annotations, we design $\mathbf{P}$ to account for within-cluster proportions and the biological correspondence between microglial transcriptomic and morphological states. See Appendix S3 for more details on this procedure.

When we applied GeoAdvAE to the microglia datasets, we observed an apparent 1-dimensional manifold, suggesting novel biological insights into microglia. Fig. 6A depicts the integrated embedding where the

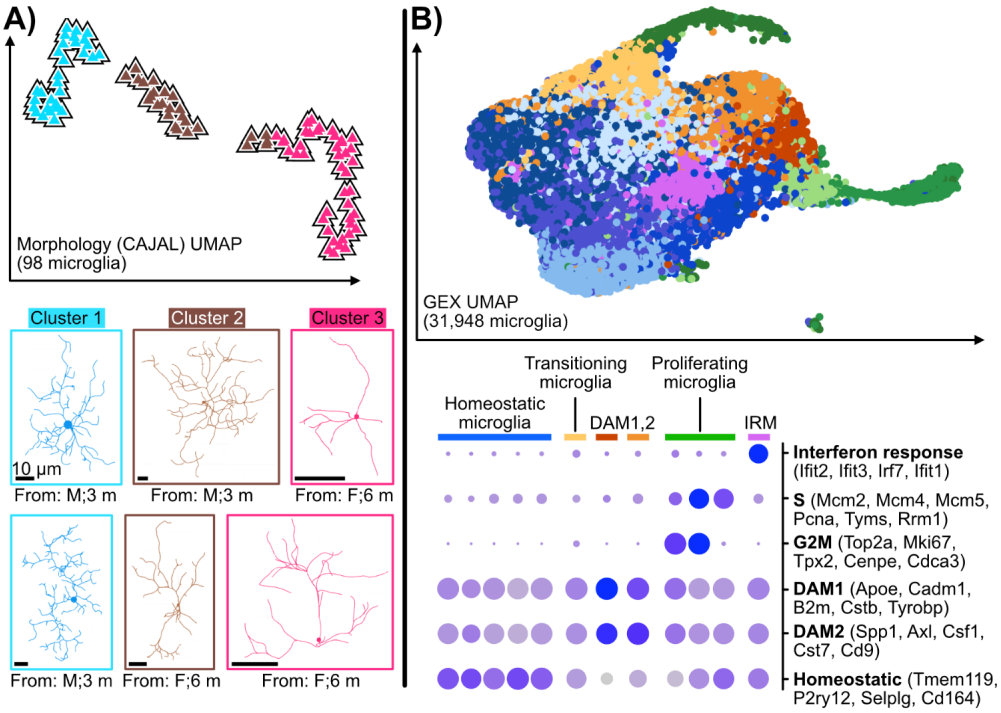

Fig. 5. 5xFAD microglia data. A) Microglia morphologies [9], quantified via CAJAL and clustered into 3 major morphological types. Exemplary microglia of each cluster are shown below. B) Microglia gene expressions [57], using the author's annotated microglial states. Marker genes for each microglia states are shown as a dot plot below.

ramified morphologies (cyan triangles) integrate with the homeostatic microglia (blue circles), while the amoeboid morphologies (pink triangles) integrate with the DAM microglia (orange circles). While this overall trend is not surprising due to GeoAdvAE's prior-guided cluster alignment, we perform downstream analyses to investigate two notable biological questions: 1) which particular genes contributed to this alignment, and 2) what is the continuum of microglial states that spanned this 1-dimensional manifold?

We used our integrated gradient framework to identify genes associated with microglial morphology. Ordering genes by importance reveals two significant, opposing pathways (Fig. 6B): DNA repair genes were enriched among ramified microglia, whereas cell killing genes were enriched among amoeboid microglia. This is consistent with prior work showing that homeostatic, ramified microglia engage in tissue-maintenance programs [45,67,49] while amoeboid microglia adopt highly phagocytic, neurotoxic phenotypes that surround and eliminate stressed or dying neurons [45,15].

To examine genes varying along this continuum, we ordered microglia (both morphology and transcriptomic) along the manifold using Slingshot [52] and correlated each gene with this ordering. We find *Ms4a6b* to be strongly associated with ramified morphologies. This is noteworthy since the closest human homolog, MS4A6A, has conflicting evidence regarding AD. MS4A6A is implicated in disease progression through TREM2 in some studies [46,29] but interpreted as a marker of accelerated aging in others [50,32]. Our results suggested *Ms4a6b* is not primarily linked to inflammatory responses in AD. Conversely, we found *Ftl1* and *Fth1*, two iron-loading genes whose upregulation is a hallmark of dystrophic microglia, and whose elevated expression localizes around AD plaques and correlates with pathological burden [22,30]. Surprisingly, several disease-associated microglia (DAM) complement markers were not correlated with this GEX-morphology axis, including *C1qa*, *C1qb*, *C1qc*, *C3*, and *C4b* [26,16]. This suggested that complement activation represents an upstream or partially orthogonal transcriptomic program that can be engaged without large shifts in soma size or process complexity, consistent with recent work showing that microglial morphology is only one particular readout of a microglia's function [20,62]. Future gene knockout experiments could be performed to validate these findings. All in all, these findings demonstrate GeoAdvAE's ability to uncover

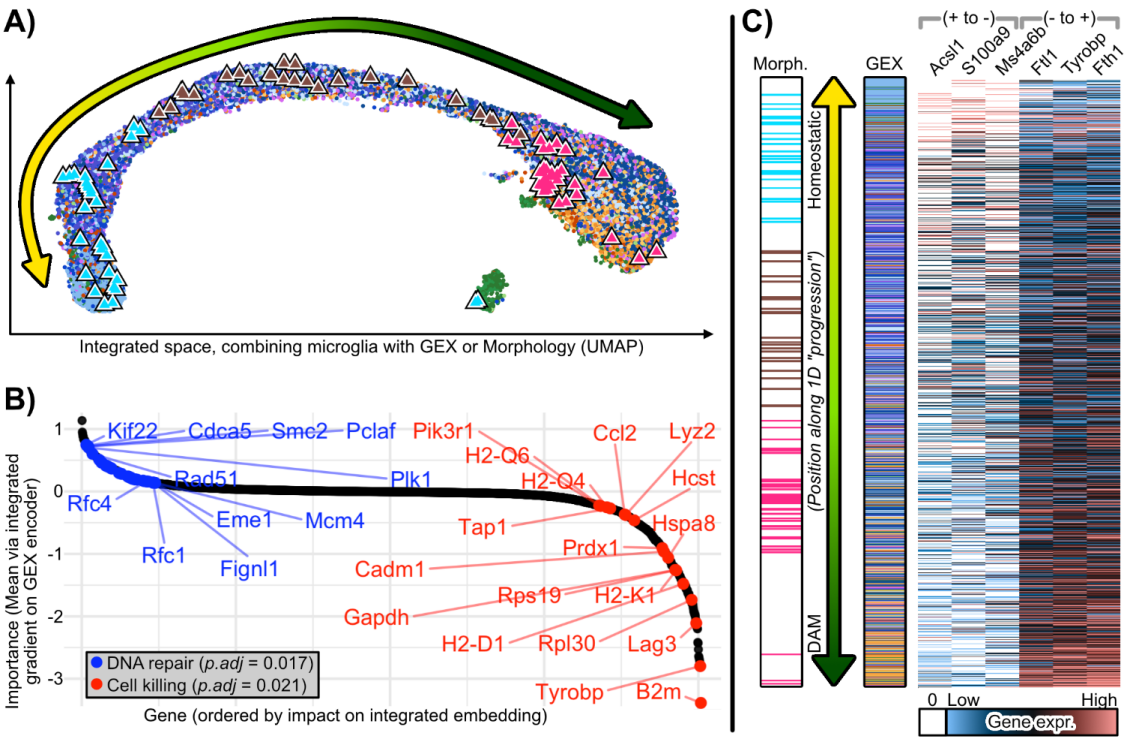

Fig. 6. Integration of GEX and morphology uncovers functional roles of microglia in 5xFAD. A) Integration of microglia GEX and morphology, uncovering a 1-dimensional manifold that correspond to a spectrum of microglial states. B) Genes ordered by their influence on a microglia's embedding in the integrated spaced as measured by integrated gradients [55], highlight two major pathways corresponding to the two different ends of the GEX-morphology integrated space. C) Ordering microglia along the 1-dimensional manifold, highlighting how microglia measured by morphology (left) or GEX (right) align along this manifold. Exemplary genes that also generally increase or decrease going from one end of the spectrum to the other are displayed as a heatmap.

promising GEX-morphology relations, enabling us to better characterize the morphological consequences of up/down-regulation of particular microglial transcriptomic programs.

5 Discussion

GeoAdvAE provides a general framework for integrating cellular morphology and gene expression from unpaired data, enabling us to link "form" and "function" even when modalities are measured in different cells. Using simulations and Patch-seq neurons, we showed that combining adversarial alignment, Gromov–Wasserstein regularization, and prior-guided cluster correspondences yields more accurate and biologically coherent cross-modal embeddings than existing integration methods. Applied to 5xFAD microglia, GeoAdvAE uncovers a one-dimensional ramified-to-amoeboid continuum, highlights DNA repair and cell-killing programs associated with morphological transitions, and reveals complement/DAM signatures that appear decoupled from visible shape changes. These results suggest that some microglial transcriptional programs lie upstream of, or orthogonal to, large-scale morphological remodeling, underscoring the limits of morphology alone as a proxy for microglial state. More broadly, GeoAdvAE offers a scalable template for integrating unpaired morphology with high-dimensional omics data in other brain and non-brain systems where joint profiling remains impractical.

References

1. Andrew, G., Arora, R., Bilmes, J., and Livescu, K. Deep canonical correlation analysis. In *International Conference on Machine Learning* (2013), PMLR, pp. 1247–1255.
2. Barshan, E., Ghodsi, A., Azimifar, Z., and Jahromi, M. Z. Supervised principal component analysis: Visualization, classification and regression on subspaces and submanifolds. *Pattern Recognition 44*, 7 (2011), 1357–1371.
3. Cao, K., Bai, X., Hong, Y., and Wan, L. Unsupervised topological alignment for single-cell multi-omics integration. *Bioinformatics 36* (2020), i48–i56.
4. Cao, K., Hong, Y., and Wan, L. Manifold alignment for heterogeneous single-cell multi-omics data integration using Pamona. *Bioinformatics 38*, 1 (2022), 211–219.
5. Chartrand, T., Dalley, R., Close, J., Goriounova, N. A., Lee, B. R., Mann, R., Miller, J. A., Molnar, G., Mukora, A., Alfiler, L., et al. Morphoelectric and transcriptomic divergence of the layer 1 interneuron repertoire in human versus mouse neocortex. *Science 382*, 6667 (2023), eadf0805.
6. Chelebian, E., Avenel, C., and Wählby, C. What makes for good morphology representations for spatial omics? *arXiv preprint arXiv:2407.20660* (2024).
7. Chen, S., Zhu, B., Huang, S., Hickey, J. W., Lin, K. Z., Snyder, M., Greenleaf, W. J., Nolan, G. P., Zhang, N. R., and Ma, Z. Integration of spatial and single-cell data across modalities with weakly linked features. *Nature Biotechnology 42*, 7 (2024), 1096–1106.
8. Claes, C., Danhash, E. P., Hasselmann, J., Chadarevian, J. P., Shabestari, S. K., England, W. E., Lim, T. E., Hidalgo, J. L. S., Spitale, R. C., Davtyan, H., et al. Plaque-associated human microglia accumulate lipid droplets in a chimeric model of Alzheimer's disease. *Molecular neurodegeneration 16*, 1 (2021), 50.
9. Colombo, G., Cubero, R. J. A., Kanari, L., Venturino, A., Schulz, R., Scolamiero, M., Agerberg, J., Mathys, H., Tsai, L.-H., Chachólski, W., et al. A tool for mapping microglial morphology, morphOMICs, reveals brain-region and sex-dependent phenotypes. *Nature Neuroscience 25*, 10 (2022), 1379–1393.
10. Daniels, R. H., Hall, P. S., and Bokoch, G. M. Membrane targeting of p21-activated kinase 1 (PAK1) induces neurite outgrowth from PC12 cells. *The EMBO journal* (1998).
11. de Vries, L. E., Huitinga, I., Kessels, H. W., Swaab, D. F., and Verhaagen, J. The concept of resilience to Alzheimer's disease: Current definitions and cellular and molecular mechanisms. *Molecular neurodegeneration 19*, 1 (2024), 33.
12. Demetci, P., Santorella, R., Sandstede, B., Noble, W. S., and Singh, R. SCOT: Single-cell multi-omics alignment with optimal transport. *Journal of Computational Biology 29*, 1 (2022), 3–18.
13. Dorman, L. C., Nguyen, P. T., Escoubas, C. C., Vainchtein, I. D., Xiao, Y., Lidsky, P. V., Nakajo, H., Silva, N. J., Lagares-Linares, C., Wang, E. Y., et al. A type I interferon response defines a conserved microglial state required for effective neuronal phagocytosis. *BioRxiv* (2022), 2021–04.
14. Gabitto, M. I., Travaglini, K. J., Rachleff, V. M., Kaplan, E. S., Long, B., Ariza, J., Ding, Y., Mahoney, J. T., Dee, N., Goldy, J., et al. Integrated multimodal cell atlas of Alzheimer's disease. *Nature Neuroscience 27*, 12 (2024), 2366–2383.
15. Galloway, D. A., Phillips, A. E., Owen, D. R., and Moore, C. S. Phagocytosis in the brain: Homeostasis and disease. *Frontiers in Immunology 10* (2019), 790.
16. Gomez-Arboledas, A., Acharya, M. M., and Tenner, A. J. The role of complement in synaptic pruning and neurodegeneration. *ImmunoTargets and Therapy* (2021), 373–386.
17. Gouwens, N. W., Sorensen, S. A., Baftizadeh, F., Budzillo, A., Lee, B. R., Jarsky, T., Alfiler, L., Baker, K., Barkan, E., Berry, K., et al. Integrated morphoelectric and transcriptomic classification of cortical GABAergic cells. *Cell 183*, 4 (2020), 935–953.
18. Govek, K. W., Nicodemus, P., Lin, Y., Crawford, J., Saturnino, A. B., Cui, H., Zoga, K., Hart, M. P., and Camara, P. G. CAJAL enables analysis and integration of single-cell morphological data using metric geometry. *Nature Communications 14*, 1 (2023), 3672.
19. Green, G. S., Fujita, M., Yang, H.-S., Taga, M., Cain, A., McCabe, C., Comandante-Lou, N., White, C. C., Schmidtner, A. K., Zeng, L., et al. Cellular communities reveal trajectories of brain ageing and Alzheimer's disease. *Nature 633*, 8030 (2024), 634–645.
20. Green, T. R., and Rowe, R. K. Quantifying microglial morphology: An insight into function. *Clinical and Experimental Immunology 216*, 3 (2024), 221–229.
21. Grubman, A., Choo, X. Y., Chew, G., Ouyang, J. F., Sun, G., Croft, N. P., Rossello, F. J., Simmons, R., Buckberry, S., Landin, D. V., et al. Transcriptional signature in microglia associated with Aβ plaque phagocytosis. *Nature Communications 12*, 1 (2021), 3015.
22. Grundke-Iqbal, I., Fleming, J., Tung, Y.-C., Lassmann, H., Iqbal, K., and Joshi, J. Ferritin is a component of the neuritic (senile) plaque in Alzheimer dementia. *Acta Neuropathologica 81*, 2 (1990), 105–110.

23. Guo, S., Wang, H., and Yin, Y. Microglia polarization from M1 to M2 in neurodegenerative diseases. *Frontiers in Aging Neuroscience 14* (2022), 815347.
24. Haghighi, M., Caicedo, J. C., Cimini, B. A., Carpenter, A. E., and Singh, S. High-dimensional gene expression and morphology profiles of cells across 28,000 genetic and chemical perturbations. *Nature Methods 19*, 12 (2022), 1550–1557.
25. Hansen, D. V., Hanson, J. E., and Sheng, M. Microglia in Alzheimer's disease. *Journal of Cell Biology 217*, 2 (2018), 459–472.
26. Hong, S., Beja-Glasser, V. F., Nfonoyim, B. M., Frouin, A., Li, S., Ramakrishnan, S., Merry, K. M., Shi, Q., Rosenthal, A., Barres, B. A., et al. Complement and microglia mediate early synapse loss in Alzheimer mouse models. *Science 352*, 6286 (2016), 712–716.
27. Huang, J., Yuan, C., Jiang, J., Chen, J., Badve, S. S., Gokmen-Polar, Y., Segura, R. L., Yan, X., Lazar, A., Gao, J., et al. Bridging cell morphological behaviors and molecular dynamics in multi-modal spatial omics with MorphLink. *Nature Communications 16*, 1 (2025), 5878.
28. Jiang, J., Liu, Y., Qin, J., Chen, J., Wu, J., Pizzi, M. P., Lazcano, R., Yamashita, K., Xu, Z., Pei, G., et al. METI: Deep profiling of tumor ecosystems by integrating cell morphology and spatial transcriptomics. *Nature Communications 15*, 1 (2024), 7312.
29. Jiao, H.-S., Ge, Y.-J., Huang, L.-Y., Liu, Y., Wu, B.-S., Lian, P.-P., Hao, Y.-N., Han, S.-S., Li, Y.-T., Wu, K.-M., et al. MS4A6A/Ms4a6d deficiency disrupts neuroprotective microglia functions and promotes inflammation in Alzheimer's disease model. *Molecular Neurodegeneration 20*, 1 (2025), 94.
30. Kenkhuis, B., Somarakis, A., de Haan, L., Dzyubachyk, O., IJsselsteijn, M. E., de Miranda, N. F., Lelieveldt, B. P., Dijkstra, J., van Roon-Mom, W. M., Höllt, T., et al. Iron loading is a prominent feature of activated microglia in Alzheimer's disease patients. *Acta neuropathologica communications 9*, 1 (2021), 27.
31. Keren-Shaul, H., Spinrad, A., Weiner, A., Matcovitch-Natan, O., Dvir-Szternfeld, R., Ulland, T. K., David, E., Baruch, K., Lara-Astaiso, D., Toth, B., et al. A unique microglia type associated with restricting development of Alzheimer's disease. *Cell 169*, 7 (2017), 1276–1290.
32. Lee, D., Vicari, J. M., Porras, C., Spencer, C., Pjanic, M., Wang, X., Kinrot, S., Weiler, P., Kosoy, R., Bendl, J., et al. Plasticity of human microglia and brain perivascular macrophages in aging and Alzheimer's disease. *medRxiv* (2024), 2023–10.
33. Lee, H., and Welch, J. D. MorphNet predicts cell morphology from single-cell gene expression. *bioRxiv* (2022), 2022–10.
34. Leng, F., and Edison, P. Neuroinflammation and microglial activation in Alzheimer disease: Where do we go from here? *Nature Reviews Neurology 17*, 3 (2021), 157–172.
35. Li, X., Li, Y., Jin, Y., Zhang, Y., Wu, J., Xu, Z., Huang, Y., Cai, L., Gao, S., Liu, T., et al. Transcriptional and epigenetic decoding of the microglial aging process. *Nature aging 3*, 10 (2023), 1288–1311.
36. Lin, G., Chancellor, S. E., Kwon, T., Woodbury, M. E., Doering, A., Abdourahman, A., Bennett, R. E., Liao, F., Pastika, T., Tamm, J., et al. Cell-death pathways and tau-associated neuronal vulnerability in Alzheimer's disease. *Cell Reports 44*, 6 (2025).
37. Lin, K. Z., and Zhang, N. R. Quantifying common and distinct information in single-cell multimodal data with Tilted Canonical Correlation Analysis. *Proceedings of the National Academy of Sciences 120*, 32 (2023), e2303647120.
38. Lin, Y., Wu, T.-Y., Wan, S., Yang, J. Y., Wong, W. H., and Wang, Y. R. scJoint integrates atlas-scale single-cell RNA-seq and ATAC-seq data with transfer learning. *Nature Biotechnology 40*, 5 (2022), 703–710.
39. Liu, J., Huang, Y., Singh, R., Vert, J.-P., and Noble, W. S. Jointly embedding multiple single-cell omics measurements. In *Algorithms in Bioinformatics* (2019), vol. 143, p. 10.
40. Luo, L. Rho gtpases in neuronal morphogenesis. *Nature Reviews Neuroscience 1*, 3 (2000), 173–180.
41. Mancuso, R., Fattorelli, N., Martinez-Muriana, A., Davis, E., Wolfs, L., Van Den Daele, J., Geric, I., Premereur, J., Polanco, P., Bijnens, B., et al. Xenografted human microglia display diverse transcriptomic states in response to Alzheimer's disease-related amyloid-β pathology. *Nature Neuroscience 27*, 5 (2024), 886–900.
42. Oakley, H., Cole, S. L., Logan, S., Maus, E., Shao, P., Craft, J., Guillozet-Bongaarts, A., Ohno, M., Disterhoft, J., Van Eldik, L., et al. Intraneuronal β-amyloid aggregates, neurodegeneration, and neuron loss in transgenic mice with five familial Alzheimer's disease mutations: Potential factors in amyloid plaque formation. *Journal of Neuroscience 26*, 40 (2006), 10129–10140.
43. Olmos-Alonso, A., Schetters, S. T., Sri, S., Askew, K., Mancuso, R., Vargas-Caballero, M., Holscher, C., Perry, V. H., and Gomez-Nicola, D. Pharmacological targeting of CSF1R inhibits microglial proliferation and prevents the progression of Alzheimer's-like pathology. *Brain 139*, 3 (2016), 891–907.
44. Qian, X., Coleman, K., Jiang, S., Kriz, A. J., Marciano, J. H., Luo, C., Cai, C., Manam, M. D., Caglayan, E., Lai, A., et al. Spatial transcriptomics reveals human cortical layer and area specification. *Nature* (2025), 1–11.

45. Reddaway, J., Richardson, P. E., Bevan, R. J., Stoneman, J., and Palombo, M. Microglial morphometric analysis: So many options, so little consistency. *Frontiers in Neuroinformatics 17* (2023), 1211188.
46. Rosner, D., Sun, J., Yee, A., Wagh, C., Rychkova, A., Cacace, R., Buonfiglioli, A., Alwahagri, M., Kong, P., Roell, M., et al. The Alzheimer's disease risk genes MS4A4A and MS4A6A cooperate to negatively regulate Trem2 and microglia states. *bioRxiv* (2024), 2024–11.
47. Scala, F., Kobak, D., Bernabucci, M., Bernaerts, Y., Cadwell, C. R., Castro, J. R., Hartmanis, L., Jiang, X., Laturnus, S., Miranda, E., et al. Phenotypic variation of transcriptomic cell types in mouse motor cortex. *Nature 598*, 7879 (2021), 144–150.
48. Singh, R., Demetci, P., Bonora, G., Ramani, V., Lee, C., Fang, H., Duan, Z., Deng, X., Shendure, J., Disteche, C., et al. Unsupervised manifold alignment for single-cell multi-omics data. In *Proceedings of the 11th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics* (2020), pp. 1–10.
49. Spittau, B. Aging microglia—phenotypes, functions and implications for age-related neurodegenerative diseases. *Frontiers in Aging Neuroscience 9* (2017), 194.
50. Srinivasan, K., Friedman, B. A., Etxeberria, A., Huntley, M. A., van Der Brug, M. P., Foreman, O., Paw, J. S., Modrusan, Z., Beach, T. G., Serrano, G. E., et al. Alzheimer's patient microglia exhibit enhanced aging and unique transcriptional activation. *Cell Reports 31*, 13 (2020).
51. Stark, S. G., Ficek, J., Locatello, F., Bonilla, X., Chevrier, S., Singer, F., Rätsch, G., and Lehmann, K.-V. SCIM: Universal single-cell matching with unpaired feature sets. *Bioinformatics 36* (2020), i919–i927.
52. Street, K., Risso, D., Fletcher, R. B., Das, D., Ngai, J., Yosef, N., Purdom, E., and Dudoit, S. Slingshot: Cell lineage and pseudotime inference for single-cell transcriptomics. *BMC Genomics 19*, 1 (2018), 477.
53. Subramanian, A., Tamayo, P., Mootha, V. K., Mukherjee, S., Ebert, B. L., Gillette, M. A., Paulovich, A., Pomeroy, S. L., Golub, T. R., Lander, E. S., et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences 102*, 43 (2005), 15545–15550.
54. Sun, N., Victor, M. B., Park, Y. P., Xiong, X., Scannail, A. N., Leary, N., Prosper, S., Viswanathan, S., Luna, X., Boix, C. A., et al. Human microglial state dynamics in Alzheimer's disease progression. *Cell 186*, 20 (2023), 4386–4403.
55. Sundararajan, M., Taly, A., and Yan, Q. Axiomatic attribution for deep networks. In *International conference on machine learning* (2017), PMLR, pp. 3319–3328.
56. Vidal-Itriago, A., Radford, R. A., Aramideh, J. A., Maurel, C., Scherer, N. M., Don, E. K., Lee, A., Chung, R. S., Graeber, M. B., and Morsch, M. Microglia morphophysiological diversity and its implications for the CNS. *Frontiers in Immunology 13* (2022), 997786.
57. Wang, S., Mustafa, M., Yuede, C. M., Salazar, S. V., Kong, P., Long, H., Ward, M., Siddiqui, O., Paul, R., Gilfillan, S., et al. Anti-human TREM2 induces microglia proliferation and reduces pathology in an Alzheimer's disease model. *Journal of Experimental Medicine 217*, 9 (2020), e20200785.
58. Wang, X., Fan, Y., Guo, Y., Fu, C., Lee, K., Dallakyan, K., Li, Y., Yin, Q., Li, Y., and Song, L. Prediction of cellular morphology changes under perturbations with a transcriptome-guided diffusion model. *Nature Communications 16*, 1 (2025), 8210.
59. Watson, C., Paxinos, G., and Puelles, L. *The mouse nervous system.* Academic Press, 2011.
60. Wegmeyer, H., Egea, J., Rabe, N., Gezelius, H., Filosa, A., Enjin, A., Varoqueaux, F., Deininger, K., Schnütgen, F., Brose, N., et al. EphA4-dependent axon guidance is mediated by the RacGAP α2-chimaerin. *Neuron 55*, 5 (2007), 756–767.
61. Whitelaw, B. S., Stoessel, M. B., and Majewska, A. K. Movers and shakers: Microglial dynamics and modulation of neural networks. *Glia 71*, 7 (2023), 1575–1591.
62. Woodburn, S. C., Bollinger, J. L., and Wohleb, E. S. The semantics of microglia activation: Neuroinflammation, homeostasis, and stress. *Journal of Neuroinflammation 18*, 1 (2021), 258.
63. Xu, Y., Begoli, E., and McCord, R. P. scican: Single-cell chromatin accessibility and gene expression data integration via cycle-consistent adversarial network. *NPJ Systems Biology and Applications 8*, 1 (2022), 33.
64. Yang, K. D., Damodaran, K., Bensaddek, D., et al. Multi-domain translation between single-cell imaging and sequencing data using autoencoders. *Nature Communications 12*, 1 (2021), 31.
65. Yu, G., Wang, L.-G., Han, Y., and He, Q.-Y. clusterProfiler: An R package for comparing biological themes among gene clusters. *Omics: A Journal of Integrative Biology 16*, 5 (2012), 284–287.
66. Zhang, J., Zhang, Y., Wang, J., Xia, Y., Zhang, J., and Chen, L. Recent advances in Alzheimer's disease: Mechanisms, clinical trials and new drug development strategies. *Signal transduction and targeted therapy 9*, 1 (2024), 211.

67. Zhang, X., Heng, Y., Kooistra, S. M., van Weering, H. R., Brummer, M. L., Gerrits, E., Wesseling, E. M., Brouwer, N., Nijboer, T. W., Dubbelaar, M. L., et al. Intrinsic DNA damage repair deficiency results in progressive microglia loss and replacement. *Glia 69*, 3 (2021), 729–745.
68. Zhang, X., Wang, X., Shivashankar, G., and Uhler, C. Graph-based autoencoder integrates spatial transcriptomics with chromatin images and identifies joint biomarkers for Alzheimer's disease. *Nature Communications 13*, 1 (2022), 7480.
69. Zhang, Z., Yang, C., and Zhang, X. scDART: Integrating unmatched scRNA-seq and scATAC-seq data and learning cross-modality relationship simultaneously. *Genome Biology 23*, 1 (2022), 139.
70. Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision* (2017), pp. 2223–2232.