NeuroLens: A Holistic Visual Analytics System for Exploring Brain Networks Across Scales

Weihan Zhang and Jun Tao, *Member, IEEE*

***Abstract*—Identifying biomarkers from human brain networks is critical in early detection of neurological disorders and understanding disease mechanisms. Existing visual analytics approaches show a remarkable ability to assist experts in discovering and validating biomarkers through exploration. However, these approaches often focus only on the diffusion features of fiber bundles and evaluate individual bundles and regions separately. This may hinder their ability to accurately represent the complex brain network for investigation from various perspectives. In this paper, we present NeuroLens, a visual analytics system that integrates comprehensive information across multiple levels, including fiber bundles, local regions and entire brain networks. Specifically, to model the bundles more precisely, we enhance the features by incorporating the joint distribution of geometric features. The bundle information is further aggregated to form representations at the region and the brain level using an attention-based graph neural network. The region-level representation describes complex structures involving multiple bundles, and the brain-level representation enables comparisons between subjects and groups. The NeuroLens interface enables comparative exploration of this multi-level and multi-faceted information. To verify the findings during exploration, NeuroLens leverages the large language model to query related information from existing literature. We collaborate with domain experts to examine the effectiveness of NeuroLens. Their exploration, findings, and feedback are discussed.**

***Index Terms*—Brain network, brain fiber, visual analytics**

I. Introduction

EXPLORING human brain networks through visual analytics can facilitate the discovery of novel medical insights, such as bio-markers for early detection of neurodegenerative diseases. It is typical that these diseases can only be mitigated, but not completely cured, highlighting the importance of early diagnosis [48]. Early clinical diagnosis often leverages biomarker detection to examine how the disease begins, offering effective disease management and early treatment opportunities [28]. In this context, a biomarker refers to an observable indicator of certain neurological diseases within brain networks. Recently, extensive research efforts, from both neuroscience and computer science, have investigated the identification of bio-markers in human brain networks. This kind of bio-marker identification research concerns a wide array of neurological diseases, especially Alzheimer's Disease (AD) [15], which affects 40 million patients globally [45]. Its findings are crucial for disease prevention and management.

W. Zhang is with the School of Computer Science and Engineering, Sun Yat-sen University, China. E-mail: zhangwh79@mail2.sysu.edu.cn.

J. Tao is with the School of Computer Science and Engineering, Sun Yat-sen University and the National Supercomputer Center in Guangzhou, China. E-mail: taoj23@mail.sysu.edu.cn. He is the corresponding author.

Biomarker detection generally targets structural brain networks, defined by nodes representing specialized brain Regions of Interest (ROIs) and edges denoted by neural fibers connecting the ROIs. To ensure biological relevance, ROIs are typically parcellated according to neuroanatomical labels or coordinated activities [56]. By comparing the brain networks of diseased subjects with control groups, researchers can pinpoint discriminative network connections, essential in identifying AD biomarkers. Such analysis requires a comprehensive representation and examination of the brain network, particularly the ROIs and their interconnections. However, traditional modeling of these connections over-simplifies fiber bundles, leading to significant information lost in the representation and diminishing their analytical utility. In addition, traditional biomarker detection often scores the ROIs and fiber bundles individually, overseeing their joint interactions. This oversight can mask collective anomalies involving multiple ROIs, hindering effective biomarker detection.

To address these issues, we propose NeuroLens, a visual analytics system that integrates deep neural representations with interactive visualization, serving as a lens for examining multiscale neuroanatomical structures in brain networks. This design enables the visual exploration and comparison of disease-relevant information across multiple scales, from the micro level of individual fiber bundles to the macro scale of whole-brain networks. At the level of individual fiber bundles, we enrich our analysis with a comprehensive feature set that captures a broad spectrum of information, including strength, Diffusion Tensor Imaging (DTI) features, geometric features, and their interactions. We use an Autoencoder to synthesize these features into a unified vector and facilitate the subsequent computation. For capturing regional information at larger scales involving multiple ROIs, we leverage the Graph Neural Network (GNN) to aggregate information from neighboring ROIs effectively. Moreover, we design a specialized detection module to extract from the GNN the significance of ROIs and their collective influence on the disease.

Our visual analytics interface enables a holistic exploration of brain networks, facilitating analysis from a broad overview down to specific details. At the global level, it allows domain experts to specify subject groups based on demographic/diagnostic labels and brain network representations for comparison. This comparison can be conducted between any two selected groups, enabling a focused investigation into the differences. The interface visually encodes the ROIs, their connections, and differences for experts to identify abnormal patterns as biomarkers. To delve deeper into these abnormalities, the interface offers tools to examine the associated fiber

bundles. Experts can analyze these bundles through the lens of their geometric feature distributions and DTI properties. Additionally, the interface provides 3D visualizations of the fiber tracts, which are critical for a thorough spatial understanding of brain connections [19], [20]. To supplement the exploration process, NeuroLens includes a literature query module that leverages large language models to access and synthesize information from existing research. This module enables users to validate their findings by generating insights based on the vast corpus of published neuroscience studies. We collaborated with experts in the field of neuroscience for an empirical expert evaluation to examine its effectiveness.

The main contributions of our work are summarized as follows:

- We develop NeuroLens, a holistic framework to represent and explore brain networks across multiple scales. Specifically, we employ a graph neural network to extract disease-relevant representations of brain ROIs and model their joint influence across neighborhoods of increasing size.
- We enrich fiber bundle representations by integrating DTI properties, geometric features, and their joint distributions, enabling more precise modeling of bundle characteristics.
- We design an analytical interface tightly coupled with these representations, allowing collective anomalies among regions and joint effects among bundle attributes to be discovered through multi-scale visual exploration.
- We introduce a large language model-driven literature query module, allowing the current findings to be validated with the existing body of knowledge in the literature. We compose a specialized dataset comprising brain network publications and associated image-text pairs, serving as a robust knowledge base for literature queries.

Taken together, these contributions highlight a development perspective in which neural networks are not used primarily as predictive tools with a visualization layered on top; instead, NeuroLens employs neural networks to systematically extract structured, multi-scale information guided by the needs of an explorative visual analytics interface. This tight coupling between representation learning and visualization enables experts to explore, compare, and validate brain connectivity patterns across scales, bridging learned representations and human-driven analysis in a way not supported by existing systems.

II. Related Work

Brain networks are typically studied at multiple levels, namely, individual level, brain connectivity level, and fiber level [12]. At the individual level, the brain networks are often examined as graphs using visual analytics interfaces. At the other two levels, the fibers are often analyzed as spatial curves. As our work explores both the fiber bundle representation and brain network comparison, this section briefly discusses both the brain network visualization in Section II-A and the brain fiber visualization in Section II-B.

A. Brain Network Visualization

Visual analytics for brain network. Visual analytics (VA) emerged as a significant tool for understanding human brain networks. The existing approaches often combine human exploration with machine learning techniques, especially feature selection. For example, Shi et al. [57] propose BrainQuest, which models brain network exploration as a multi-objective feature selection problem, jointly considering data mining and visual features for comparative visualization. Shi et al. [56] use the lasso model to select discriminative features for prediction. Their system also features adjustment mechanisms to leverage insights obtained through exploration. MV^{2}Net [55] also uses the lasso model for feature selection. This approach further considers the geometric features of fibers and the uncertainty in the data. Xu et al. [70] use a custom machine learning pipeline to guide the exploration. Beyond fibers and their derived features, their exploration includes the analysis space and patient groups. Similarly, Fujiwara et al. [23] propose a comparative visualization system, focusing on the difference between the individual and group mean.

Connectivity visualization of brain networks. Brain networks are commonly considered as graphs and visually represented by node-link diagrams and adjacency matrices. For node-link diagrams, various research efforts have been dedicated to adapting traditional network layouts to brain network scenarios. Node-link diagrams have also been adopted by Al-Awami et al. [3] and Troidl et al. [65] for visualizing large-scale brain connectivity. In this context, several interactive techniques [43], [54] have been developed to facilitate exploration. John et al. [41] introduce the ChordMap metaphor, an innovative approach for displaying hierarchical connectivity in brain networks. Salvador et al. [50] employ dimension reduction algorithms to uncover the functional similarity between ROIs. For adjacency matrices, they are widely used because brain networks are often large and dense. Alper et al. [4] conduct research on the impact of node-link and matrix representations in comparing weighted brain graphs. They recommended a specific overlaid matrix design in their findings. Wang et al. [68] apply multi-graph techniques to brain connectome datasets for classifying Alzheimer's Disease patients. This approach showed significant improvement by incorporating features selected through interactive means. Yang et al. [71] apply the NodeTrix representation to visual comparison of brain networks. Their approach extracts blocks from brain networks and represents the blocks using adjacency matrices. The blocks are further considered as nodes, and their connections are displayed as links.

Our system integrates a visual analytics interface for exploring the holistic information across scales. Similar to FiberStar [22] and MV^{2}Net [55], our approach targets visual comparison between multiple subjects. However, unlike the existing connectivity-based approaches, which often emphasize edge-centered examination and exploration, our approach aggregates connectivity information into region-level representations using GNNs, enabling exploration at different levels of abstraction. These GNN-based aggregation incorporates comprehensive features to support more structured and scalable

analysis beyond individual edges.

B. Brain Fiber Visualization

Brain fibers exhibit complex 3D structures, leading to severe visual clutter in rendering and serious challenges in interactive exploration. To address these issues, significant research efforts have been devoted to analyzing, visualizing, and exploring brain fibers [39].

Brain fiber clustering. Traditional fiber clustering methods leverage proximity measures between fibers [75] for clustering. Everts et al. [19] compute the similarity between streamlines based on the fiber tract segments. Goceri et al. [24] cluster fibers by matching distributions produced by the Gaussian mixed model. Franke et al. [22] propose FiberStars, allowing users to interactively differentiate between fiber clusters through the distance in 2D space. Deep learning-based approaches were also developed to group fibers for exploring connections between brain networks and neurodegenerative diseases. Gupta et al. [25], [26] propose FiberNet, using a convolution-based network to extract fiber shape features for clustering these fibers. In addition to the unsupervised frameworks, Lam et al. [36] design TRAFIC, which utilizes supervised signals to classify the brain fibers into anatomically significant groups. It outperformed previous approaches in accuracy and efficiency on multi-class tasks. Chen et al. [13] use a convolution network to embed brain fibers. Instead of the conventional reconstruction loss, their approach used the distance between fiber pairs as labels.

Brain fiber representation. Benou et al. [7] use a recurrent network to map the diffusion-weighted imaging sequences to probability distributions, simplifying the fiber tractography. Nath et al. [44] applied a residual network to learn the transformation between diffusion-weighted MRI signals and real structure. Hainline et al. [27] further introduce the bias and variance of generalized anisotropy fractions in a fully connected network. Li et al. [37] propose SuperDTI, which transfers a diffusion image to fiber tractography via a convolution network. SuperDTI can efficiently perform ultrafast diffusion tensor imaging and fiber tractography. Research on fiber orientation is an effective method for exploring brain connections. Zen et al. [74] propose FodNet, which consists of convolution layers and spherical harmonic coefficient blocks. FodNet obtains super-resolution fiber orientation direction images from DMRI images. It is the first work to generate super-resolution fiber orientation distribution from low-resolution images.

3D brain fiber visualization. For rendering brain fibers, Petrovic et al. [47] introduce a GPU-accelerated white matter tractography rendering system that generates fully shaded and labeled stream tube simulations. Brecheisen et al. [8] introduce an early stopping strategy using a specified threshold to track fibers without affecting their shapes. Merhof et al. [42] employ textured triangular strips for rendering large collections of dense fiber bundles. This approach supports the real-time rendering of high-quality and interactive fiber bundles. Researchers also embed brain fibers into a feature space for analysis. Brecheisen et al. [8] develop an interface with multiple coordinated views for selecting and visualizing fibers in both spatial and feature spaces. Similarly, Chen et al. [11] embed DTI fibers into a unified space, enabling the comparison of multiple datasets within this embedded space using heatmaps. Hurter et al. [32] introduce FiberClay, visualizing DTI fiber tracts in immersive environments.

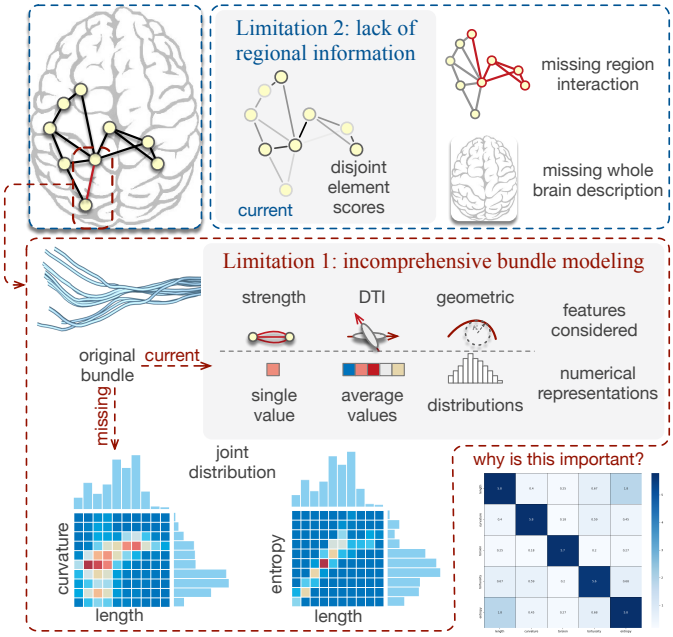

Fig. 1: Limitations of the existing approaches for extracting information from brain networks. The red rectangle shows the incomprehensive information used to represent the fiber bundles, and the blue rectangle shows the lack of joint interactions among ROIs and bundles for representing regional or whole brain information.

In terms of bundle modeling, our approach considers the geometric features beyond the traditional strength or DTI features [70] and further incorporates the connections among features, which are not available in existing approaches using geometric features [55]. By leveraging the GNN, it extracts the interrelations between brain regions, allowing the detection of collective anomalies involving multiple regions. Leveraging the rich information, our approach can support comparison at multiple scales, including subject groups, collective brain regions and bundles.

III. Overview

A. Background and Motivations

In this paper, we focus on the structural brain network, where nodes are the anatomical regions of the cerebral cortex. These nodes are interconnected by white matter fibers, which consist of clusters of axons that establish the network's edges. The regions of the brain, identified as Regions of Interest (ROIs), are consistently aligned across all subjects within a specific cohort, facilitating uniform analysis. The connectivity between these ROIs is predominantly estimated using diffusion-weighted magnetic resonance imaging, with

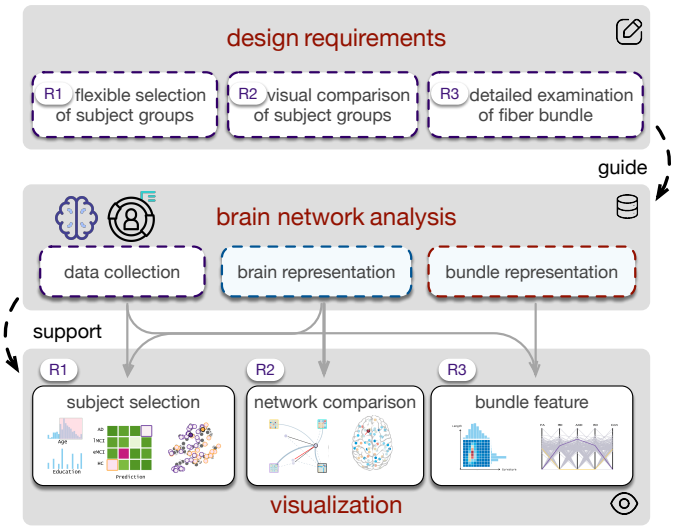

Fig. 2: Overview of the NeuroLens framework. Guided by three design requirements (R1–R3), the framework consists of brain network analysis and visualization components. Latent representations derived from raw data support three views: subject selection, network comparison, and bundle feature exploration.

Diffusion Tensor Imaging (DTI) serving as the primary method for this estimation [5].

The exploration of brain networks requires accurate representations at multiple levels, which may not be fully available. In Fig. 1, we illustrate two key limitations in existing representation approaches for brain networks.

First, the existing fiber bundle modeling is incomprehensive, as illustrated in the red rectangle. Early approaches [56], [71] only use one scalar number (the number of fibers) to represent the connection strength between two ROIs. Later approaches introduce DTI features [55], [70] as a vector and distributions of geometric features [55] derived from fiber tracts. However, these approaches ignore the interrelations among features, leading to unrevealed information. We provide a quantitative analysis of feature dependencies based on mutual information in Appendix E.

Second, the existing approaches do not extract holistic regional information at multiple scales. As shown in the blue rectangle in Fig. 1, the elements (i.e., fiber bundles and ROIs) are often rated individually, leading to a score of significance corresponding to one element. This element-wise evaluation may overlook the interactions among elements, hindering the discovery of collective anomalies. Additionally, missing the representation of the entire brain network may fail to provide hints for identifying subjects of interest during exploration.

B. Design Requirements

To better understand the research problem of brain network exploration, we teamed up with domain experts to identify design requirements. We invited six experts ($E1$–$E6$) with different research domains and experiences. $E1$ and $E2$ are neurology clinicians with sixteen and six years of experience in the clinical diagnosis of neurological diseases, respectively. Their research primarily focuses on AD, epilepsy, and Parkinson's disease. $E3$ is a neurologist with eighteen years of experience in cerebrovascular diseases and cognitive disorders. $E4$–$E6$ are senior PhD students with up to four years of research experience in neurological diseases.

To elicit design requirements, we conducted interviews with domain experts. The interviews began with an introduction of the motivations and preliminary design concepts to the domain experts. After we explained the brain network modeling process, we conducted open discussions to collect initial feedback. The experts then completed a questionnaire and provided additional elaboration on their responses.

We summarized the key user requirements for the brain network exploration task as follows. These user requirements are also aligned with classic interaction principles in information visualization [58].

R1: Flexible selection of subject groups. To identify unique connectivity bio-markers, experts often compare the group of subjects with the healthy controls. The subject groups with changing severity of disease are also compared with the healthy controls to understand the disease progression. The experts want to group the subjects based on various criteria such as demographics and diagnostic labels. They would also like to select individual subjects in these groups for further investigation. To support these tasks, the system needs to enable flexible selection by combining multiple criteria.

R2. Visual comparison of subject groups. In clinical practice, identifying connectivity bio-markers often relies on inspecting differences across subjects with the naked eye, which is subjective and prone to oversight. This limitation indicates the need for effective support to localize significant variations and provide interpretable insights into their underlying causes. Experts are therefore interested in understanding differences between subject groups from multiple perspectives. For instance, they seek to identify which regions or fiber bundles exhibit significant variations, and more importantly, to examine the underlying reasons why these regions or fibers are considered distinct.

R3. Detailed examination of fiber bundles. The experts believe that the brain fibers represent more details of brain connectivity, compared to the brain network with a single edge weight of connectivity strength. The analysis of brain connectivity should be conducted from multiple perspectives, on diffusion features at the voxel level and on geometric features representing the physical shape of brain fibers. Therefore, the domain experts want to examine the details of fiber bundles of interest. They would like to observe the features of a specified bundle and view the 3D bundles in the context of all brain fibers.

C. Design of System Framework

Driven by the requirements, we design the framework of the system, consisting of two major components: brain network analysis and visualization, as shown in Fig. 2. First, at the brain network analysis level, we organize the raw data and generate latent representations at both the brain and bundle

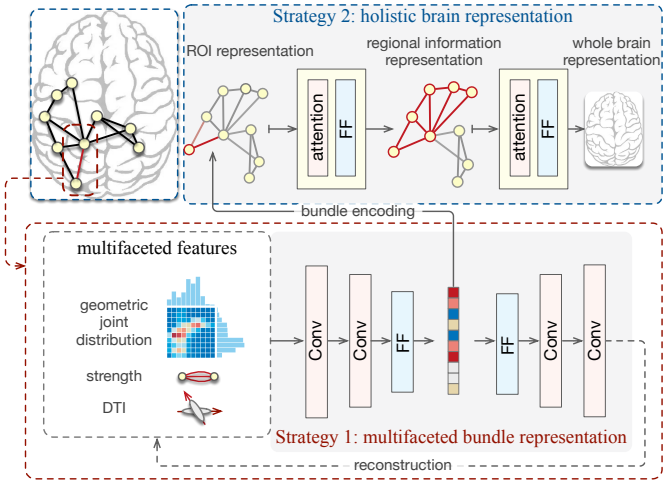

Fig. 3: An overview of our representation framework. The red rectangle shows the fiber bundle representation that takes the strength, DTI properties, and joint distributions of geometric features as input and outputs a latent vector as the representation. The blue rectangle shows the holistic brain representation, which aggregates the information over the brain network to form a representation at multiple scales, starting from the bundle representation.

levels. The brain representation provides an additional feature dimension to support subject selection (R1), while the bundle representation enables detailed analysis at multiple scales (R2). Second, at the visualization level, we support the design requirements through three main views: a subject selection view that combines demographic information with brain representations to specify subjects of interest for comparison (R1); a network comparison view that identifies regions of interest across scales (R2); and a bundle feature exploration view that facilitates fine-grained examination of fiber bundles (R3).

IV. Approach

Our approach employs a multi-scale representation strategy using the GNN to provide a foundation for exploration, as illustrated in Fig. 3. At the bundle scale, we integrate DTI and geometric features using an autoencoder. To capture information at larger scales, we leverage the GNN to iteratively collect and aggregate features from progressively larger neighborhoods, with the highest-level representation corresponding to the entire brain. In addition to these representations, we include a literature query module to support validation and contextualization of exploration results. In this section, we focus on the overall representation pipeline, while deferring model architectures, training details, and ablation studies to the appendices: Appendix A covers the bundle-level representation learned with the autoencoder, and Appendix B describes the brain region–level representation learned with the GNN.

A. Data Representation

Bundle representation. We consider a rich set of bundle features and transform the features into latent vectors to form bundle representations using an autoencoder [30].

(1) *Bundle data collection.* Both the diffusion features and geometric features are incorporated into our system. The diffusion features include fractional anisotropy (FA), mean diffusivity (MD), axial diffusivity (AxD), and radial diffusivity (RD), which are known to be strongly connected to neurological pathology [33], [46] and used in many existing visual analytics systems [8], [53], [55], [70]. We derive the geometric features from individual fibers, which are traced using the tractography algorithm [34]. These features reveal bundle characteristics beyond connection strength.

We consider five geometric features, namely, the length, curvature, torsion, tortuosity, and entropy, following [55]. While the curvature, torsion, and tortuosity convey shape-related information, the length and entropy are often correlated with uncertainty. Instead of representing each bundle with mean values over its fibers, we extend the existing approaches [55] in considering both the distributions of individual features and the joint distributions of pairs of features over the fibers of a bundle. As shown in Fig. 1, using the joint distribution to characterize the features of a bundle can reveal more information. According to correlation analysis, the distribution of features on specific bundles may have a strong correlation, requiring integration of these features. The example in Fig. 1 shows that while fibers generally have a longer distribution and smaller curvature, the curvature of a portion of the longer fibers has already begun to increase (within the red box). In our implementation, the joint distribution of two features is represented by a histogram of 20×20 bins and the mean and covariance matrix produced by the Gaussian mixture model (GMM). GMM captures the macroscopic statistical information of the bundle to ensure the overall pattern is captured.

(2) *Bundle encoding.* To integrate the multifaceted data collected from a fiber bundle, we represent the data in a unified latent space using an autoencoder, as shown in Fig. 3. The latent space captures essential structural patterns with a lower-dimensional representation, enabling more effective group discrimination compared to traditional metrics [30]. The autoencoder consists of an encoder and a decoder: convolutional layers are used to extract spatial structural features from the joint histograms, and linear layers embed macroscopic statistical information (e.g., the covariance matrix). The encoder progressively compresses these features to produce compact latent vectors, which reside in what is known as a bottleneck layer. The decoder then reconstructs the original input from this latent representation using upsampling linear layers. Through this reconstruction objective, the autoencoder learns to preserve meaningful features in the latent space that support comparison and downstream analysis. Full details of the data preprocessing, network architecture, training strategy, and ablation studies are provided in Appendix A.

Brain representation. The brain network of each subject is represented as a graph, where the nodes are brain regions of interest (ROIs) and the edges are fiber bundles between ROIs. This structure allows us to aggregate the bundle information to describe the individual ROIs, features across scales, and the entire brain Fig. 3. Given the graph structure, GNNs become a natural choice for such aggregation.

(1) *Representation loss.* Instead of conventional reconstruction loss, we adopt a classification loss as the training target for our representation model. The classification task requires GNNs to predict the diagnosis class of a subject given her brain network. The diagnosis classes include healthy controls (HC), early Mild Cognitive Impairment (eMCI), late Mild Cognitive Impairment (lMCI), and Alzheimer's Disease (AD) subjects. By focusing on classification loss, the network is forced to learn features critical to the disease diagnosis, thereby enhancing the diagnostic relevance of the regional representations.

(2) *Brain region encoding.* To represent information across scales, we first aggregate the bundle (edge) information to the corresponding ROIs (nodes), and then use the GNN to propagate information among ROIs. Different pooling strategies are used to summarize neighborhood information: mean pooling captures global trends, while max and min pooling emphasize locally salient signals. This section outlines the key design choices, while details of the GNN architecture, training procedure, and ablation studies are provided in Appendix B.

Our encoding model for brain regions leverages multiple graph attentional layers [9], [67] to aggregate information across scales. The input to graph attentional layers is a set of node features, $\mathbf{H} = \{\mathbf{h}_1^l, \mathbf{h}_2^l, \ldots, \mathbf{h}_N^l\}, \mathbf{h}_i^l \in \mathbb{R}$, where N is the number of nodes, and l is the index of the graph attention layer. The output is a new set of node features $\mathbf{H}^{l+1}$ by aggregating the information to each node from its neighbors. This can be viewed as an information collection procedure for representing features at various scales. Meaningful aggregation should weight the neighbors based on the relationships among nodes (ROIs), which is captured by learnable linear transformations, parameterized by weight matrices $\mathbf{W}$. We then perform self-attention on the nodes to obtain the importance weights of each node within its neighborhood:

$$\alpha_{ij} = \text{Softmax}_j\left((\mathbf{W}_q\mathbf{h}_i)^T \cdot \mathbf{W}_k\mathbf{h}_j\right), \tag{1}$$

where α_{ij} assesses the significance of node j's features with respect to node i. With these attention coefficients, a weighted linear combination of neighbor features produces the output features for each node:

$$\mathbf{h}_i^{l+1} = \sigma\left(\sum_{j\in N_i} \alpha_{ij}\mathbf{W}_v\mathbf{h}_j^l\right), \tag{2}$$

where N_i denotes the set of neighbors of node i. By applying a series of attention layers, the latent vector $\mathbf{h}_i^l$ gradually expands its vision to capture information at larger scales. By choosing the latent vectors produced at different layers, we can extract neighborhood information at different scales. The latent vector $\mathbf{h}^l$ is referred to as the *regional representation* at scale l. This representation aggregates information from neighborhoods of different extents (determined by l) around each ROI, encoding collective structural patterns rather than individual ROI properties. The visualization preserves the same anatomical ROIs across scales, while their representations reflect an increasingly global context. Finally, we can obtain the representation of the entire brain through global pooling followed by a simple linear classification head. This classification head enforces the extracted information to be closely related to the diagnosis of the disease.

Detection Module. Our detection module evaluates the global significance of the embedding features at various scales. This module uses the attention mechanism to capture the underlying interconnections among regional representations and their impact on the entire brain network. Inspired by Li et al. [38], we believe that the regional representation is associated with dense semantic information pertinent to disease diagnosis. Therefore, we utilize the detection module to recover diagnostic categories directly from the regional representations, noting that regions with higher attention weights contain more diagnostic information.

Specifically, to describe the significance of the regional representations and their interrelations, we use the attention mechanism from the transformer [66]. As illustrated in Fig. 3, the detection module consists of an attention layer and a feed-forward (FF) layer. Each of them is followed by a layer normalization (LayerNorm) and corresponds to a skip connection. In the attention layer, the interaction between regional representations is blended through an attention module, which updates the representations using:

$$\mathbf{h}_i' = \sum_j \text{attention}(\mathbf{h}_i, \mathbf{h}_j) \cdot \text{val}(\mathbf{h}_j). \tag{3}$$

The attention mechanism resembles the propagation of information along connected pathways in the brain network. To ensure the extracted interrelations are physically meaningful, we use the adjacency matrix of the brain network to extend the mask matrix used in the attention mechanism. The relationships between disconnected brain regions are masked, similar to Graphormer [72]. Moreover, we also leverage a classification token [18] connected to all brain regions, in order to extract a global embedding of the entire brain network for the subject.

Then, we utilize layer normalization. Since all tokens are normalized to a Gaussian distribution, the discrepancies resulting from inconsistent measurements can be diminished. In the FF layer, we use a multi-layer perceptron (MLP) to implement the forward embedding process. The features embedded through the FF layer are added to the original features, proceeding to the subsequent layer normalization calculation, i.e., $\text{FF}(\mathbf{H}') = \text{MLP}(\mathbf{H}) + \mathbf{H}$.

B. Literature Query

The literature query module allows users to immediately acquire related information from existing literature, so that they can verify their findings and refine their exploration directions. Toward this goal, the query module leverages a large language model (LLM) to process diverse natural language queries and retrieve relevant information from the literature. In our implementation, the LLM is used to fetch information from a knowledge base. Specifically, given the queries, the LLM is used to identify the most relevant sentences and figures in

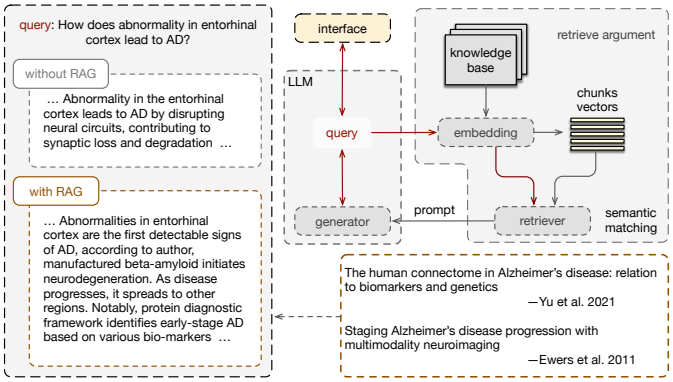

Fig. 4: The workflow of Literature Query. A query is first represented as a vector by the embedding module for semantic matching the vectors representing information of text chunks in the knowledge base. The matched literature chunks are passed as prompts into the LLMs, which compose an answer based on the extracted information and sends the answer to the NeuroLens interface.

existing publications. This strategy grounds the LLM's output in a curated knowledge base, helping to reduce the risk of hallucinated or fabricated information. This section explains how we collect the relevant papers and build the knowledge base, and how we align the LLM with the knowledge base.

Knowledge base construction. The knowledge base comprises over three hundred papers on brain fibers. We collected these papers from Google Scholar using a crawler script. To ensure relevance, we applied keyword filtering, selecting articles containing at least one of the terms: "brain fiber", "fiber bundle", "brain network", and "Alzheimer". We examined the publication years of the top 100 search results per keyword on Google Scholar and found that most papers were published after 2010 (e.g., 83 out of 100 for "brain network"). Based on this trend, we restricted the main publication range to 2010–present to focus on recent advances. To include foundational work, we also added highly cited papers from the 1990s with more than 2,000 citations; apart from the seminal study by Basser et al. [5] with over 5,000 citations, this threshold captures other influential early DTI studies that provide important theoretical and methodological context. To ensure the accuracy of figure-related queries, we manually process the papers to extract figures and captions. In the query stage, the LLM uses the captions to match, and the system can show both the caption and the corresponding figure.

Information Retrieval and answer generation. We employ the Retrieve Argument Generation (RAG) workflow for this task. RAG combines information retrieval and generative models, empowering the LLM to extract knowledge from extensive texts and incorporate the knowledge into the generated responses. This strategy has proven to be highly beneficial for LLMs. RAG is built upon the joint usage of two essential components: the *retriever* and the *generator*. The retriever first identifies and fetches the pertinent information from an extensive corpus of textual data, and then the generator produces answers based on the information supplied by the retriever.

In our current system, MPNet [59] is used as the retriever, and Llama2 [64] is used as the generator. MPNet is an innovative pre-trained language model. It incorporates both masked language modeling and permutation language modeling to develop a deeper understanding of sentence structures and word relationships than conventional BERT models [49]. With improved language comprehension abilities, MPNet is an ideal choice for understanding complex queries and finding the most relevant sentences in our knowledge base. Llama2 synthesizes the retrieved information into comprehensive answers. Llama2 exhibits two key advantages as a generator in our query module. Firstly, it adopts group query attention to accelerate inference speed, providing interactive performance for exploration. Secondly, Llama2 has 7 billion parameters and a context size of more than four thousand tokens, enabling the understanding of dialog context and supporting multiple subsequent queries. To improve efficiency, we use a virtual LLM (vLLM) tool [35] as an LLM serving engine. vLLM improves the throughput of Llama2, while maintaining the same level of latency. It also provides efficient memory management to handle the attention keys and values.

V. Visualization

A. Interface Design

Targeting the aforementioned requirements, we design NeuroLens to facilitate holistic exploration and comparison of brain networks of subjects or subject groups. Fig. 5 shows the interface of NeuroLens, consisting of five views, namely the subject selection view, network comparison view, bundle feature view, bundle rendering view, and literature query view. We start by introducing the typical workflow and then elaborate on each view in this section.

Typical workflow. NeuroLens adopts a hierarchical exploration workflow that guides users from a high-level overview of subject groups down to detailed analysis of specific fiber bundles. The process begins in the subject selection view, where users can define two groups based on their labels or brain network representations (R1), as shown in Fig. 5 (a). Users then explore and compare ROIs at different scales in the network comparison view Fig. 5 (b), identifying those that exhibit notable differences between groups and selecting related fiber bundles for further investigation (R2). Detailed information about the selected bundle is displayed in the bundle feature view and the bundle rendering view Fig. 5 (b) (d), providing both quantitative and anatomical insights (R3). Based on these insights, users can iteratively guide their exploration or validate findings using the literature query view Fig. 5 (e), which supports natural language search over a publication database.

Subject selection view. This view supports the flexible selection of subject groups (R1) in a cross-filtering manner, based on demographic attributes, diagnostic labels, and their global embedding positions, as shown in Fig. 5 (a). It summarizes subject-level information. On the left, two histograms display demographic distributions of age and years of education, supporting exploration of population-level trends. In

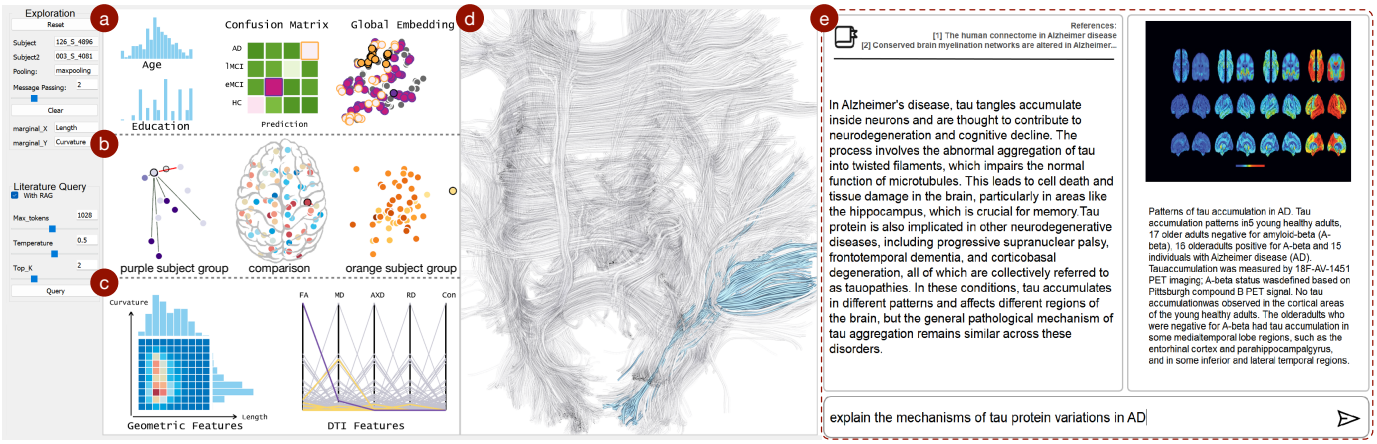

Fig. 5: The interface of our NeuroLens. (a) Subject selection view for sampling subjects from the dataset. (b) Network comparison view for comparing the embedded features of subjects at different scales. (c) The bundle feature view shows detailed information about fiber bundles. (d) Fiber render view for visualizing the fibers in physical space. (e) Literature query view for querying and answering questions using natural language.

the center, a confusion matrix encodes diagnostic prediction results, where each cell shows the number of subjects for a given combination of diagnostic labels and predicted labels. This design facilitates the selection of subject groups at different stages of disease progression. On the right, a scatterplot visualizes the learned brain embeddings, revealing the distribution of brains in the representation space. All charts support interactive brushing and linking. Users may brush the two histograms and the global embedding or click the cells in the confusion matrix for cross-filtering. Users may select two groups of subjects in this view, denoted by the purple and orange colors, respectively. This view helps to explore how demographic and diagnostic patterns relate to brain representations.

Network comparison view. This view displays the ROI representations of the two selected subject groups at a specified level (R2), as shown in Fig. 5 (b). Note that the level indicates the size of a neighborhood being considered. Therefore, each node in this view may correspond to a structure involving multiple ROIs in the brain. The two plots on the sides visualize node distributions for the purple and orange subject groups, respectively. These nodes are embedded into 2D screen space using t-SNE to highlight structural differences between groups. Node brightness encodes attention weights from the detection module, with darker nodes indicating stronger contributions to the diagnostic prediction. The middle plot overlays node-level differences onto a brain illustration. Each node depicts the difference between the same node in the two groups, with red nodes indicating large differences. The nodes in this plot are placed according to the locations of respective ROIs. This helps the experts build connections between the two plots on the sides. These plots together support the identification of potential bio-markers by localizing differences within the physical brain regions, bridging abstract model representations with interpretable anatomical insights.

We further incorporate a design inspired by NodeTrix [29] to support concise visualization and aligned comparison, as illustrated in Fig. 6. This hybrid representation retains node-link diagrams for visualizing local neighborhoods of interest, while embedding matrix plots to represent relationships within each of the four functional brain regions. Each matrix encodes the pairwise distances between ROI embeddings, providing a compact summary of structural similarity within a region. The solid links between nodes and the matrix reflect the connectivity strength between the ROI neighborhood and the functional regions, and the dashed links suggest the brain functional regions to which the nodes belong. We emphasize the functional regions because their damage is often associated with neurodegenerative diseases. By preserving both local connectivity and regional internal structure, this design enables inspection of ROIs within and across functional areas. This design facilitates comparison of patterns across subject groups, and may help reveal regional changes in brain organization that are associated with disease progression.

Bundle feature view. This view visualizes the feature distribution of a selected fiber bundle, supporting direct comparison of fiber-specific properties between subject groups (R3), as illustrated in Fig. 5 (c). On the left, the joint histogram encodes the bivariate distribution of two geometric features, complemented by marginal histograms that reveal the univariate distributions along each axis. This design captures the joint variation of features, providing a more intuitive characterization of fiber shape distributions. Users can identify structural trends that may not be apparent when features are considered in isolation. On the right, a parallel coordinates plot visualizes DTI features for the two selected subject groups. Lines are colored purple and orange to represent each group, respectively, enabling direct comparison of diffusion feature patterns. This design supports rapid identification of dominant differences in diffusion properties. These plots together support detailed pattern recognition of brain fiber bundles, facilitating the identification of structural differences within fiber bundles.

Fiber rendering view. This view presents a 3D rendering of all fibers, allowing domain experts to explore the data within its original anatomical context, as shown in Fig. 5 (d). The

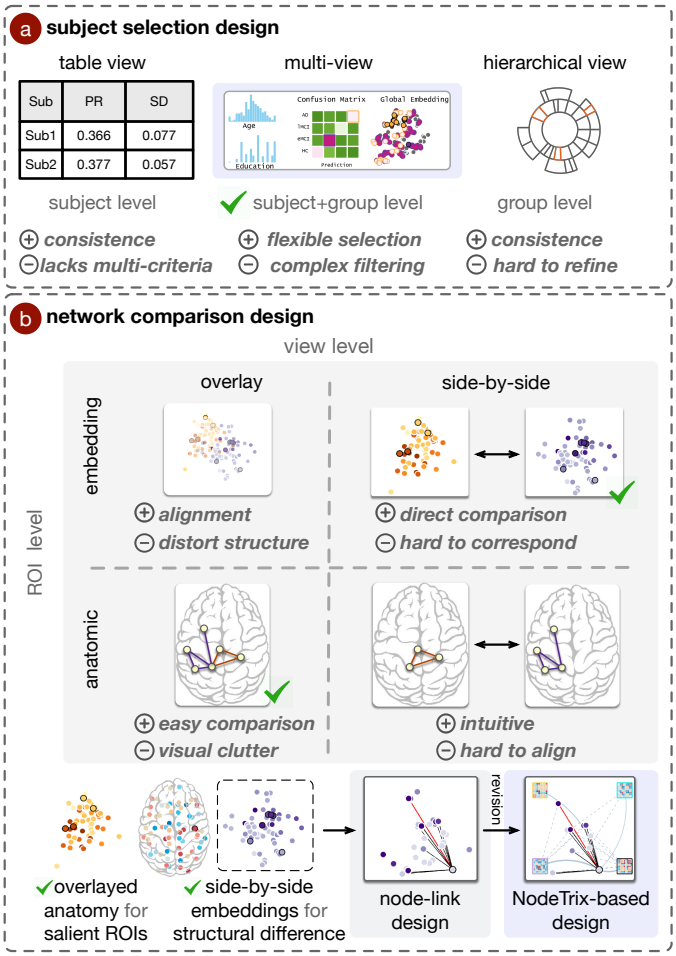

Fig. 6: Design alternatives and our corresponding choices for subject selection and network comparison.

selected fiber bundle is rendered in blue, while all other fibers are shown in semi-transparent gray to provide spatial reference without overwhelming visual attention. This contrastive design highlights the spatial extent and orientation of the selected bundle, facilitating detailed inspection of its shape, trajectory, and anatomical relationships. This visual contrast emphasizes the selected bundle, enabling clearer inspection and supporting detailed structural analysis (R3).

Literature query view. This view is only displayed on demand, when users click the "Query" button in the panel, as shown in Fig. 5 (e). Users may enter their query in natural language. The literature query module matches the relevant papers from our database, displaying their titles in the top area. The corresponding figure and its caption appear on the right. Additionally, the query module summarizes the relevant paragraphs in the matched papers and composes an answer, as displayed on the left of this view.

B. Alternative Designs

We discuss alternative designs and our rationales for critical components in our visual analytics interface. Since fiber bundle visualization follows established designs, we do not further elaborate on its alternatives here. Instead, we focus on the design choices for subject selection and network comparison, which play a central role in supporting flexible cohort definition and cross-scale analysis.

Subject selection. Common designs in existing visual analytics systems adopt tabular views to select individual subjects [56], [70], or apply hierarchical visualization designs for group selection based on diagnostic labels and demographic thresholds [55], as illustrated in Fig. 6 (a). While effective for targeted lookup and consistent cohort construction, these unified designs provide limited support for exploratory analysis. However, these unified designs visualize all attributes in the same manner, which is often ineffective for heterogeneous data types involving both scalar attributes and latent representations. Additionally, they either support group selection without sufficient ability for individual refinement, or focus on subject selection while making it difficult to express combined selection criteria. To avoid these drawbacks, we adopt a multi-view design for cross-filtering. This design allows users to select subject groups based on demographic information and diagnostic labels, and further refine selections using brain-level representations of individual subjects.

Network comparison. Our design rationale considers two factors: view-level layout (overlay versus side-by-side) and ROI-level representation (embedding versus anatomic space), as illustrated in Fig. 6 (b). Each combination has trade-offs: embedding views reveal relational patterns but not anatomic semantics, whereas anatomic views preserve spatial context but not group-specific structures. Overlay facilitates correspondence but can introduce clutter for anatomic views or damage structures in individual groups for embedding views, while side-by-side improves separation but requires extra alignment effort. No single combination suffices, motivating a hybrid design. Therefore, we adopt a hybrid anatomical design in which an overlay view at the center serves as a stable reference to connect the two groups, while the groups themselves are visualized separately to observe their internal structures. This hybrid design provides both a coherent reference for comparison and the flexibility to explore group-specific details.

Additionally, for embedding visualization, expert feedback indicates that detailed analysis typically focuses on selected ROIs and their immediate neighbors rather than the full network. To support this analytic pattern, we employ a hybrid design based on NodeTrix: functional regions are shown as matrices for easy correspondence between groups, while selected ROIs and their neighbors are visualized as nodes and links to highlight important connectivity structures. By combining these representations, the hybrid design allows users to simultaneously capture high-level patterns across functional regions and inspect detailed local connectivity, providing a balance between overview and detail that would not be possible with a single representation.

VI. Evaluation

We collaborated with six domain experts in neuroscience to examine the effectiveness of NeuroLens. The same experts described in Section 3 participated in this evaluation. This study focuses on expert feedback on a visualization system

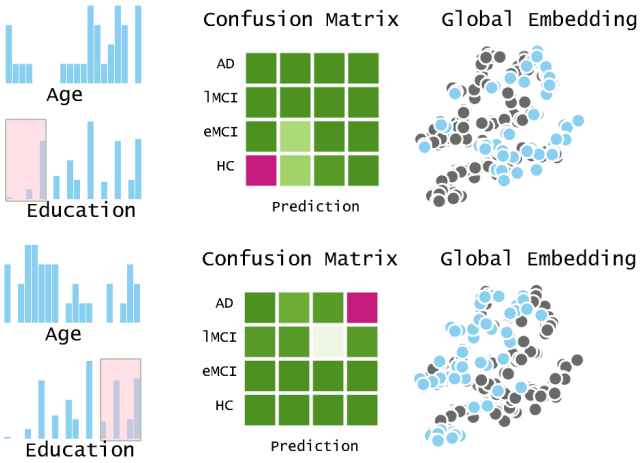

Fig. 7: Selecting the subjects with shorter and longer years of education in the subject selection view. The top row shows the diagnostic label and global embedding of the subjects with shorter education experience, and the bottom row shows the information of the subjects with longer education.

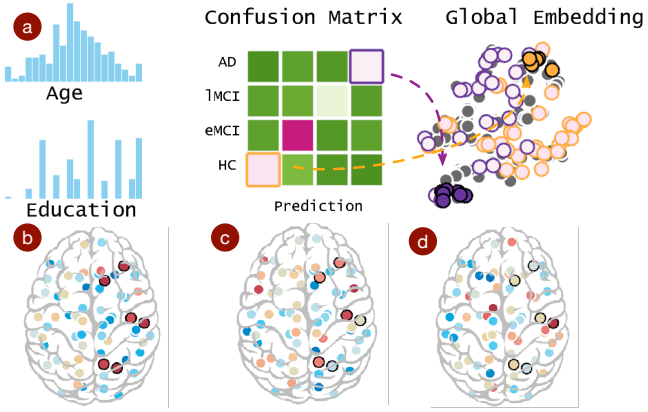

Fig. 8: Comparison of AD and HC subjects at different scales. (a) shows the subject selection view where the HC and AD subjects are selected. (b) shows the network comparison view at scale two, where the regions exhibiting significant differences (red nodes) can be observed. (c) and (d) compare the differences observed at scale one and two, respectively, where the differences are less pronounced.

and does not involve patient participation or intervention. The dataset used in this study is publicly available and de-identified. All domain experts provided informed consent prior to participation. We present their exploration process, findings, and feedback in this section.

Dataset. We use a public dataset released by the Alzheimer's Disease Neuroimaging Initiative (ADNI), a public consortium focused on the collection, validation, and use of AD data [2]. This dataset includes structural MRIs from a varied group of subjects, who were recruited and scanned at 16 different sites throughout North America. FreeSurfer [1] provides the segmentation method to generate the structural brain network for each subject. The segmentation follows the Desikan-Killiany parcellation template [17]. The parcellation template defines 70 ROIs across the cortex of each subject's whole brain, with 35 ROIs in each cerebral hemisphere [15]. The fiber bundles are generated between the ROIs and the bundle features are derived from corresponding fibers. The default edge weight between two ROIs is defined as the number of white matter fiber tracts connecting them, estimated from each subject's DTI data.

Evaluation procedure. We invited domain experts to explore the ADNI dataset using NeuroLens for exploratory analysis. To guide the exploration, we suggested two focal comparisons: one at the group level between AD and HC subjects, and one at the individual level between an lMCI subject and a similar subject from the healthy control. We encouraged the experts to freely explore the dataset without time constraints. We recorded their exploration processes, observations, and insights and later organized the case studies presented below.

A. Case Study 1: Comparing AD and HC Groups

Exploration goals. Experts $E1$ and $E2$ were interested in the relationships between age, years of education, and AD diagnostic outcomes, as research indicates that age and educational attainment significantly impact AD progression [21]. They explained that older individuals are more likely to show cognitive symptoms at a lower level of molecular pathology, attributed to their reduced functional reserve. Similarly, lower educational levels are associated with a heightened risk of cognitive decline, underscoring the combined effect of biological aging and socio-environmental factors such as education on AD's development and progression. While aging inevitably leads to functional degradation of the brain, higher education levels can fortify internal brain connections, potentially mitigating brain aging effects. Leveraging the cross-filtering selection, they could examine the interaction among these factors.

Relationships between age, education, and diagnostic labels. As shown in Fig. 7, when selecting individuals with higher levels of education, the confusion matrix indicated that these individuals were predominantly diagnosed with AD and lMCI. In contrast, when selecting individuals with lower levels of education, the confusion matrix showed that these individuals were more likely to belong to the healthy control group and were diagnosed with eMCI. Sliding the selection window on the education histogram also reveals different distributions of age, indicating that most subjects from the ADNI datasets are either older with lower levels of education or younger with longer education experience. Fig. 7 demonstrates the minor correspondence between years of education and the global embedding of brains as well.

Identifying the most distinct ROIs between AD and HC subjects. The experts then selected the AD and HC subjects for comparison in the confusion matrix, as shown in Fig. 8 (a). To emphasize the difference, they brushed the global embedding to select the two most distant groups. Each group corresponds to a distinct label, as indicated by the orange and purple nodes, respectively. In Fig. 8 (b), (c), and (d),

the network comparison view showed the group differences at scales two, one, and three, respectively. The expert found that the difference became obvious at scale two, as denoted by several dark red nodes in (b). This indicated that the collective anomaly appeared at this scale. In (c) and (d), the abnormal pattern was either not formed at scale one or smoothed out at scale three.

Therefore, the experts used scale two for their analysis, which showed that abnormal regions appeared on the right, which corresponded to the left hemisphere as the brain was viewed from bottom to top. They commented that this asymmetry was more common in neurodegenerative diseases. The asymmetric progression of AD was previously documented in clinical research, indicating that pathological changes manifest earlier in the left hemisphere than in the right [63]. $E1$ hypothesized that the asymmetry phenomenon might be attributed to the left brain dominance, as the majority of people are right-handed. Therefore, abnormalities in the left hemisphere become more visible. While $E4$ and $E5$ agreed that the difference in the decline of the left and right brain is common in group statistics, they also stated that occasional exceptions exist. $E1$ further explained that the abnormal brain regions in Fig. 8 (b) were mostly distributed in areas such as the frontal and temporal lobes, and pointed out that these brain regions were particularly rich in projection neurons to other brain locations, which could easily lead to global brain functional damage.

Validation from literature query module. The experts further examined the discriminative ROIs (red nodes) in the left hemisphere in Fig. 8 (b). They explained that the left precuneus and superior parietal brain regions exhibited white matter abnormalities in AD, which correlated with the intra-hemispheric connectivity changes observed in our network comparison view. The experts also observed a decrease in connectivity between the precuneus and superior parietal regions. $E1$ searched this phenomenon with the literature query module, and found that it was reported in prior research [61]. $E1$ found the literature query to be particularly useful, and used this module to further confirm the significant involvement of the precuneus and superior parietal regions [14] and the alterations of fiber connectivity in the temporal and parietal lobes, predominantly in the left hemisphere [62], [71].

B. Case Study 2: Comparing lMCI and HC Subjects

Exploration goals. The experts wanted to compare the lMCI subjects with the healthy control group. They explained this might identify candidates of indicative bio-markers for early diagnosis and intervention strategies, which were critical to delay the transition to AD [69].

Identifying the regional differences. The experts examined the regional differences at various scales by adjusting the graph convolutional level. In the network comparison view in Fig. 9, they found that the difference becomes more obvious at scale-three (c) than scale-two (a), as more red nodes appear in the scatterplot in the middle view. The experts selected a red node exhibiting significant differences for further investigation. $E4$ stated that cognitive functions typically involve multiple brain

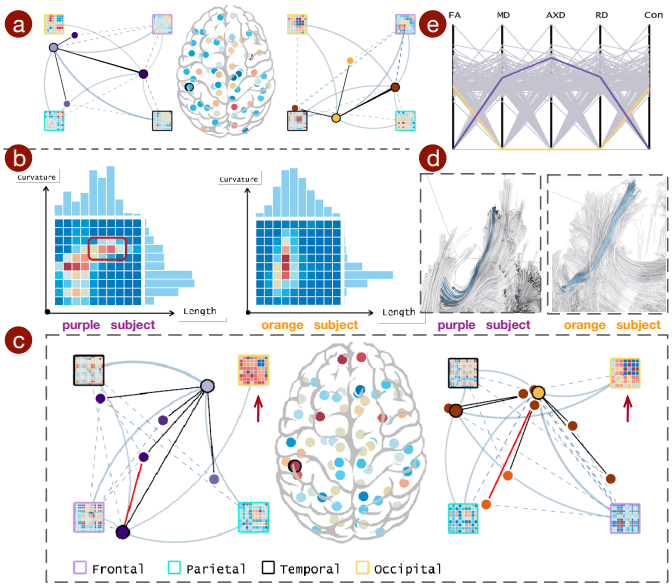

Fig. 9: Comparison between an lMCI subject and an HC subject. (a) and (c) show the network comparison view at scale-two and scale-three, respectively. A bundle connected to the ROI is selected in (c). (b) and (e) show the selected bundle's geometric features and DTI features, respectively, where the lMCI subject is shown in purple and the HC subject is shown in orange. (e) renders the selected fiber bundle of the lMCI subject.

regions affected by distant damage, and considering non-local features is necessary.

The distant correspondence can also be observed from Fig. 9 (c), where a large red block appears in the matrix of the occipital lobe for the lMCI subject, as indicated by the red arrows. The occipital lobe is adjacent to the temporal function region containing the selected ROI, forming a collective abnormality. $E1$ believed that abnormalities in the temporal and occipital lobe regions usually manifest as visuospatial impairments. She presented us with a similar case where a patient holding his phone could not articulate what it was. This is because the patient's visual pathway is unable to transmit signals through the occipital lobe. $E4 - E6$ explained that the occipital lobe typically causes different types of dementia in addition to AD, such as Lewy body dementia. Using the literature query module, the experts found that this potential bio-marker can be linked to several clinical evidences [31], [51]. These papers suggest that the number of synapses is highly correlated with the Mini-Mental State Examination (MMSE) score, a commonly used test for evaluating cognitive function in older adults. The experts examined the MMSE score of this lMCI subject and found her score (25) was below the normal level, indicating cognitive impairment.

Identifying the fibers' subtle change in the bundle. After exploring the white matter regions, experts wished to study the detailed characteristics of the fiber bundles linking the ROIs. They examined the brain fibers located in the white matter of the temporal functional region. They selected the

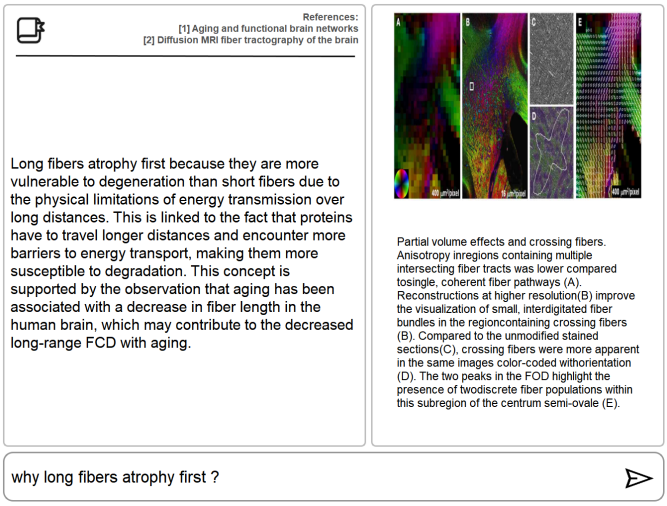

Fig. 10: The literature query view in NeuroLens. Retrieving expressions related to the atrophy of long fibers in the literature and the literature query view returns the relevant literature found and generates an answer. It suggests that the changes in long fibers may be due to energy transfer.

fiber bundle between the deep purple nodes in the comparison view, as shown in Fig. 9 (c). The bundle views in Fig. 9 (b) and (d) indicate that the length of these fibers is moderately long and their curvature mainly stays at a slightly low level. However, in the joint distribution of the lMCI subjects, several anomalous bins appear in the upper right region (within the red solid box), indicating long fibers with large curvature. This could be considered as atrophy occurring in the selected fiber bundle. Experts examined the brain fibers corresponding to these bins for both subjects. The fiber render view in Fig. 9 (d) showed that the bundle of the purple subject exhibited atrophy at one end, which did not exist in the orange subject. The experts further examined the DTI features of this bundle. In Fig. 9 (e), the atrophied fibers (the purple line) exhibit an observed decline in the FA value, suggesting potential damage to white matter integrity. These facts confirmed that the atrophy is primarily attributed to developmental processes rather than focal lesions or age-related degeneration, in line with the existing finding [52]. As the purple subject is the lMCI, the experts suspected that the longer fibers were more vulnerable to atrophy in the early stages of AD.

Explanation from literature query module. The experts were intrigued by the long fiber atrophy phenomenon and used the literature query module for further investigation, as shown in Fig. 10. They confirmed that typical connectivity damage was due to the development of abnormalities in the cell bodies of upstream neurons. In this process, abnormal accumulation of τ protein forms neurofibrillary tangles, thereby damaging the microtubule structure of neuronal cell bodies. $E3$ pointed out that when the high signal in white matter was not prominent, it was usually accompanied by atrophy of nerve fibers and microstructural damage (revealed by DTI features).

C. Expert Feedback and Discussion

After the exploration and evaluation of case studies, we interviewed the experts for feedback on the current functionality of NeuroLens and potential expansions.

The experts found that NeuroLens is effective primarily as a tool for exploratory sense-making. $E1$ and $E2$ noted that the multi-scale network comparison view enabled them to observe structural differences from multiple perspectives, which inspired hypotheses about potentially abnormal regions. They emphasized that even when the learned representations were not fully interpretable, the observed differences still served as useful hints for further investigation, which they could then verify through the fiber rendering view or by consulting existing literature. The joint distribution of geometric features was similarly valued as a cue for directing attention to candidate fiber bundles. $E3$ further commented that the system's ability to surface abnormal structural patterns at the network level could support early warning analyses by highlighting regions and connections for closer inspection. Compared to conventional MRI-based analysis focused on localized regions, the experts highlighted that NeuroLens enables exploration across scales, supporting hypothesis generation that is difficult to achieve with existing tools.

In terms of usability, the experts noted that although they did not fully understand the underlying mechanisms of graph neural networks, they could easily identify suspicious bundles through visual representations. For example, $E6$ mentioned that the interface was intuitive and allowed her to conveniently compare brain networks between subjects. She explained that she followed the same procedure across case studies: selecting subjects, identifying abnormal ROIs, and examining bundles. As a result, she felt comfortable using NeuroLens after completing the first case study. However, she also noted that assistance was needed during the initial learning phase to fully appreciate the system's features. This suggests the value of a guided example embedded in the system to support users' onboarding.

Limitations. The experts identified several limitations of the current system. First, they noted that geometric features might not always be effective in distinguishing abnormal bundles. This limitation arises because geometric features are computed at the individual fiber level and aggregated at the bundle level, which may not fully capture the overall bundle shape. Despite this, the experts acknowledged that these features still provided valuable spatial context and were helpful in specific cases. Second, they emphasized that identifying reliable biomarkers generally requires long-term studies that go beyond the scope of our tool. Although our system reveals detailed structural differences not available in existing clinical tools, determining whether those differences carry biological significance remains challenging. Moreover, although the literature query module helped verify findings against existing research, it could not lead to discoveries that had not yet been discussed.

We also identify several limitations in the current system ourselves. Some of these have been partially addressed through iterative design refinement. For example, early versions of the network comparison view exposed the full embedding

structure, although experts primarily focused on selected ROIs and their immediate neighbors. Based on observations, we revised the design to use a node-link representation for selected ROIs and their neighbors, while summarizing other regions using the NodeTrix-based representation. This hybrid design better aligns with expert analytic practices and reduces visual complexity.

Other limitations remain. A key challenge is the lack of explicit semantic grounding in the current representations. Although GNN-based embeddings capture complex structural patterns, their semantic meaning is implicit, which limits interpretability. This limitation also constrains the integration of large language models. Currently, LLMs are mainly used for literature queries and do not have access to the learned data representations. As a result, they can support conceptual questions but cannot reason directly over the data. Bridging this semantic gap would enable more data-aware question answering and deeper integration of LLMs into the visual analytics workflow.

Revision suggestions. The experts also offered suggestions for refining the tool from clinical practice perspectives. $E1$ and $E2$ noted that filtering by age and educational background is not commonly used in clinical settings. Such information is typically referenced only for cognitive ability assessments. However, we have chosen to retain these components in the interface, as they do not interfere with clinical diagnosis and may be valuable in research contexts for exploring the influence of demographic factors. $E4$–$E6$ recommended incorporating additional brain region segmentation schemes, as segmentation criteria can vary across studies or institutions. This can be addressed by re-labeling the fibers into different sets of bundles, and we will continue collaborating with the experts to implement these revisions.

VII. Conclusion

Identifying bio-markers in the brain network is an intricate task. Our work introduces NeuroLens, a visual analytics system for the identification and analysis of biomarkers within the human brain network. By addressing the limitations of existing methodologies that primarily focus on the single features of fiber bundles in isolation, NeuroLens offers a more holistic and integrated approach. It leverages comprehensive techniques, including the joint distribution of geometric features and graph neural networks, coupled with attention mechanisms, to model and analyze brain networks at various levels of complexity. The literature query view powered by an LLM helps align findings with the relevant research, reinforcing the validity and relevance of the exploration process.

In the future, first, we will continue to explore how NeuroLens can embrace the advancements in diffusion MRI, such as multiple fiber tractography, Kurtosis imaging, and fixed-based analysis. We will explore incorporating these methods at the data level by generating more accurate fiber models or converting their outputs into fiber attributes. We will also develop novel visual analytics interfaces and exploration workflows tailored to these specific approaches. Second, we aim to more tightly integrate LLM with the representation and exploration process. With improved semantic alignment between our current representation and domain concepts, LLMs could reason over the data itself rather than serving solely as a literature query tool. For instance, the system could highlight task-relevant ROIs based on natural language queries, suggest next-step exploration paths, or identify semantic patterns from embeddings. Such integration could enhance interpretability, reduce cognitive load, and support more data-aware visual analysis.

Acknowledgements

This research was supported in part by the National Natural Science Foundation of China through grants 62372484 and 62172456. The authors would like to thank the anonymous reviewers for their insightful comments.

References

[1] FreeSurfer, 2012. [Online]. Available: http://surfer.nmr.mgh.harvard.edu/.

[2] Alzheimer's Disease Neuroimaging Initiative, 2020. [Online]. Available: http://adni.loni.usc.edu/.

[3] A. K. Al-Awami, J. Beyer, H. Strobelt, N. Kasthuri, J. W. Lichtman, H. Pfister, and M. Hadwiger. NeuroLines: A subway map metaphor for visualizing nanoscale neuronal connectivity. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2369–2378, 2014. doi: 10.1109/TVCG.2014.2346312

[4] B. Alper, B. Bach, N. Henry Riche, T. Isenberg, and J. Fekete. Weighted graph comparison techniques for brain connectivity analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 483–492, 2013. doi: 10.1145/2470654.2470724

[5] P. J. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66(1):259–267, 1994. doi: 10.1016/S0006-3495(94)80775-1

[6] C. Beaulieu. The basis of anisotropic water diffusion in the nervous system — a technical review. *NMR in Biomedicine*, 15(7-8):435–455, 2002. doi: 10.1002/nbm.782

[7] I. Benou and T. Riklin-Raviv. DeepTract: A probabilistic deep learning framework for white matter fiber tractography. In *Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 626–635, 2019. doi: 10.1007/978-3-030-32248-9_70

[8] R. Brecheisen, A. Vilanova, B. Platel, and B. ter Haar Romeny. Parameter sensitivity visualization for DTI fiber tracking. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1441–1448, 2009. doi: 10.1109/TVCG.2009.170

[9] S. Brody, U. Alon, and E. Yahav. How attentive are graph attention networks? *Proceedings of International Conference on Learning Representations*, 2022. doi: 10.48550/arXiv.2105.14491

[10] T. Cai, S. Luo, K. Xu, D. He, T. Liu, and L. Wang. GraphNorm: A principled approach to accelerating graph neural network training. In *Proceedings of the International Conference on Machine Learning*, pp. 1204–1215, 2021. doi: 10.5555/3461702.3461745

[11] H. Chen, H. Mei, Z. Liu, S. Yan, and W. Chen. Visualizing differences of DTI fiber models using 2D normalized embeddings. In *Proceedings of the IEEE Pacific Visualization Symposium*, pp. 350–351. IEEE, 2014. doi: 10.1109/PacificVis.2014.68

[12] W. Chen, L. Shi, and W. Chen. A survey of macroscopic brain network visualization technology. *Chinese Journal of Electronics*, 27(5):889–899, 2018. doi: 10.1049/cje.2018.04.007

[13] Y. Chen, C. Zhang, T. Xue, Y. Song, N. Makris, Y. Rathi, W. Cai, F. Zhang, and L. J. O'Donnell. Deep fiber clustering: Anatomically informed fiber clustering with self-supervised deep learning for fast and effective tractography parcellation. *NeuroImage*, 273:120086, 2023. doi: 10.1016/j.neuroimage.2023.120086

[14] M. Dadar, A. L. Manera, S. Ducharme, and D. L. Collins. White matter hyperintensities are associated with grey matter atrophy and cognitive decline in Alzheimer's disease and frontotemporal dementia. *Neurobiology of Aging*, 111:54–63, 2022. doi: 10.1016/j.neurobiolaging.2021.11.007

[15] M. Daianu, N. Jahanshad, T. M. Nir, A. W. Toga, C. R. J. Jack, M. W. Weiner, P. M. Thompson, and A. D. N. Initiative. Breakdown of brain connectivity between normal aging and Alzheimer's disease: A structural k-core network analysis. *Brain Connectivity*, 3(4):407–422, 2013. doi: 10.1089/brain.2012.0137

[16] R. Della Nave, A. Ginestroni, S. Diciotti, E. Salvatore, A. Soricelli, and M. Mascalchi. Axial diffusivity is increased in the degenerating superior cerebellar peduncles of friedreich's ataxia. *Neuroradiology*, 53(5):367–372, 2011. doi: 10.1007/s00234-010-0807-1

[17] R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, et al. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, 31(3):968–980, 2006. doi: 10.1016/j.neuroimage.2006.01.021

[18] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16×16 words: Transformers for image recognition at scale. In *Proceedings of the International Conference on Learning Representations*, pp. 1–17, 2021. doi: 10.48550/arXiv.2010.11929

[19] M. H. Everts, E. Begue, H. Bekker, J. B. T. M. Roerdink, and T. Isenberg. Exploration of the brain's white matter structure through visual abstraction and multi-scale local fiber tract contraction. *IEEE Transactions on Visualization and Computer Graphics*, 21(7):808–821, 2015. doi: 10.1109/TVCG.2015.2403323

[20] M. H. Everts, H. Bekker, J. B. T. M. Roerdink, and T. Isenberg. Depth-dependent halos: Illustrative rendering of dense line data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1299–1306, 2009. doi: 10.1109/TVCG.2009.138

[21] M. Ewers, G. B. Frisoni, S. J. Teipel, L. T. Grinberg, E. J. Amaro, H. Heinsen, P. M. Thompson, and H. Hampel. Staging Alzheimer's disease progression with multimodality neuroimaging. *Progress in Neurobiology*, 95(4):535–546, 2011. doi: 10.1016/j.pneurobio.2011.06.004

[22] L. Franke, D. K. I. Weidele, F. Zhang, S. Cetin-Karayumak, S. Pieper, L. J. O'Donnell, Y. Rathi, and D. Haehn. FiberStars: Visual comparison of diffusion tractography data between multiple subjects. In *Proceedings of the IEEE Pacific Visualization Symposium*, pp. 116–125, 2021. doi: 10.1109/PacificVis52677.2021.00023

[23] T. Fujiwara, J.-K. Chou, A. M. McCullough, C. Ranganath, and K.-L. Ma. A visual analytics system for brain functional connectivity comparison across individuals, groups, and time points. In *Proceedings of the IEEE Pacific Visualization Symposium*, pp. 250–259, 2017. doi: 10.1109/PACIFICVIS.2017.8031601

[24] E. Goceri. Intensity normalization in brain MR images using spatially varying distribution matching. In *Proceedings of the International Conference on Computer Graphics, Visualization, Computer Vision and Image Processing*, pp. 300–304, 2017.

[25] V. Gupta, S. I. Thomopoulos, C. K. Corbin, F. Rashid, and P. M. Thompson. FiberNET 2.0: An automatic neural network-based tool for clustering white matter fibers in the brain. In *Proceedings of the IEEE International Symposium on Biomedical Imaging*, pp. 708–711, 2018. doi: 10.1109/ISBI.2018.8363672

[26] V. Gupta, S. I. Thomopoulos, F. M. Rashid, and P. M. Thompson. FiberNET: An ensemble deep learning framework for clustering white matter fibers. In *Medical Image Computing and Computer-Assisted Intervention*, pp. 548–555. Springer, 2017.

[27] A. E. Hainline, V. Nath, P. Parvathaneni, K. G. Schilling, J. A. Blaber, A. W. Anderson, H. Kang, and B. A. Landman. A deep learning approach to estimation of subject-level bias and variance in high angular resolution diffusion imaging. *Magnetic Resonance Imaging*, 59:130–136, 2019. doi: 10.1016/j.mri.2019.03.021

[28] O. Hansson. Biomarkers for neurodegenerative diseases. *Nature Medicine*, 27(6):954–963, 2021. doi: 10.1038/s41591-021-01382-x

[29] N. Henry, J. Fekete, and M. J. McGuffin. NodeTrix: A hybrid visualization of social networks. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1302–1309, 2007. doi: 10.1109/TVCG.2007.70582

[30] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006. doi: 10.1126/science.1127647

[31] P. R. Hof, J. H. Morrison, and K. Cox. Quantitative analysis of a vulnerable subset of pyramidal neurons in Alzheimer's Disease: I. superior frontal and inferior temporal cortex. *Journal of Comparative Neurology*, 301(1):44–54, 1990. doi: 10.1002/cne.903010105

[32] C. Hurter, N. H. Riche, S. M. Drucker, M. Cordeil, R. Alligier, and R. Vuillemot. FiberClay: Sculpting three dimensional trajectories to reveal structural insights. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):704–714, 2019. doi: 10.1109/TVCG.2018.2865191

[33] Y. Jin, C. Huang, M. Daianu, L. Zhan, E. L. Dennis, R. I. Reid, C. R. Jack Jr., H. Zhu, P. M. Thompson, and A. D. N. Initiative. 3D tract-specific local and global analysis of white matter integrity in Alzheimer's disease. *Human Brain Mapping*, 38(3):1191–1207, 2017. doi: 10.1002/hbm.23448

[34] D. K. Jones. Tractography gone wild: Probabilistic fibre tracking using the wild bootstrap with diffusion tensor MRI. *IEEE Transactions on Medical Imaging*, 27(9):1268–1274, 2008. doi: 10.1109/TMI.2008.922191

[35] W. Kwon, Z. Li, S. Zhuang, Y. Sheng, L. Zheng, C. H. Yu, J. Gonzalez, H. Zhang, and I. Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the Symposium on Operating Systems Principles*, pp. 611–626, 2023. doi: 10.1145/3600006.3613165

[36] P. D. N. Lam, G. Belhomme, J. Ferrall, B. Patterson, M. Styner, and J. C. Prieto. TRAFIC: Fiber tract classification using deep learning. In *Proceedings of the SPIE International Society for Optical Engineering Medical Imaging – Image Processing*, pp. 257–265, 2018. doi: 10.1117/12.2293931

[37] H. Li, Z. Liang, C. Zhang, R. Liu, J. Li, W. Zhang, D. Liang, B. Shen, X. Zhang, Y. Ge, et al. SuperDTI: Ultrafast DTI and fiber tractography with deep learning. *Magnetic Resonance in Medicine*, 86(6):3334–3347, 2021. doi: 10.1002/mrm.28937

[38] T. Li, D. Katabi, and K. He. Return of unconditional generation: A self-supervised representation generation method. In *Proceedings of the Conference on Neural Information Processing Systems*, pp. 125441–125468, 2024. doi: 10.52202/079017-3985

[39] Y. Liu, C. Xu, Z. Jiang, L. Jiang, Y. Feng, and R. Liang. A survey on brain fiber visualization. *Journal of Computer-Aided Design and Computer Graphics*, 30(1):9–19, 2018. doi: 10.3724/SP.J.1089.2018.16927

[40] X. Ma, X. Dai, Y. Bai, Y. Wang, and Y. Fu. Rewrite the stars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5694–5703, 2024.

[41] J. McGonigle, M. Mirmehdi, et al. Visualising functional connectivity in fMRI using hierarchical edge bundles. In *Proceedings of the Annual Meeting of the Organization for Human Brain Mapping*, 2011.

[42] D. Merhof, M. Sonntag, F. Enders, C. Nimsky, P. Hastreiter, and G. Greiner. Hybrid visualization for white matter tracts using triangle strips and point sprites. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1181–1188, 2006. doi: 10.1109/TVCG.2006.151

[43] E. Moerth, B. Beinder, R. Raidou, and N. Gehlenborg. Nexus: Interactive visual exploration of 3D spatial connectivity for tissue analysis. *OSF Preprints*, 2025. doi: 10.31219/osf.io/nypmf_v1

[44] V. Nath, K. G. Schilling, P. Parvathaneni, C. B. Hansen, A. E. Hainline, Y. Huo, J. A. Blaber, I. Lyu, V. Janve, Y. Gao, et al. Deep learning reveals untapped information for local white-matter fiber reconstruction in diffusion-weighted MRI. *Magnetic Resonance Imaging*, 62(6):220–227, 2019. doi: 10.1016/j.mri.2019.07.012

[45] E. Nichols, C. E. I. Szoeke, S. E. Vollset, N. Abbasi, F. Abd-Allah, J. Abdela, M. T. E. Aichour, R. O. Akinyemi, F. Alahdab, S. W. Asgedom, et al. Global, regional, and national burden of Alzheimer's disease and other dementias, 1990–2016: A systematic analysis for the global burden of disease study 2016. *The Lancet Neurology*, 18(1):88–106, 2019. doi: 10.1016/S1474-4422(18)30403-4

[46] T. M. Nir, N. Jahanshad, J. E. Villalon-Reina, A. W. Toga, C. R. Jack, M. W. Weiner, P. M. Thompson, and A. D. N. Initiative. Effectiveness of regional DTI measures in distinguishing Alzheimer's disease, MCI, and normal aging. *NeuroImage: Clinical*, 3:180–195, 2013. doi: 10.1016/j.nicl.2013.07.006

[47] V. Petrovic, J. Fallon, and F. Kuester. Visualizing whole-brain DTI tractography with GPU-based tuboids and LoD management. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1488–1495, 2007. doi: 10.1109/TVCG.2007.70532

[48] W. Poewe, K. Seppi, C. M. Tanner, G. M. Halliday, P. Brundin, J. Volkmann, A. Schrag, and A. E. Lang. Parkinson's disease. *Nature Reviews Disease Primers*, 3(1):17013, 2017. doi: 10.1038/nrdp.2017.13

[49] N. Reimers and I. Gurevych. Sentence-BERT: Sentence embeddings using siamese bert-networks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, pp. 3980–3990, 2019. doi: 10.18653/v1/D19-1410

[50] R. Salvador, J. Suckling, M. R. Coleman, J. D. Pickard, D. Menon, and E. D. Bullmore. Neurophysiological architecture of functional magnetic resonance images of the human brain. *Cerebral Cortex*, 15(9):1332–1342, 2005. doi: 10.1093/cercor/bhi069

[51] S. W. Scheff, D. A. Price, F. A. Schmitt, M. A. Scheff, and E. J. Mufson. Synaptic loss in the inferior temporal gyrus in mild cognitive impairment and Alzheimer's disease. *Journal of Alzheimer's Disease*, 24(3):547–557, 2011. doi: 10.3233/JAD-2011-101782

[52] V. J. Schmithorst and W. Yuan. White matter development during adolescence as shown by diffusion MRI. *Brain and Cognition*, 72(1):16–25, 2010. doi: 10.1016/j.bandc.2009.06.005

[53] A. Sherbondy, D. Akers, R. Mackenzie, R. Dougherty, and B. Wandell. Exploring connectivity of the brain's white matter with dynamic queries. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):419–430, 2005. doi: 10.1109/TVCG.2005.59

[54] M. F. Shewarega, J. Troidl, O. Alvarado Rodriguez, M. Dindoost, P. Harth, H. Haberkern, J. Stegmaier, D. Bader, and H. Pfister. MoMo – combining neuron morphology and connectivity for interactive motif analysis in connectomes. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):2025–07, 2025. doi: 10.1109/TVCG.2023.3327388

[55] L. Shi, J. Hu, Z. Tan, J. Tao, J. Ding, Y. Jin, Y. Wu, and P. M. Thompson. Mv^{2}net: Multi-variate multi-view brain network comparison over uncertain data. *IEEE Transactions on Visualization and Computer Graphics*, 28(12):4640–4657, 2022. doi: 10.1109/TVCG.2021.3098123

[56] L. Shi, H. Tong, M. Daianu, F. Tian, and P. M. Thompson. Visual analysis of brain networks using sparse regression models. *ACM Transactions on Knowledge Discovery from Data*, 12(1):1–30, 2018. doi: 10.1145/3023363

[57] L. Shi, H. Tong, and X. Mu. BrainQuest: Perception-guided brain network comparison. In *Proceedings of the IEEE International Conference on Data Mining*, pp. 379–388, 2015. doi: 10.1109/ICDM.2015.135

[58] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages*, pp. 336–343. IEEE, 1996. doi: 10.1109/VL.1996.545307

[59] K. Song, X. Tan, T. Qin, J. Lu, and T. Liu. MPNet: Masked and permuted pre-training for language understanding. In *Proceedings of the Conference on Neural Information Processing Systems*, pp. 16857–16867, 2020. doi: 10.5555/3495724.3497138

[60] S. Song, J. Yoshino, T. Q. Le, S. Lin, S. Sun, A. H. Cross, and R. C. Armstrong. Demyelination increases radial diffusivity in corpus callosum of mouse brain. *NeuroImage*, 26(1):132–140, 2005. doi: 10.1016/j.neuroimage.2005.01.028

[61] S. Srivishagan, L. Kumaralingam, K. Thanikasalam, U. Pinidiyaarachchi, and N. Ratnarajah. Discriminative patterns of white matter changes in Alzheimer's. *Psychiatry Research: Neuroimaging*, 328:111576, 2023. doi: 10.1016/j.pscychresns.2022.111576

[62] P. M. Thompson, M. S. Mega, R. P. Woods, C. I. Zoumalan, C. J. Lindshield, R. E. Blanton, J. Moussai, C. J. Holmes, J. L. Cummings, and A. W. Toga. Cortical change in Alzheimer's disease detected with a disease-specific population-based brain atlas. *Cerebral Cortex*, 11(1):1–16, 2001. doi: 10.1093/cercor/11.1.1

[63] A. W. Toga and P. M. Thompson. Mapping brain asymmetry. *Nature Reviews Neuroscience*, 4(1):37–48, 2003. doi: 10.1038/nrn1009

[64] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al. Llama2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. doi: 10.48550/arXiv.2307.09288

[65] J. Troidl, S. Warchol, J. Choi, J. Matelsky, N. Dhanyasi, X. Wang, B. Wester, D. Wei, J. W. Lichtman, H. Pfister, and J. Beyer. ViMO - visual analysis of neuronal connectivity motifs. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):748–758, 2024. doi: 10.1109/TVCG.2023.3327388

[66] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Proceedings of the International Conference on Neural Information Processing Systems*, vol. 30, pp. 5998–6008. Curran Associates, Inc., 2017. doi: 10.48550/arXiv.1706.03762

[67] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. *Proceedings of International Conference on Learning Representations*, 2018. doi: 10.17863/CAM.48429

[68] J. Wang. *Multigraph Visualization for Feature Classification of Brain Network Data*. Ph.d. dissertation, Purdue University, 2016.

[69] J. Xiang, H. Guo, R. Cao, H. Liang, and J. Chen. An abnormal resting-state functional brain network indicates progression towards Alzheimer's disease. *Neural Regeneration Research*, 8(30):2789–2799, 2013. doi: 10.3969/j.issn.1673-5374.2013.30.009

[70] C. Xu, T. Neuroth, T. Fujiwara, R. Liang, and K.-L. Ma. A predictive visual analytics system for studying neurodegenerative disease based on DTI fiber tracts. *IEEE Transactions on Visualization and Computer Graphics*, 29(4):2020–2035, 2023. doi: 10.1109/TVCG.2021.3137174

[71] X. Yang, L. Shi, M. Daianu, H. Tong, Q. Liu, and P. Thompson. Blockwise human brain network visual comparison using nodetrix representation. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):181–190, 2017. doi: 10.1109/TVCG.2016.2598472

[72] C. Ying, T. Cai, S. Luo, S. Zheng, G. Ke, D. He, Y. Shen, and T. Liu. Do transformers really perform badly for graph representation? In *Proceedings of the Conference on Neural Information Processing Systems*, pp. 28877–28888, 2021. doi: 10.48550/arXiv.2106.05234

[73] M. Yu, O. Sporns, and A. J. Saykin. The human connectome in Alzheimer disease — relationship to biomarkers and genetics. *Nature Reviews Neurology*, 17(9):545–563, 2021. doi: 10.1038/s41582-021-00529-1

[74] R. Zeng, J. Lv, H. Wang, L. Zhou, M. Barnett, F. Calamante, and C. Wang. FOD-Net: A deep learning method for fiber orientation distribution angular super resolution. *Medical Image Analysis*, 79:102431, 2022. doi: 10.1016/j.media.2022.102431

[75] S. Zhang, S. Correia, and D. H. Laidlaw. Identifying white-matter fiber bundles in DTI data using an automated proximity-based fiber-clustering method. *IEEE Transactions on Visualization and Computer Graphics*, 14(5):1044–1053, 2008. doi: 10.1109/TVCG.2008.52

Weihan Zhang is a Ph.D. student of computer science at Sun Yat-sen University. He received a Bachelor's degree in software engineering from Chongqing University in 2022. His research interests include scientific visualization and visual analytics.

Jun Tao is an associate professor of computer science at Sun Yat-sen University and National Supercomputer Center in Guangzhou. He received a Ph.D. degree in computer science from Michigan Technological University in 2015. His major research interest is scientific visualization, especially in applying information theory, deep learning, and optimization techniques to interactive flow visualization and multivariate data exploration.

APPENDIX A
FIBER BUNDLE REPRESENTATION

We consider a rich set of bundle features and transform the geometric features into latent vectors to form bundle representations using an autoencoder. Using an appropriate autoencoder structure, we can compress the joint distribution of the features into latent vectors.

Data preparation. Brain fibers are modeled as streamlines traced from the diffusion tensor field. Based on a predefined set of 70 ROIs, we segment each streamline into sub-portions that connect specific ROI pairs. For each fiber segment, we compute geometric features such as length, curvature, and torsion. These features are then aggregated according to their associated ROI pair to construct bundle-level representations. The geometric features are paired to form joint distributions, where each joint distribution is represented as a 20×20 histogram. All histograms corresponding to a bundle form a single feature map of the bundle. This allows us to analyze and represent the brain networks using bundles instead of individual fibers, whose number could be overwhelming. Using 70 brain regions, this strategy reduces the data of each subject from hundreds of thousands of fibers to several thousand bundles. To ensure reliability, we sample fiber bundles with high connectivity strength, resulting in a dataset of approximately 13,000 samples. Each sample is a joint distribution feature map for a bundle. We use this data to sufficiently train an autoencoder until the geometric features can be reconstructed from the latent vectors. The trained model is used to encode all the fiber bundles.

Model design. Our representation model is similar to typical autoencoders for images, consisting of blocks that blend linear projection (i.e., convolution and linear layers) with non-linear activations. The autoencoder transforms a feature map into a latent vector, and reconstructs the map from the vector. When the map can be accurately reconstructed, the vector is considered to contain the same information as the map. In our scenario, the autoencoder should capture the *affinity of neighboring bins* and the *non-linear behavior* in the maps.

For the bin affinity, we use 2D convolution to preserve the neighboring relationships in the feature maps. While attention mechanisms are considered as powerful convolutional encoders in computer vision tasks as well, their quadratic complexity is unnecessary for our task because the joint distributions are relatively small.

For the non-linear behavior, we use the star operation [40], which maps inputs to a high-dimensional non-linear feature space to enhance the representation power. Our experiment shows that the star operation represents the feature maps more accurately than the conventional activations (e.g., ReLU and Sine), as evidenced by Table I and Fig. 1. Please refer to the following paragraph for the detailed performance discussion.

Implementation, training, and performance. We use a four-layer autoencoder network, where each layer consists of a convolutional layer, a linear layer, and a star operation. We utilized the Adam optimizer for parameter updates. The batch size was set to 2048 feature maps. The learning rate started at 10^{-3} with 0.1 decrease scheduler, and was decayed by 10^{-5} using L2 regularization.

TABLE I: The average RMSE, PSNR, and SSIM with different activation functions in the fiber bundle representation.

Activation	RMSE ↓	PSNR ↑	SSIM ↑
ReLU	0.24	12.32	0.057
Sine	0.12	18.75	0.315
Star	**0.06**	**24.77**	**0.782**

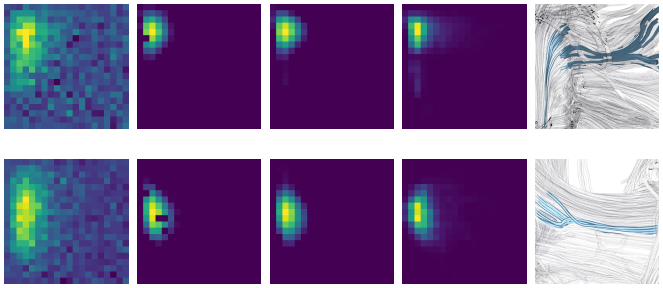

(a) Sine (b) ReLU (c) Star (d) GT (e) Fibers

Fig. 1: A qualitative comparison of reconstructed joint histograms between length and curvature, using different activation functions. (a) to (c) show the reconstructed histograms using Sine, ReLU, and star operations, respectively. (d) shows the original histograms as the ground truth, and (e) shows the corresponding fiber bundles.

We examine the performance of our autoencoder using the star operation, and compare it with the autoencoders using ReLU and Sine activation functions. As shown in Table I, our model achieves the smallest reconstruction loss (RMSE) and the highest peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM). For the other two activation functions, Sine seems to consistently outperform ReLU for the three metrics. We further study the reconstruction performance by examining their visualization results. We use the joint histograms of length and curvature. The reconstructed joint histograms of two example bundles are shown in Fig. 1. We find that Sine often creates high-frequency signals that are not prominent in the original histogram, and ReLU may miss some values. In contrast, using the star operation, our autoencoder can mostly resemble the original histograms.

APPENDIX B
BRAIN NETWORK REPRESENTATION

Data preparation. The brain network of a subject is a graph, where each node is a brain region and an edge is the fiber bundle connecting two regions. Each edge is associated with a latent vector, created from the fiber bundle representation, containing both DTI and geometric features. We construct a dataset of 198 brain networks, each corresponding to one subject. The networks are fed into a graph neural network for a classification task, where the subjects are classified as healthy, mild cognitive impairment, severe cognitive impairment, and AD. This forces the network to focus on AD-related information.

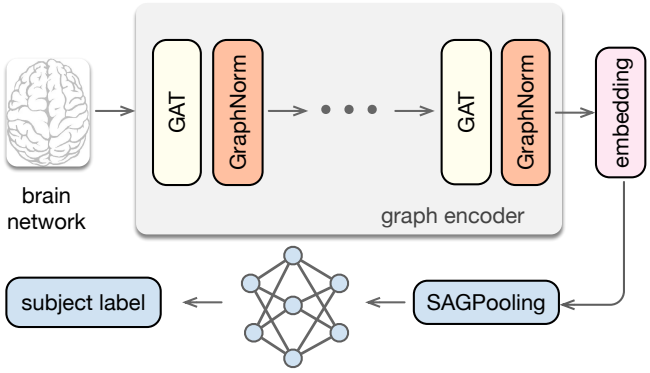

Fig. 2: Illustration of our GNN structure. It includes a multi-layer graph encoder to extract node embeddings at different scales; SAGPooling is then used to obtain graph-level representations, followed by an MLP classification head to predict the subject's diagnostic class.

Model design. Our brain network representation model follows the typical graph classification network structure, as shown in Fig. 2. It consists of three modules: the graph convolution layer for aggregating neighborhood information and updating node features, the pooling layer to generate graph-level representations, and the linear head to classify the graphs.

For aggregation of node features, we use attention-based graph convolutional layers (GAT) to pass information between nodes. To capture structures across scales, our graph encoder consists of multiple GAT layers. Following each GAT layer, we use an additional GraphNorm layer to prevent nodes with higher degrees from dominating the aggregation. GraphNorm is proven to improve training stability and accelerate model convergence in graph neural networks [10].

For learning graph-level representations, we use a self-attention-based pooling (SAGPool) to collect node features. Pooling is a process of graph coarsening, where similar nodes are clustered together. Therefore, it requires a proper score to identify the prominent nodes that remain in the coarsened graph. SAGPool leverages a network to learn the self-attention score for each node, considering both node features and graph topology information, using the following equation:

$$Z = \sigma\left(\tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}X\Theta_{att}\right), \tag{4}$$

where $\tilde{A}$ and $\tilde{D}$ are the normalized adjacency matrix and the degree matrix, respectively, and Θ_{att} is a learnable attention parameter. The nodes are ranked by scores, and the nodes with higher scores are kept after the coarsening.

For classification, we use a simple linear classification head, because the graph encoder already extracts the important information for the classification. Skip connections are used to propagate classification information.

Training and Performance During training, we split the data with 20% as the testing set. We used a cross-entropy loss because predicting the subject labels was a typical classification task. We included an L2 regularization term with a decay rate of 10^{-5} to prevent overfitting. We used the Adam optimizer for parameter updates. The batch size was set to 64 brain networks and the model was trained for 200 epochs. The learning rate started at 10^{-3} with a 0.1 decrease scheduler.

TABLE II: Performance comparison of our GNN and without SAGPool and GraphNorm. The precision, recall, and F1 scores are computed using the macro average.

method	accuracy ↑	F1 ↑	precision ↑	recall ↑
our GNN	**0.77**	**0.79**	**0.83**	**0.77**
w/o SAGPool	0.67	0.60	0.60	0.67
w/o GraphNorm	0.64	0.54	0.52	0.64

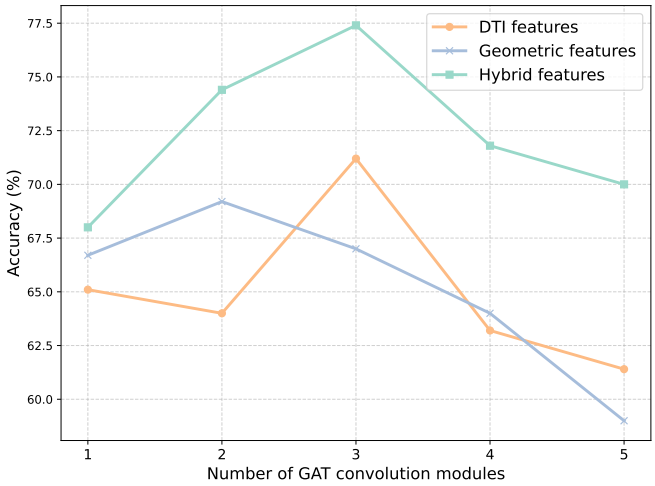

Fig. 3: Performance comparison using three types of features (DTI features, geometric features, and hybrid features), along with the performance of the graph encoder containing multiple GAT modules.

We examined our model design with an ablation study to assess the impact of individual components. The results in Table II confirmed the effectiveness of both SAGPooling and GraphNorm. With these two modules, our model outperformed the alternatives without them in all four classification metrics, including accuracy, recall, precision, and F1-score.

We then used the full model to study the impact of the feature type and network depth. For the feature, we experimented with three types of features as the input: DTI features, geometric features, and hybrid features. The hybrid features included both the DTI features and geometric features. For the depth, we assessed the performance of graph encoders from a single layer to five layers.

Fig. 3 confirmed the effectiveness of combining the DTI and geometric features, because the model trained with the hybrid features consistently outperformed the models trained with individual types of features. The models trained with individual types of features delivered similar overall performance: the models of 1, 2, and 4 layers performed better with geometric features, and the models of 3 and 5 layers performed better with DTI features.

In terms of depth, we found the models generally performed the best with a moderate depth. The model trained with the hybrid features and the DTI features delivered the best

performance at three layers, and the model trained with the geometric features reached the best performance at two layers. A possible explanation is that the network may not collect enough regional information when it is too shallow, and it may smooth out the information when it is too deep.

Based on the above results, we train a three-layer network with SAGPooling and GraphNorm modules using both DTI and geometric features, for the exploration presented in the case studies.

Appendix C
System Implementation

Our system framework consists of two components: a data processing engine and a user interface. The data engine extracts geometric features from fibers and represents fiber bundles and brain networks as latent vectors. The user interface reads and renders the fibers and associated features.

Data processing and timing performance. The data processing engine consists of a CUDA module for geometric feature extraction and two neural networks implemented using PyTorch for the representation. Our data has 43,327 fibers and 14,124,602 control points on average for each subject. The geometric feature extraction requires 10.03 seconds to process each subject's fibers, using an NVidia GeForce RTX 4090 graphics card. The fiber bundle representation network reads the geometric and DTI features and produces the bundle representations, which takes 2.91 seconds for each subject. The brain network representation reads the bundle representations and outputs the network representation of a subject in 1.18 seconds. Please refer to the repository https://github.com/sysuvis/neurolens for the implementation of the representation networks and user interface. Overall, the data engine can prepare all data for an individual subject in 14.09 seconds, which is efficient enough for clinical use. Under research settings, researchers usually need to study many subjects simultaneously, typically several hundred. Our data engine can still generate the required data in hours, which should meet the time requirement as well.

User interface. Our user interface is a standalone desktop application built with the Qt framework in C++. All visual components, including 2D charts and the 3D fiber view, are rendered using OpenGL. User interactions are handled by Qt, which may trigger renderer updates as needed. The most computationally intensive operation is refreshing the brain fiber renderer, which takes an average of 0.39 seconds, ensuring the system remains largely interactive using an NVidia RTX 4090 graphics card. However, when the user interface is run on a less powerful machine (e.g., a laptop), it may need to rely on remote rendering, transmitting either a video stream or rendered images to the interface. Given that our exploration primarily relies on the 2D charts, we do not anticipate frequent updates to the fiber rendering. Therefore, minor rendering lag should be acceptable. During runtime, the application consumes approximately 2.7 GB of memory, which includes both data management and rendering buffers for one subject.

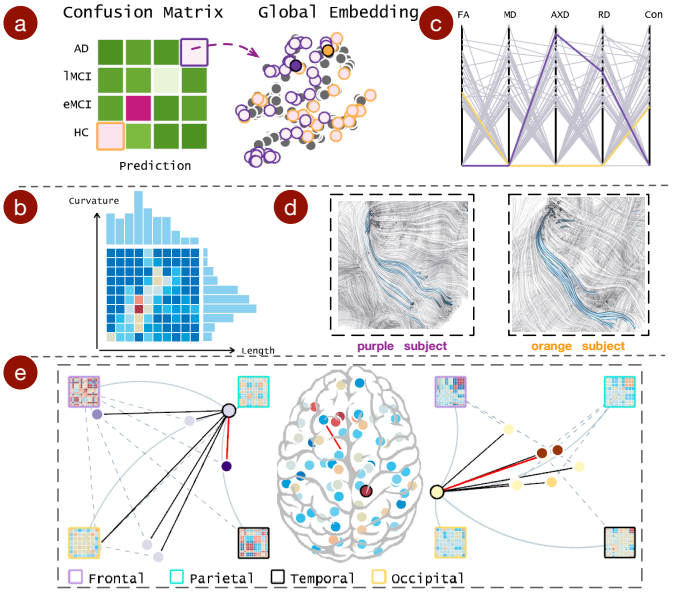

Fig. 4: Comparison between an AD subject and an HC subject. (a) displays the exploration process of the selected observation object. (b) shows the joint distribution of fiber bundle geometric features corresponding to the AD subject. (c) displays the DTI features of the AD subject (purple line) and the HC subject (orange line). (d) highlights the corresponding fiber bundle of the AD (purple) and HC (orange) subjects. (e) displays their brain region embeddings and differences in the network comparison view.

Appendix D
Additional Case Study: Comparing AD and HC Individuals

Exploration goals. Following the group-level comparison, the experts aimed to further investigate the detailed differences between two individuals with AD and HC. Toward this goal, they selected two individual nodes within the global embedding. In contrast to the previously analyzed distant groups, these nodes were located in close proximity to reveal more subtle differences. The purple node represented an AD subject, while the orange node represented an HC subject.

Examining the regional differences. $E1$ selected the darkest red node in the network comparison view, as shown in Fig. 4 (e). The position of the node indicated an ROI in the medial temporal lobe, and its color indicated a marked difference between the two subjects. By hovering over the node, the expert found the node to be the left parahippocampal gyrus, which plays a critical role in cognitive processing and memory functions. $E1$ and $E2$ explained that the hippocampal region is capable of autonomous firing, making it particularly susceptible to focal brain damage.

Furthermore, the experts observed several red nodes in the upper region of the comparison view, corresponding to the frontal lobe. This observation was further supported by comparing the matrices in the two embeddings, where the matrix of the frontal region (upper left) and the matrix of the temporal region (lower left) exhibited more pronounced

differences. These two matrices of the AD subject showed a higher number of red cells, indicating a broader distribution of nodes in the embedding. $E1$ and $E2$ noted that both the frontal and temporal lobes contain projection neurons and are densely connected with other brain regions.

Investigating the bundle difference. The experts decided to further examine the fiber bundle connecting the parahippocampal region to the frontal lobe. $E2$ further explained that this pathway is implicated in various neurodegenerative diseases, including AD and epilepsy. They selected an edge connected to the dark purple node, as indicated by the red line in Fig. 4 (e). This edge represents a fiber bundle connecting to the frontal lobe in the right hemisphere.

Their analysis began with the geometric feature distributions, as shown in Fig. 4 (b). The experts found that the fibers exhibited moderate length and curvature in both subjects, without indicating any overt abnormalities. They further verified this in the fiber rendering view, as shown in Fig. 4 (d). While the overall shape of individual fibers appeared similar across subjects, the experts noted that the fiber bundle terminals in the AD subject were more dispersed, suggesting possible neuronal degeneration.

Subsequently, they examined the DTI features of the fiber bundle, as depicted in Fig. 4 (c), where the purple line represented the AD subject and the orange line represented the HC subject. The experts observed that the AD subject exhibited lower FA but higher AxD and RD values. They explained that FA reflects the integrity of white matter [6], and a reduction in FA might indicate progressive white matter damage in AD. Furthermore, they noted that RD tends to increase due to demyelination associated with AD [60], and AxD has also been reported to increase with white matter neurodegeneration [16]. Based on these findings, the experts suggested that the observed reduction in white matter integrity may contribute to increased fiber dispersion, causing the fiber pathways to appear more divergent.

Validation from literature query module. Based on these observations, $E1$ and $E2$ attributed the abnormal fiber connectivity in the parahippocampal region of the temporal lobe to pathological changes in τ proteins. This hypothesis was validated using the literature module. They found that τ protein initially accumulates in the temporal lobe and spreads from the entorhinal cortex to the parahippocampal cortex, eventually reaching other brain regions [73].

APPENDIX E
QUANTITATIVE ANALYSIS OF FEATURE DEPENDENCIES VIA MUTUAL INFORMATION

To support the observation discussed in Section 3.1, we provide a quantitative analysis of the dependencies among geometric features using mutual information. As illustrated in the right bottom corner of Fig. 1, we analyze the pairwise relationships between geometric features. Mutual information $I(X;Y)$ measures the amount of information obtained about one random variable X given another variable Y, defined as:

$$I(X;Y) = \sum_{x \in X} \sum_{y \in Y} p(x,y) \log \left(\frac{p(x,y)}{p(x)p(y)} \right), \quad (5)$$

where x and y are observed outcomes of the variables X and Y. In Fig. 1, we find that the mutual information between features is small, compared to their own information (ranging from 5.6 to 5.8). The mutual information between the entropy and the fiber length (1.8) is the highest, as short fiber tracts are often stable exhibiting fewer variations; while all other variables share little information (< 0.68), indicating their interactions cannot be revealed by individual variables. These joint interactions suggest that analyzing their joint distribution is necessary.