STAG: Smart Tools and Applications in Graphics (2025)
M. Comino Trinidad, C. Mancinelli, D. Cabiddu, and D. Giorgi (Editors)

NeuroShape: exploiting neural architectures for shape analysis of ultrastructural 3D neuroscience morphologies

Humaira Shaffique[1] Uzair Shah[1] Mahmood Alzubaidi[1] Jens Schneider[1] Pierre Julius Magistretti[2] Corrado Calì[3] Mowafa Househ[1] Marco Agus[1]

[1]College of Science and Engineering, Hamad Bin Khalifa University, Qatar
[2]Biological, Environmental Sciences and Engineering, King Abdullah University of Science and Technology, Saudi Arabia
[3]Neuroscience Institute Cavalieri Ottolenghi, University of Turin, Italy

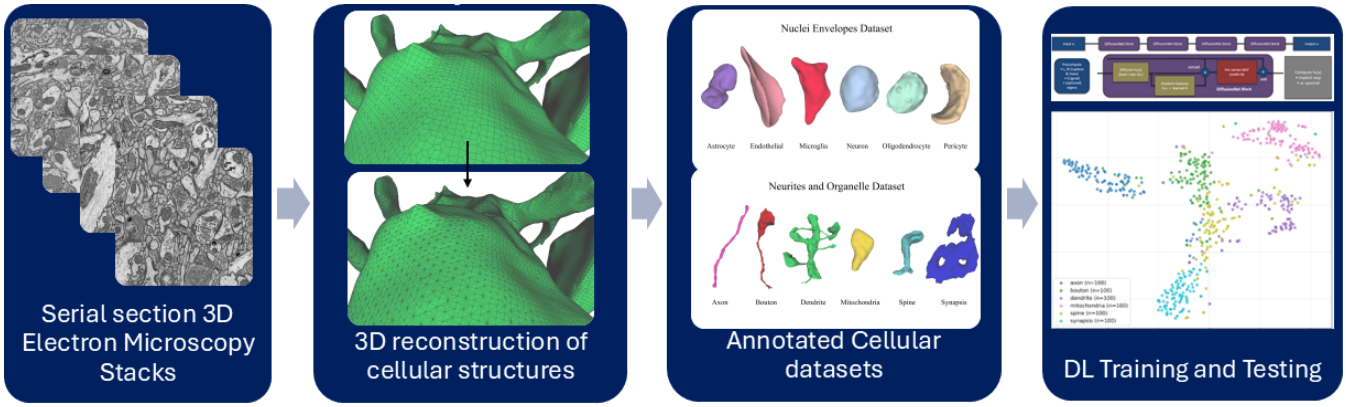

Figure 1: ***NeuroShape****. From 3D Serial Section EM stacks we derived two different 3D morphology dataset (nuclei envelopes and neurites with organelles), and analyzed them with geometric deep networks (DiffusionNet [SACO22], Laplacian2Mesh [DWL*24]). The figure highlights our pipeline; datasets are publicly released to support future studies.*

Abstract

Recent advances in volume electron microscopy (EM) enable nanometric-scale 3D reconstructions of neural tissue, providing unprecedented opportunities for studying cellular and subcellular morphology in neuroscience. The geometry of structures such as nuclei, neurites, and organelles can encode phenotypic information relevant to both functional specialization and pathological conditions, and thus represents a valuable complement to connectivity-based approaches in connectomics. While previous studies relied on handcrafted descriptors and classical machine learning for morphology analysis, recent progress in deep learning for 3D shape understanding offers new opportunities to learn robust, task-specific representations directly from geometric data. In this work we present NeuroShape*, a first exploration of modern deep learning methods for shape analysis of ultrastructural 3D neuroscience morphologies. We introduce two annotated datasets derived from EM reconstructions: one of nuclei envelopes, and one of neurites and neural organelles. We benchmark two state-of-the-art neural architectures for 3D geometry (DiffusionNet [SACO22] and Laplacian2Mesh [DWL*24]) and compare them against traditional feature-based descriptors previously used in neural morphology analysis. Our preliminary results highlight both the feasibility and the challenges of applying deep learning shape analysis techniques in this domain, and we release the datasets as a reference resource for future studies.*

CCS Concepts

• ***Computing methodologies*** → ***Shape analysis;*** *Image segmentation; Point-based models; Mesh models;* ***Supervised learning by classification; Neural networks;*** • ***Applied computing*** → *Computational biology;*

DOI: 10.2312/stag.20251333

1. Introduction

Recent advances in volume electron microscopy (EM) have enabled the acquisition of ultrastructural data at nanometric resolution, yielding three-dimensional (3D) reconstructions of neural tissue with unprecedented detail. Such datasets provide a foundation for connectomics research, where the reconstruction of cellular morphologies and their substructures is essential for mapping synaptic circuits and deciphering brain organization at the micro- and nanoscale [CAK*19,BTB*22]. The resulting volumetric stacks allow for systematic exploration of fine cellular components, including nuclei, neurites, mitochondria, and synaptic organelles, and thus represent an invaluable resource for both basic neuroscience and clinical investigations [CCB*25].

A central challenge lies in the analysis of the 3D shapes of these structures. The morphology of nuclei, neurites, and other organelles encodes phenotypic signatures that can reflect both functional specializations and pathological alterations [TMB*22]. For example, nuclei envelope deformations have been linked to neurodegenerative disorders [WXC*25], while neurite arborization patterns contribute to functional connectivity and plasticity and can differentiate between different species [KSA*25]. Beyond pathology, the spatial arrangement and geometry of subcellular components provide clues about mechanisms of spatial organization, compartmentalization, and cellular interactions in the brain. Shape analysis in this context supports quantitative characterizations that complement connectivity-focused approaches, offering a richer understanding of cellular and tissue-level function.

The availability of 3D reconstructions has stimulated a variety of computational approaches for geometry understanding and shape analysis. Early methods relied on handcrafted shape descriptors, such as curvature-based signatures, spherical harmonics, or skeleton-based representations, often combined with classical machine learning classifiers [LNJJ*24]. Different modalities have been exploited, including voxel grids, point clouds, meshes, and graph-based abstractions, each offering a distinct trade-off between geometric fidelity and computational cost. While these methods provided the first quantitative insights into neural morphologies, they often suffered from sensitivity to noise, limited generalization across datasets, and dependence on carefully engineered features.

In recent years, the proliferation of deep learning technologies has transformed 3D shape analysis across computer graphics, computer vision, and biomedical imaging [BBL*17,GAC*23]. Neural architectures such as point-based networks, graph neural networks, and mesh-based operators have demonstrated the ability to learn task-specific geometric representations directly from data, potentially surpassing the limitations of traditional descriptors. In particular, models that leverage intrinsic geometric operators—such as diffusion layers or Laplacian embeddings—provide robust ways to encode shape structure while respecting surface geometry [FLJ*25]. Despite their success in other domains, the application of such neural architectures to ultrastructural 3D reconstructions in neuroscience remains largely unexplored.

In this work, we take a first step toward filling this gap. We propose *NeuroShape*, a study of the feasibility of modern deep learning methods for 3D geometry applied to ultrastructural data derived from nanometric-scale EM reconstructions. Specifically, we focus on two classification problems: (i) brain cell recognition from the 3D shapes of nuclei envelopes; and (ii) classification of partial reconstructions of neural organelles and neurites. To support this study, we introduce two annotated datasets and evaluate two representative deep learning baselines—DiffusionNet [SACO22], which employs diffusion layers on discretized surfaces, and Laplacian2Mesh [DWL*24], a data-driven method for Laplacian-based mesh understanding—comparing them against traditional feature-based descriptors previously explored for neural morphology analysis.

Our main contributions are as follows:

- Two annotated datasets: one comprising nuclei envelopes, and one comprising neural substructures such as neurites and organelles.
- A preliminary assessment of two modern deep learning–based shape analysis baselines (DiffusionNet [SACO22] and Laplacian2Mesh [DWL*24]) for neuroscience morphologies.

To the best of our knowledge, this study represents the first application of deep learning 3D geometry methods to ultrastructural EM-based neural reconstructions. We believe that the proposed datasets and baseline evaluations will provide a valuable benchmark for practitioners and foster future research at the intersection of connectomics, geometry processing, and neural shape understanding.

2. Related Work

Our work deals with shape analysis for neuroscience, and deep learning in geometry. For space limitations, we do not aim to provide here an extensive review of the related literature: we refer interested readers to the most recent surveys in these topics, namely morphometric methods [LNJJ*24], 3D shape analysis [MKKS22], and visual analytics and visualization [BTB*22] targeting neuroscience, connectomics data, and ultrastructural analysis [CCB*25], or the most recent deep learning solutions for 3D point clouds [VKA*23], or 3D segmentation [HYL*25], or for task-oriented graph analysis [CWD*24]. In the following paragraphs we discuss the methods most closely related to our work.

Shape analysis in ultrastructural neuroscience data With the recent evolution of volume scanning and automatized serial scanning imaging technologies, various investigation efforts have been dedicated to 3D analysis of neural cells and morphologies at nanometric scale level [CCB*25]. Previous studies regarded quantitative analysis [CAK*19] and immersive exploration analysis [SAB*25] for understanding neuroenergetics mechanisms [CCK*18], spatial distributions [TCG*22], and layouts [BGC*25, TWC*24], with the ultimate goal of reverse engineering brain organization and functioning [MMR*15], and support simulation and modeling efforts [CCK*18]. To this end, various 3D representations have been considered and exploited along the last decade, ranging from segmented volume data [ACAA*19, PD24], up to triangular mesh representations [ACG*23, TCG*22], and medial axis representations (skeletons) [MUM*21, KDS*18, BAS*20, SZHQ25]. Recently, some efforts have been oriented to leveraging machine learning towards classification and recognition of 3D ultrastructural reconstructions. To this end, Calì, Agus et al. considered

3D reconstructions of nuclei envelopes extracted from cortical layers [CAK*19], and developed various shape descriptors to be incorporated in support vector machines for classification, by considering implicit surface representations [AVG*19], spherical harmonics decompositions [AGP*20], and curvature descriptors [ATAG*21]. In this work, we extend the nuclei classification framework to the analysis of neurites reconstructions, and we preliminary assess modern deep-learning based frameworks in comparison to traditional machine learning approaches. To the best of our knowledge, deep learning frameworks have not been considered yet for shape analysis of 3D nanometric scale neural reconstructions.

Deep learning for shape analysis With the explosion of deep learning technologies, the computational geometry field started to experience a transformative process [BBL*17], with various attempts of generalizing data-driven techniques to non-Euclidean domains, like geometric data without regular structure, that is challenging to process using traditional deep learning operations such as convolution and pooling. The simplest and traditional way to deal with 3D data is through multi-view projection, that obtained by projecting 3D shapes into 2D images, which can be processed using standard image-based architectures using convolutional neural networks [HGG21, JBC*19] or transformers [GXG*23]. On the other hands, point clouds are one of the most studied representations in geometric deep learning, for their diffusion and because they can be directly obtained by scanning devices. Along the last decade, a huge number of deep learning frameworks have been proposed for various tasks in point clouds [CSKG17], ranging from classification and retrieval [ZX19], up to completion [YRW*21], filtering [ZLQH21], generation [LH21, YHH*19, SWL*20], and spectral analysis [PZL*24]. For what concerns 3D mesh analysis, the most direct method consists on converting 3D mesh in voxels and extending 2D convolutions to 3D [WL20]: however, these methods are limited to low resolution grids and can be applied to moderately simple shapes. Another strategy consists of extracting traditional shape features like Heat Kernel Signature (HKS) [SOG09] or Wave Kernel Signature [ASC11], and exploiting them to learn intrinsic geometric structures, through specialized CNNs involving convolution and pooling operations to handle the irregularities of meshes [DWL*24, LSG*25, SS21]. In this category, DiffusionNet [SACO22] is a general-purpose approach to deep learning on 3D surfaces, that is agnostic with respect to 3D representation (point clouds or triangular meshes), while Laplacian2Mesh [DWL*24] is a flexible convolutional neural network (CNN) that can deal with irregular triangle meshes by mapping the input mesh surface to the multi-dimensional Laplacian-Beltrami space. In this work, we assess both DiffusionNet [SACO22] and Laplacian2Mesh [DWL*24] capabilities in shape classification of neuroscience nanometric scale morphologies. To the best of our knowledge, this is the first attempt of adapting mesh-based 3D surface learning methodologies to the neuroscience domain.

3. Background

In the following, we provide contextual information about ultrastructural analysis in connectomics, as well as the main computational methods for performing shape analysis, ranging from traditional shape descriptors to modern deep learning frameworks.

3.1. 3D connectomics ultrastructure analysis

Connectomics aims to reconstruct the wiring diagram of the brain by mapping neurons, glia, and their synaptic connections at nanometer resolution. Recent advances in volume electron microscopy (EM) provide dense image stacks that enable the three-dimensional (3D) reconstruction of cellular and subcellular structures, including nuclei, dendrites, axons, boutons, spines, synapses, and organelles such as mitochondria [AVG*19, BTB*22]. These reconstructions form the basis for high-resolution connectomics, where geometry is not only a matter of visualization but also a carrier of biological meaning.

The importance of 3D reconstruction lies in its ability to capture morphological detail that is invisible or ambiguous in 2D slices. For example, nuclei envelopes can deviate strongly from simple spherical models, and their shape variability has been linked to both functional specialization and pathological conditions [DM21]. Similarly, the branching patterns of neurites and the spatial organization of organelles affect neuronal communication, energy metabolism, and plasticity [KSA*25, CAK*19]. Quantitative analysis of these structures provides essential clues for understanding the cellular basis of brain computation and disease mechanisms.

Shape analysis in this context is therefore critical. By transforming raw reconstructions into measurable geometric objects—such as meshes, skeletons, or graphs—researchers can classify cell types, compare phenotypes, and detect structural signatures of pathology. Traditional approaches relied on handcrafted shape descriptors combined with machine learning, but recent advances in geometric deep learning offer the promise of learning such features directly from the data. In the following subsections, we summarize these two lines of work.

3.2. Shape descriptors

Traditional shape analysis methods for neural morphologies rely on handcrafted geometric descriptors that capture salient properties of 3D structures. These descriptors have been widely used in computer graphics and geometry processing, and they have also been adapted to ultrastructural neuroscience tasks. Below we describe three representative families: spherical harmonics, Laplace–Beltrami spectra, and diffusion-based heat kernel signatures.

Spherical harmonics. Spherical harmonics (SH) provide a natural basis for representing functions defined on the unit sphere. By mapping a 3D surface onto a spherical domain, its geometry can be decomposed into frequency components invariant to rotation and scaling [KFR03]. The spherical harmonic expansion of a surface function $f(\theta,\phi)$ is

$$f(\theta,\phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_{lm} Y_l^m(\theta,\phi),$$

where $Y_l^m(\theta,\phi)$ are the spherical harmonic basis functions and c_{lm} are the coefficients. Low-order terms capture the global shape,

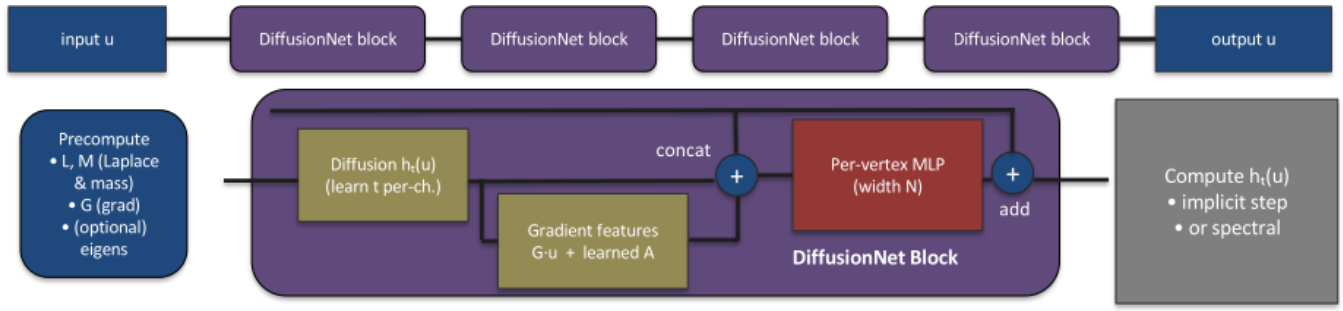

Figure 2: ***DiffusionNet:*** *The architecture learns diffusion processes followed by multilayer perceptrons (MLPs) for feature transformation, effectively combining intrinsic geometric features such as coordinates, normals, or spectral signatures with data-driven learning. Simplified scheme adapted from [SACO22].*

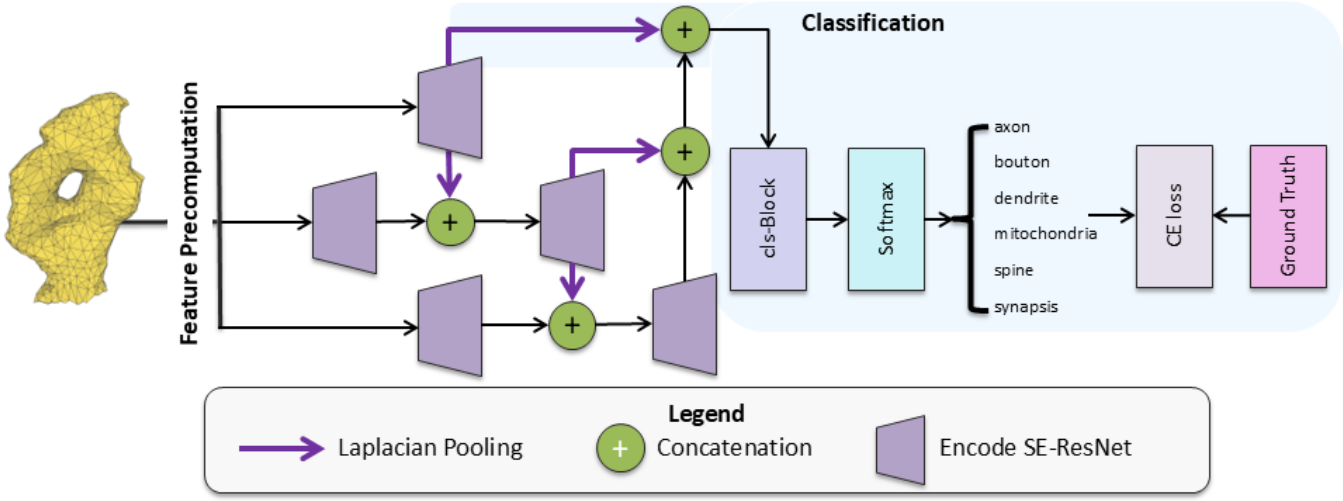

Figure 3: ***Laplacian2Mesh:*** *Given an input 3D mesh as the input, the geometric features are projected into the spectral domain at three different resolutions; then a U-Net architecture [AAR*24] uses Squeeze and Excitation ResNet blocks [HSS18] with small-sized convolution kernels to fuse the nearby-frequency features, and Laplacian pooling/unpooling fuse the spectral features of different resolutions. Simplified scheme adapted from [DWL*24].*

while higher-order terms encode local detail. Such descriptors have been used for tasks such as classifying neural nuclei shapes and characterizing envelope variability [AGP*20, ATAG*21].

Laplace–Beltrami decomposition. The Laplace–Beltrami operator (LBO) generalizes the Laplacian to curved manifolds, providing an intrinsic description of shape geometry. For a function f defined on a surface,

$$\Delta f(x) = -\mathrm{div}(\nabla f(x)).$$

Solving the eigenvalue problem

$$\Delta \phi_i(x) = \lambda_i \phi_i(x),$$

yields eigenpairs (λ_i, ϕ_i) that form a spectral basis of the surface [RWP06]. The sequence of eigenvalues, also known as the *Shape DNA*, is invariant to isometric deformations and encodes fundamental shape properties. In connectomics, this approach can be used to compare morphologies of reconstructed organelles or neurites across samples.

Heat Kernel Signatures. The Heat Kernel Signature (HKS) is a diffusion-based descriptor that characterizes the way heat propagates across a surface over time [SOG09]. The heat kernel is defined as

$$h_t(x,y) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(x)\phi_i(y),$$

where λ_i and ϕ_i are the eigenvalues and eigenfunctions of the LBO. The diagonal of the kernel provides the HKS at point x:

$$\mathrm{HKS}(x,t) = \sum_{i=0}^{\infty} e^{-\lambda_i t}\, \phi_i(x)^2.$$

By evaluating the HKS across a range of diffusion times t, one obtains a multi-scale descriptor that captures both local and global geometry. In the context of ultrastructural reconstructions, HKS can support classification of subcellular structures by encoding their intrinsic geometric variability in a robust, scale-aware manner.

3.3. Deep learning for shape analysis

Recent advances in geometric deep learning have transformed 3D shape analysis in computer graphics, computer vision, and biomedical imaging [BBL*17]. Unlike traditional pipelines based on handcrafted shape descriptors, neural networks can learn task-specific representations directly from data while still exploiting geometric features such as those described in Section 3.2. In practice, these frameworks often incorporate spectral decompositions, Laplace–Beltrami eigenfunctions, or heat kernel descriptors as input features, but enhance them through learned nonlinear operators that generalize convolution or diffusion to non-Euclidean domains.

DiffusionNet. DiffusionNet [SACO22] is a discretization-agnostic deep learning framework that operates directly on 3D surfaces represented as meshes or point clouds. It leverages the diffusion operator derived from the Laplace–Beltrami spectrum to propagate information across the surface. The core layer applies learned diffusion processes followed by multilayer perceptrons (MLPs) for feature transformation, effectively combining intrinsic geometric features such as coordinates, normals, or spectral signatures with data-driven learning. Fig. 2 depicts the architecture, that is composed by multiple DiffusionNet blocks. A key advantage of DiffusionNet is its invariance to mesh resolution and triangulation quality, making it particularly suitable for biological datasets where reconstructions may vary significantly in sampling and resolution.

Laplacian2Mesh. Laplacian2Mesh [DWL*24] is a convolutional neural network designed specifically for irregular triangle meshes. It maps the input surface into the Laplacian–Beltrami eigenspace, using the first k eigenfunctions as a spectral basis for learning. These eigenfunctions, together with geometric descriptors such as curvature or dihedral angles, form a feature vector that is processed by hierarchical neural blocks. The network learns to extract multi-scale geometric patterns by combining information from different eigenspaces, enabling robust classification of complex shapes. The architecture is depicted in Fig. 3: given an input 3D mesh as the input, the geometric features are projected into the spectral domain at three different resolutions; then a U-Net architecture [AAR*24] uses Squeeze and Excitation ResNet blocks [HSS18] with small-sized convolution kernels to fuse the nearby-frequency features, and Laplacian pooling/unpooling fuse the spectral features of different resolutions. Laplacian2Mesh is particularly relevant for ultrastructural neuroscience data, as it can capture both global morphology (e.g., nuclei envelope geometry) and local details (e.g., neurite surface folding) within the same framework. In the following sections, we assess the applicability of DiffusionNet [SACO22] and Laplacian2Mesh [DWL*24] to ultrastructural EM reconstructions, benchmarking their performance against traditional descriptor-based approaches for the classification of nuclei envelopes and neural substructures.

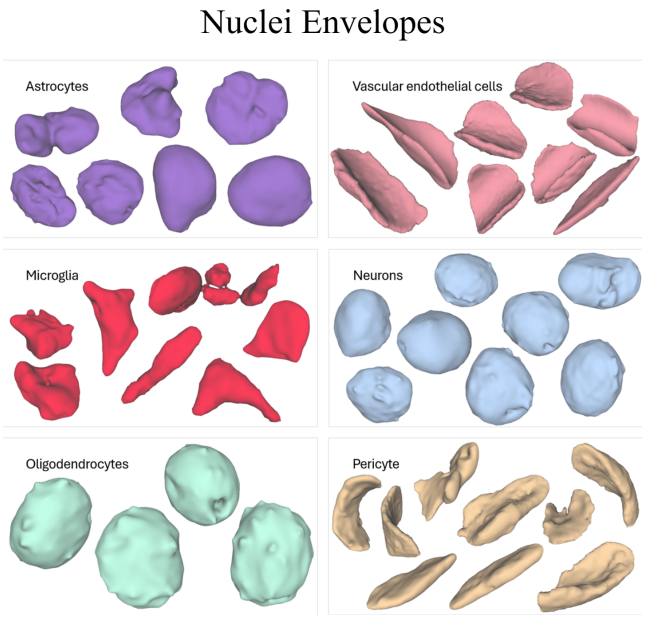

Figure 4: ***Nuclei Envelopes dataset.*** *Class exemplars shown in different colors. Types of nuclei: astrocyte (purple), endothelial (light coral), microglia (scarlet), neuron (baby blue), oligodendrocyte (mint Green) and pericyte (peach). All models were rendered with MeshLab [CCC*08].*

4. Methodology

In this section, we describe the creation of two annotated datasets, and outline the subsequent shape analysis carried out with neural architectures.

4.1. Dataset creation

A central contribution of this work is the creation of two curated datasets derived from ultrastructural reconstructions. The full datasets can be downloaded from the following web resource: NeuroShape. These datasets correspond to (i) nuclei envelopes from cortical cells and (ii) neurites and organelles from neuropil volumes. Both were constructed through data extraction, cleaning, and annotation from high-resolution electron microscopy (EM) studies. Both datasets were constructed to preserve the morphological fidelity of the original reconstructions while ensuring consistency for computational geometry workflows. Together, they provide a benchmark for exploring the feasibility of modern deep learning methods in ultrastructural neuroscience.

Table 1: ***Nuclei Envelopes dataset****. Class names with the associated split for training and testing*

Class	Train	Test	Total
Astrocyte	17	6	23
Endothelial	16	3	19
Microglia	13	6	19
Neuron	64	32	96
Oligodendrocyte	2	2	4
Pericyte	8	4	12
Total	120	53	173

Nuclei envelopes. The first dataset was derived from Serial Block-Face Electron Microscopy (SBEM) reconstructions of juvenile rat somatosensory cortex at postnatal day 14, originally published by Calì et al. [CAK*19]. The original study imaged a 750,000 μm^3 cortical volume and reconstructed 186 nuclei spanning neurons, astrocytes, microglia, pericytes, endothelial cells, and oligodendrocytes. From these reconstructions, we extracted and cleaned the nuclei envelope meshes, correcting segmentation artifacts using MeshLab. Six annotated classes were established: astrocyte, endothelial, microglia, neuron, oligodendrocyte, and pericyte. The final dataset consisted of 173 meshes, with a 70:30 train-test split per class, except for oligodendrocytes (4 samples total) which were divided 50:50 due to sparsity (Table 1). Representative examples of the six nuclei envelope classes are shown in Figure 4.

Table 2: ***Neurites & Organelles dataset****. Class names with the associated split for training and testing.*

Class	Train	Test	Total
Axon	70	30	100
Bouton	70	30	100
Dendrites	70	30	100
Mitochondria	70	30	100
Spine	70	30	100
PSD	70	30	100
Total	420	180	600

Neurites and organelles. The second dataset was obtained from dense FIB-SEM reconstructions of neuropile in layer I of the somatosensory cortex, acquired from six mice (three adults, three aged), as reported in the aging study of Calì et al. [CWB*18]. From these volumes, we extracted meshes representing six annotated classes: axon, bouton, dendrite, mitochondria, spine, and Postsynaptic Density (PSD). Cleaning steps involved mesh simplification, artifact removal, and class balancing. During mesh simplification using quadric edge collapse decimation [GH97], we observed varying degrees of geometric degeneration illustrated in Figure 6 across different face counts (4000→2000→1000→500→250), highlighting its effect on smoothness and biological fidelity. The curated dataset contains 600 meshes evenly distributed across classes, with a 70:30 train-test split (Table 2). Representative examples are illustrated in Figure 5.

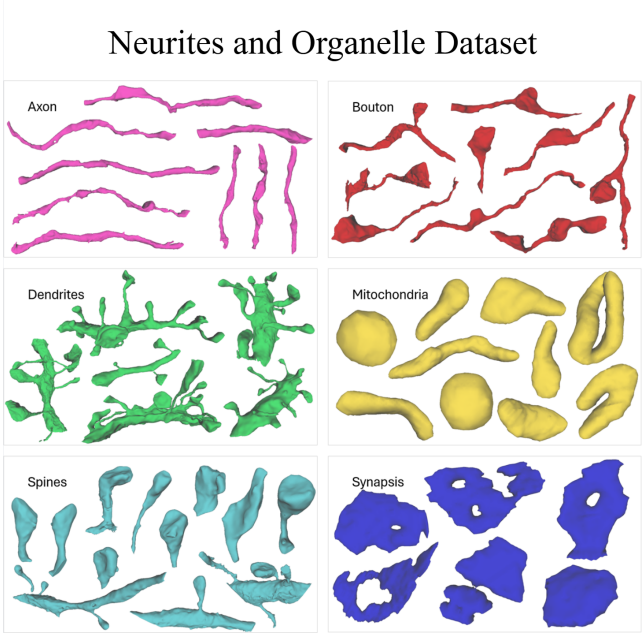

Figure 5: ***Neurites and Organelle dataset****. Class exemplars shown in different colors. Neurites: axon (pink), bouton (red), dendrite (green) and spine (turquoise). Neurites junction: Postsynaptic density (PSD; dark blue). Organelle: mitochondria (yellow). All models were rendered with MeshLab [CCC*08].*

4.2. Experimental setup

We evaluated two state-of-the-art neural architectures for shape analysis: DiffusionNet [SACO22] and Laplacian2Mesh [DWL*24]. Both were adapted to our biological datasets through careful hyperparameter tuning. In the following, we summarize the training setup and dataset-specific adjustments. All experiments were performed on a workstation equipped with NVIDIA Geforce RTX 4090 with 24GB RAM. For fair comparison of the models, all meshes were reduced to 4000 faces.

DiffusionNet For both datasets, DiffusionNet was configured with four diffusion blocks and trained for 200 epochs. The heat kernel signature was set to 16, matching the original model design. Model width was chosen equal to k_{eig} to fully exploit the spectral basis, and a smoothing rate of 0.2 was applied to improve training stability. Dropout and gradient-based features (`dropout_state`, `grad_feat`, `grad_rot`) were disabled, as their benefits are limited for smaller, noisy biological datasets. For the *nuclei envelope dataset*, parameter search showed that lower spectral bases ($k_{\text{eig}} = 32, 64$) degraded performance, while $k_{\text{eig}} = 128$ achieved peak accuracy. A learning rate of 3×10^{-3} with decay factor 0.2 every 30 epochs provided stable convergence, while smaller rates (e.g., 1×10^{-5}) led to slow learning and early plateau. For the *neu-*

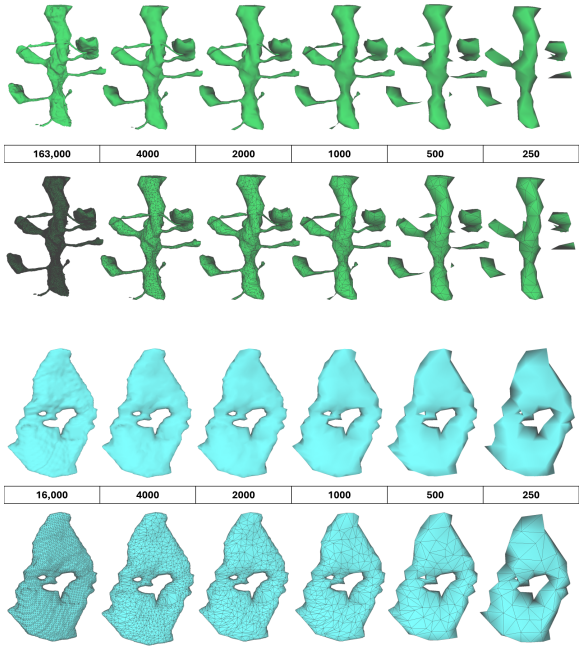

Figure 6: ***Mesh degeneration examples.*** *Top: Dendrite mesh degeneration loses connection to spines. Bottom: Synapse mesh wihholds its structure but decimation removes surface details and smoothness.*

rites and organelles dataset, hyperparameters were again tuned by trial and error. Initial training with learning rate 1×10^{-3} caused unstable accuracy. Reducing the rate to 5×10^{-4} stabilized convergence but slowed progress. A final choice of 9×10^{-4} struck a balance between stability and accuracy. Increasing k_{eig} to 128 improved accuracy at moderate cost, while 256 eigenvalues were computationally expensive without significant gains. To evaluate the impact of feature representation, DiffusionNet was trained separately using raw XYZ coordinates and Heat Kernel Signature (HKS) features. The resulting performance differences are reported in Section 5..

Laplacian2Mesh Each mesh was projected onto three spectral layers (32, 16, and 8 eigenfunctions), and the batch size was set to 32. For the *nuclei envelope dataset*, Laplacian2Mesh was trained with cosine learning-rate restarts. Initial settings ($\text{lr} = 5 \times 10^{-2}$, $\eta_{\min} = 5 \times 10^{-9}$) proved unstable, so the range was reduced to 10^{-2}–10^{-4}, yielding improved stability. Weight decay and loss weights were reduced to 10^{-3} and 0.9, respectively. To further stabilize accuracy, the number of epochs was increased to 225, with restart intervals set to 15 epochs. For the *neurites and organelles dataset*, the larger sample size facilitated rapid learning, so training was limited to 70 epochs to mitigate overfitting. Learning rates in the range 5×10^{-5} to 10^{-6} yielded stable convergence. Weight decay and loss weights were kept as in the nuclei dataset, but restart intervals were shortened to 10 epochs to match the smaller training duration.

Table 3: ***Nuclei Envelopes datasets****. Accuracy Comparison of Diffusion Net and Lap2mesh for different classification tasks.*

No.	Dataset	DiffusionNet	Lap2mesh
1	**FULL**	81.13%	75.50%
2	**AENOP**	87.23%	76.60%
3	**AEMNP**	83.31%	76.50%
4	**ANOP**	95.45%	77.08%
5	**AENP**	95.56%	84.40%
6	**AENM**	89.36%	87.20%
7	**AENO**	95.34%	86.00%
8	**EMOP**	60.00%	46.70%
9	**AEN**	100.00%	87.80%
10	**ANO**	95.00%	82.50%
11	**ANM**	97.73%	84.10%
12	**EOP**	55.56%	55.60%
13	**AN**	100.00%	92.10%
14	**EM**	88.89%	100.00%
15	**EO**	100.00%	80.00%

Abbreviations: **FULL** = Astrocyte, Endothelial, Microglia, Neuron, Oligodendrocyte & Pericyte; ALL REMAINING DATASETS LABELED WITH THE FIRST LETTER OF ALL CLASSES PRESENT: **A** = Astrocyte; **E** = Endothelial; **M** = Microglia; **N** = Neuron; **O** = Oligodendrocyte; **P** = Pericyte.

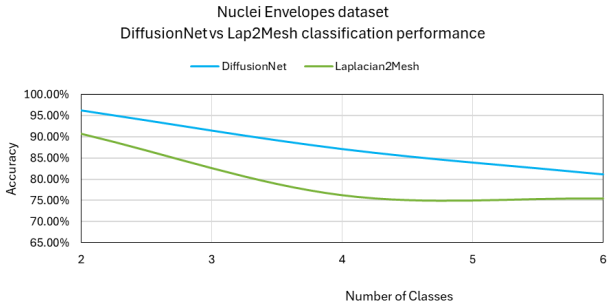

Figure 7: ***Accuracy for Nuclei Envelopes dataset****. Comparison between DiffusionNet & Lap2mesh with respect to the complexity of classification task.*

5. Results

We report experimental results on the two datasets introduced in Section 4.1, evaluating classification accuracy and analyzing the learned feature embeddings. Results are organized by dataset.

Nuclei envelopes Table 3 and Figure 7 summarize classification performance across different class subsets. When considering the full six-class problem (astrocytes, endothelial cells, microglia, neurons, oligodendrocytes, pericytes), both networks struggled due to strong class imbalance and visual similarity among rare categories, with oligodendrocytes and microglia being most frequently misclassified. Removing these challenging classes increased accuracy substantially, reaching above 90% for subsets such as AENP and ANOP. The best results were obtained on balanced subsets

Table 4: ***Neurites and Organelles datasets***. *Accuracy Comparison of Diffusion Net and Lap2mesh for different classification tasks.*

No.	Dataset	DiffusionNet	Lap2mesh
1	**BMSySp**	96.60%	81.50%
2	**ADMSy**	100.00%	94.10%
3	**ADwithSpine**	92.22%	87.70%
4	**ADwithBouton**	92.22%	85.30%
5	**ADonly**	100.00%	98.30%
6	**ALL_SIX_CLASSES**	91.11%	80.90%
7	**Five_Classes_No_Spine**	93.21%	87.80%
8	**Five_Classes_No_Bouton**	94.66%	85.70%

Abbreviations: BMSySp = Bouton, Mitochondria, PSD & Spine; **ADMSy** = Axon, Dendrite, Mitochondria & PSD; **ADwithSpine** = Axon & Dendrite with Spine; **ADwithBouton** = Axon & Dendrite with Bouton; **ADonly** = Axon & Dendrite only; **ALL_SIX_CLASSES** = Axon, Bouton, Dendrite, Mitochondria, Spine & PSD; **Five_Classes_No_Spine** = Axon, Dendrite, Bouton, Mitochondria & PSD; **Five_Classes_No_Bouton** = Axon, Dendrite, Mitochondria, Spine & PSD.

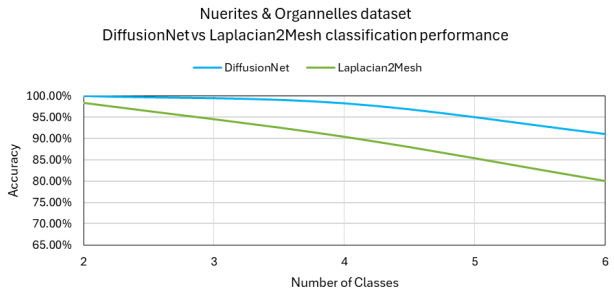

Figure 8: ***Accuracy for Neurites and Organelles dataset***. *Comparison between DiffusionNet & Lap2mesh with respect to the complexity of classification task.*

(e.g., AN and AEN), where DiffusionNet achieved perfect accuracy (100% in AEN and AENP), while Laplacian2Mesh plateaued at lower values (around 87.8%). As comparison with traditional shallow machine learning frameworks with Spherical Harmonics shape descriptors, InShade [ATAG*21] and Wish [AGP*20] could reach accuracies around 80% for the case number 3, but the composition of the dataset was different as well as the testing conditions (for the limited sample availability, only cross-validation accuracies were computed). Figure 7 shows the accuracy trend as a function of class subsets. DiffusionNet exhibits a smooth degradation as class complexity increases, whereas Laplacian2Mesh stabilizes early. This difference reflects the design of the models: DiffusionNet learns robust global representations, while Laplacian2Mesh is more sensitive to class imbalance. Overall, these results highlight the importance of balanced datasets for ultrastructural classification and confirm that DiffusionNet is better suited for discriminating nuclei morphologies.

Neurites and organelles Results for the neurites dataset are reported in Table 4 and Figure 8. Both methods achieved solid performance across the six-class problem (82.1% for DiffusionNet, 80.9% for Laplacian2Mesh). Subsets of higher-level structures (ADMSy, ADonly) reached very high accuracies, with DiffusionNet exceeding 99%. Interestingly, classification of spines reduced performance due to their geometric similarity to dendrites and boutons, while boutons were easier to separate due to their rounded morphology. This effect is visible in subsets such as Five_Classes_No_Spine, where accuracy increased by nearly 7%. Figure 8 shows accuracy trends across subsets. Both methods drop in performance as class complexity increases, but DiffusionNet consistently surpasses Laplacian2Mesh. The differences are especially pronounced when discriminating fine-grained categories (spines vs. dendrites). A slight drop in accuracy was observed when using lower-resolution meshes (2000–500 faces), confirming that excessive simplification and geometric degeneration negatively affect classification performance (Fig. 10).

Impact of HKS vs. XYZ on performance In addition to comparing neural architectures, we evaluated the effect of input feature representation on DiffusionNet by training it separately using raw XYZ vertex coordinates and the Heat Kernel Signature (HKS) on both datasets. As shown in Figure 11, HKS consistently outperformed XYZ in the 6-class and 5-class settings. In the Nuclei Envelopes dataset, the difference between XYZ and HKS was relatively small, likely due to the limited dataset size. In contrast, the Neurites and Organelles dataset showed a much clearer advantage for HKS, with XYZ inputs yielding noticeably lower accuracy. This demonstrates that spectral diffusion-based descriptors capture global surface morphology more effectively than purely Euclidean coordinates, particularly when fine-grained geometric differences are critical for class discrimination.

Analysis of projected embeddings To further interpret the learned representations, we projected the features from both models using PCA [AW10], t-SNE [MH08], and UMAP [SMG21](Fig. 9). DiffusionNet embeddings formed more compact and well-separated clusters, particularly in UMAP space, reflecting the network's ability to capture global shape differences. Laplacian2Mesh embeddings, while structured, exhibited stronger overlaps between morphologically similar classes such as spines and dendrites, consistent with its local feature bias. These observations align with the quantitative results: DiffusionNet, by leveraging diffusion-based propagation across the entire surface, provides more discriminative embeddings for classification. In contrast, Laplacian2Mesh, originally designed for mesh segmentation tasks, emphasizes local geometry, which may be advantageous for fine-grained surface analysis but less effective for whole-shape classification.

Summary Across both datasets, DiffusionNet demonstrated higher and more consistent classification accuracy compared to Laplacian2Mesh. Nuclei envelope experiments highlight the sensitivity of both models to dataset imbalance, while neurite experiments underline the difficulty of separating fine-grained, morphologically similar structures. Embedding analyses confirm that DiffusionNet yields more separable latent spaces, supporting its superior performance. These findings suggest that diffusion-based neu-

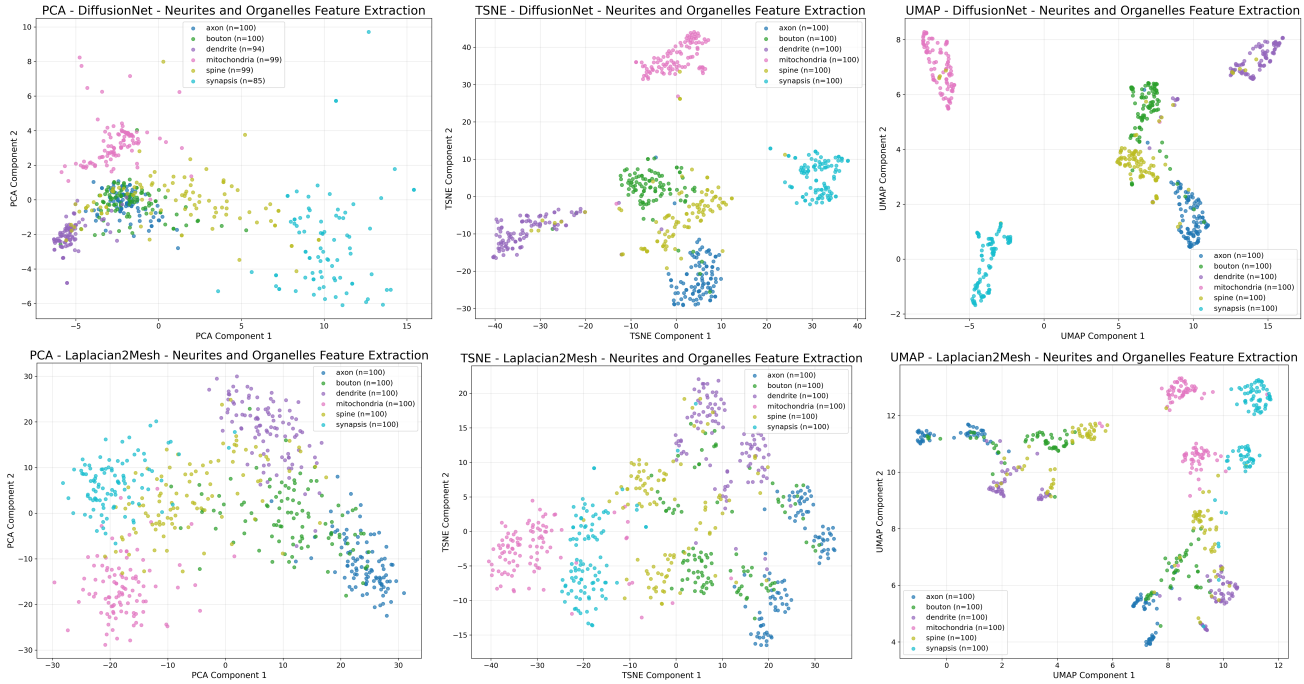

Figure 9: ***Embedding comparison.*** *We compared the projected feature embeddings of DiffusionNet [SACO22] (top) and Laplacian2Mesh [DWL*24] (bottom) through various dimension reduction method on Neurites and Organelles dataset. From left to right, the projected features obtained with PCA, tSNE, and UMAP. The clusters obtained with DiffusionNet (on top) exhibit higher compactness and degree of separation, consistently with the reported accuracies.*

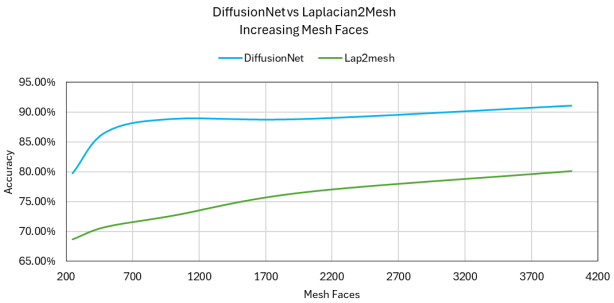

Figure 10: ***Effect of mesh resolution on classification performance.*** *Accuracy of DiffusionNet and Laplacian2Mesh on the Neurites and Organelles dataset as the number of mesh faces increases.*

ral architectures are better suited for classification of ultrastructural morphologies, whereas Laplacian-based models may be more effective for tasks emphasizing local shape segmentation.

6. Conclusions

In this paper we introduced *NeuroShape*, the first systematic study of deep learning methods for 3D shape analysis of ultrastructural morphologies reconstructed from nanometric-scale electron microscopy. Our contributions are threefold: (i) the creation of two curated datasets for ultrastructural analysis, one focusing on nuclei envelopes across six cortical cell types and another comprising neurites and organelles; (ii) the benchmarking of two state-of-the-art neural architectures, DiffusionNet [SACO22] and Laplacian2Mesh [DWL*24]; and (iii) a preliminary comparative evaluation showing that diffusion-based approaches provide more discriminative embeddings and higher classification accuracy, particularly for complex ultrastructural categories. Our findings highlight the potential of deep learning–based shape analysis in neuroscience, while also pointing to current limitations. Some annotations in the datasets may be noisy due to variability in the expertise of domain scientists and the intrinsic difficulty of differentiating fine-grained neurite structures. Moreover, the datasets represent partial reconstructions, which constrain the scope of shape analysis. As future work, we plan to extend shape analysis to segmentation tasks on full 3D reconstructions of entire cells, including neurons and glia, in order to capture complete morphological signatures. In particular, regarding glia, more effort is required, considering their high fractality and shape heterogeneity [VMA*23]. Another research direction is the integration of learned shape features into visual analytics frameworks that support the exploration of dense reconstructions and the spatial organization of cells. Recent advances such as the Mixture Graph [ATAS21] and Volume Puzzle [AAAT*22] demonstrate how efficient data structures and multivariate visualization can enable scalable analysis of large segmented volumes, and we envision combining such frameworks

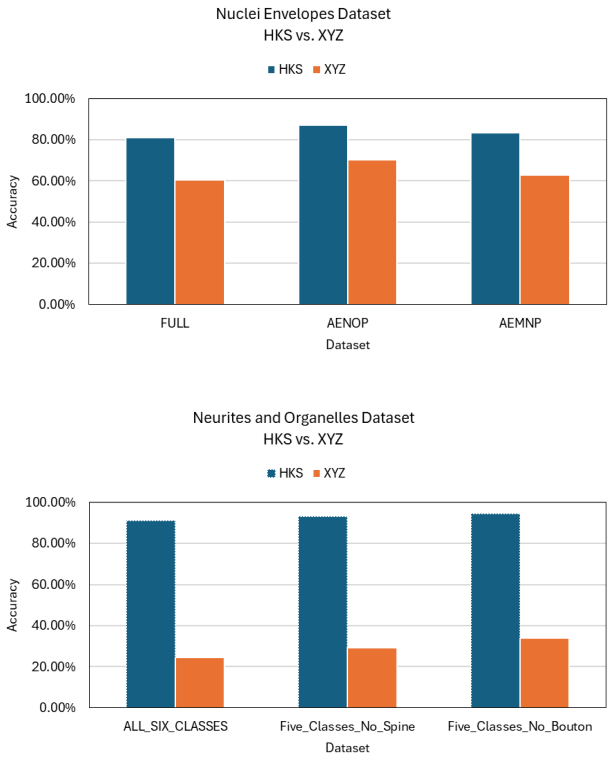

Figure 11: ***HKS and XYZ comparison.*** *Top: Nuclei dataset with closer XYZ and HKS accuracies. Bottom: Neurites and organelles dataset with substantial advantage for HKS.*

with learned shape descriptors to provide interactive, task-driven exploration of connectomics data.

Acknowledgments

This publication was funded by the PPM-7th Cycle grant (PPM 07-0409-240041, AMAL-For-Qatar) from the Qatar National Research Fund, a member of the Qatar Foundation. The findings herein reflect the work and are solely the responsibility, of the authors. We thank the PICO imaging facility at the NICO institute, Roberta Parolisi for the imaging support, and Christel Genoud at UNIL for the help in sample preparation. Open Access funding provided by the Qatar National Library.

References

[AAAT*22] AGUS M., ABOULHASSAN A., AL THELAYA K., PINTORE G., GOBBETTI E., CALÌ C., SCHNEIDER J.: Volume puzzle: visual analysis of segmented volume data with multivariate attributes. In *2022 IEEE Visualization and Visual Analytics (VIS)* (2022), pp. 130–134. doi:10.1109/VIS54862.2022.00035. 9

[AAR*24] AZAD R., AGHDAM E. K., RAULAND A., JIA Y., AVVAL A. H., BOZORGPOUR A., KARIMIJAFARBIGLOO S., COHEN J. P., ADELI E., MERHOF D.: Medical image segmentation review: The success of u-net. *IEEE Transactions on Pattern Analysis and Machine Intelligence 46*, 12 (2024), 10076–10095. doi:10.1109/TPAMI.2024.3435571. 4, 5

[ACAA*19] AGUS M., CALÌ C., AL-AWAMI A., GOBBETTI E., MAGISTRETTI P., HADWIGER M.: Interactive volumetric visual analysis of glycogen-derived energy absorption in nanometric brain structures. *Computer Graphics Forum 38*, 3 (2019), 427–439. doi:10.1111/cgf.13700. 2

[ACG*23] ABDELLAH M., CANTERO J. J. G., GUERRERO N. R., FONI A., COGGAN J. S., CALÌ C., AGUS M., ZISIS E., KELLER D., HADWIGER M., ET AL.: Ultraliser: a framework for creating multiscale, high-fidelity and geometrically realistic 3d models for in silico neuroscience. *Briefings in bioinformatics 24*, 1 (2023). doi:10.1093/bib/bbac491. 2

[AGP*20] AGUS M., GOBBETTI E., PINTORE G., CALÌ C., SCHNEIDER J.: Wish: efficient 3d biological shape classification through willmore flow and spherical harmonics decomposition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* (2020), pp. 4184–4194. doi:10.1109/CVPRW50498.2020.00494. 3, 4, 8

[ASC11] AUBRY M., SCHLICKEWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)* (2011), pp. 1626–1633. doi:10.1109/ICCVW.2011.6130444. 3

[ATAG*21] AL-THELAYA K., AGUS M., GILAL N. U., YANG Y., PINTORE G., GOBBETTI E., CALÍ C., MAGISTRETTI P. J., MIFSUD W., SCHNEIDER J.: Inshade: Invariant shape descriptors for visual 2d and 3d cellular and nuclear shape analysis and classification. *Computers Graphics 98* (2021), 105–125. doi:10.1016/j.cag.2021.04.037. 3, 4, 8

[ATAS21] AL-THELAYA K., AGUS M., SCHNEIDER J.: The mixture graph-a data structure for compressing, rendering, and querying segmentation histograms. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (2021), 645–655. doi:10.1109/TVCG.2020.3030451. 9

[AVG*19] AGUS M., VELOZ CASTILLO M., GARNICA MOLINA J. F., GOBBETTI E., LEHVÄSLAIHO H., MORALES TAPIA A., MAGISTRETTI P. J., HADWIGER M., CALÍ C.: Shape analysis of 3d nanoscale reconstructions of brain cell nuclear envelopes by implicit and explicit parametric representations. *Computers Graphics: X 1* (2019), 100004. doi:10.1016/j.cagx.2019.100004. 3

[AW10] ABDI H., WILLIAMS L. J.: Principal component analysis. *Wiley interdisciplinary reviews: computational statistics 2*, 4 (2010), 433–459. doi:10.1002/wics.101. 8

[BAS*20] BOGES D., AGUS M., SICAT R., MAGISTRETTI P. J., HADWIGER M., CALÌ C.: Virtual reality framework for editing and exploring medial axis representations of nanometric scale neural structures. *Computers Graphics 91* (2020), 12–24. doi:10.1016/j.cag.2020.05.024. 2

[BBL*17] BRONSTEIN M. M., BRUNA J., LECUN Y., SZLAM A., VANDERGHEYNST P.: Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine 34*, 4 (2017), 18–42. doi:10.1109/MSP.2017.2693418. 2, 3, 5

[BGC*25] BARTESAGHI L., GIANGRECO B., CHIAPPINI V., VELOZ CASTILLO M. F., MONACO M., MÉDARD J.-J., GAMBAROTTA G., AGUS M., CALÌ C.: Tight spaces, tighter signals: Spatial constraints as drivers of peripheral myelination. *Cells 14*, 12 (2025). doi:10.3390/cells14120926. 2

[BTB*22] BEYER J., TROIDL J., BOORBOOR S., HADWIGER M., KAUFMAN A., PFISTER H.: A survey of visualization and analysis in high-resolution connectomics. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 573–607. doi:10.1111/cgf.14574. 2, 3

[CAK*19] CALÌ C., AGUS M., KARE K., BOGES D. J., LEHVÄSLAIHO H., HADWIGER M., MAGISTRETTI P. J.: 3d cellular reconstruction of cortical glia and parenchymal morphometric analysis from

serial block-face electron microscopy of juvenile rat. *Progress in Neurobiology 183* (2019), 101696. doi:https://doi.org/10.1016/j.pneurobio.2019.101696. 2, 3, 6

[CCB*25] CHIAPPINI V., CASTILLO M. F. V., BIANCARDI F., DI CUNTO F., MAGISTRETTI P. J., VERCELLI A., AGUS M., CALÌ C.: Ultrastructure of astrocytes using volume electron microscopy: A scoping review. *The Journal of Physiology* (2025). doi:10.1113/JP287455. 2

[CCC*08] CIGNONI P., CALLIERI M., CORSINI M., DELLEPIANE M., GANOVELLI F., RANZUGLIA G.: MeshLab: an Open-Source Mesh Processing Tool. In *Eurographics Italian Chapter Conference* (2008), Scarano V., Chiara R. D., Erra U., (Eds.), The Eurographics Association. doi:10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136. 5, 6

[CCK*18] COGGAN J. S., CALÌ C., KELLER D., AGUS M., BOGES D., ABDELLAH M., KARE K., LEHVÄSLAIHO H., EILEMANN S., JOLIVET R. B., ET AL.: A process for digitizing and simulating biologically realistic oligocellular networks demonstrated for the neuro-glio-vascular ensemble. *Frontiers in neuroscience 12* (2018), 664. doi:10.3389/fnins.2018.00664. 2

[CSKG17] CHARLES R. Q., SU H., KAICHUN M., GUIBAS L. J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017), pp. 77–85. doi:10.1109/CVPR.2017.16. 3

[CWB*18] CALÌ C., WAWRZYNIAK M., BECKER C., MACO B., CANTONI M., JORSTAD A., NIGRO B., GRILLO F., DE PAOLA V., FUA P., KNOTT G. W.: The effects of aging on neuropil structure in mouse somatosensory cortex—a 3d electron microscopy analysis of layer 1. *PLOS ONE 13*, 7 (07 2018), 1–21. URL: https://doi.org/10.1371/journal.pone.0198131, doi:10.1371/journal.pone.0198131. 6

[CWD*24] CHEN C., WU Y., DAI Q., ZHOU H.-Y., XU M., YANG S., HAN X., YU Y.: A survey on graph neural networks and graph transformers in computer vision: A task-oriented perspective. *IEEE Transactions on Pattern Analysis and Machine Intelligence 46*, 12 (2024), 10297–10318. doi:10.1109/TPAMI.2024.3445463. 2

[DM21] DEOLAL P., MISHRA K.: Regulation of diverse nuclear shapes: pathways working independently, together. *Communicative & integrative biology 14*, 1 (2021), 158–175. doi:10.1080/19420889.2021.1939942. 3

[DWL*24] DONG Q., WANG Z., LI M., GAO J., CHEN S., SHU Z., XIN S., TU C., WANG W.: Laplacian2mesh: Laplacian-based mesh understanding. *IEEE Transactions on Visualization and Computer Graphics 30*, 7 (2024), 4349–4361. doi:10.1109/TVCG.2023.3259044. 1, 2, 3, 4, 5, 6, 9

[FLJ*25] FEI Y., LIU Y., JIA C., LI Z., WEI X., CHEN M.: A survey of geometric optimization for deep learning: From euclidean space to riemannian manifold. *ACM Comput. Surv. 57*, 5 (Jan. 2025). doi:10.1145/3708498. 2

[GAC*23] GERKEN J. E., ARONSSON J., CARLSSON O., LINANDER H., OHLSSON F., PETERSSON C., PERSSON D.: Geometric deep learning and equivariant neural networks. *Artificial Intelligence Review 56*, 12 (2023), 14605–14662. doi:10.1007/s10462-023-10502-7. 2

[GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), ACM Press/Addison-Wesley Publishing Co., pp. 209–216. doi:10.1145/258734.258849. 6

[GXG*23] GOYAL A., XU J., GUO Y., BLUKIS V., CHAO Y.-W., FOX D.: Rvt: Robotic view transformer for 3d object manipulation. In *Proceedings of The 7th Conference on Robot Learning* (06–09 Nov 2023), Tan J., Toussaint M., Darvish K., (Eds.), vol. 229 of *Proceedings of Machine Learning Research*, PMLR, pp. 694–710. doi:10.48550/arXiv.2306.14896. 3

[HGG21] HAMDI A., GIANCOLA S., GHANEM B.: Mvtn: Multi-view transformation network for 3d shape recognition. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)* (2021), pp. 1–11. doi:10.1109/ICCV48922.2021.00007. 3

[HSS18] HU J., SHEN L., SUN G.: Squeeze-and-excitation networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018), pp. 7132–7141. doi:10.1109/CVPR.2018.00745. 4, 5

[HYL*25] HE Y., YU H., LIU X., YANG Z., SUN W., ANWAR S., MIAN A.: Deep learning based 3d segmentation in computer vision: A survey. *Information Fusion 115* (2025), 102722. doi:10.1016/j.inffus.2024.102722. 2

[JBC*19] JIANG J., BAO D., CHEN Z., ZHAO X., GAO Y.: Mlvcnn: Multi-loop-view convolutional neural network for 3d shape retrieval. *Proceedings of the AAAI Conference on Artificial Intelligence 33*, 01 (Jul. 2019), 8513–8520. doi:10.1609/aaai.v33i01.33018513. 3

[KDS*18] KANARI L., DŁOTKO P., SCOLAMIERO M., LEVI R., SHILLCOCK J., HESS K., MARKRAM H.: A topological representation of branching neuronal morphologies. *Neuroinformatics 16*, 1 (2018), 3–13. doi:10.1007/s12021-017-9341-1. 2

[KFR03] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Rotation invariant spherical harmonic representation of 3d shape descriptors. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing* (Goslar, DEU, 2003), SGP '03, Eurographics Association, p. 156–164. 3

[KSA*25] KANARI L., SHI Y., ARNAUDON A., BARROS-ZULAICA N., BENAVIDES-PICCIONE R., COGGAN J. S., DEFELIPE J., HESS K., MANSVELDER H. D., MERTENS E. J., MEYSTRE J., DE CAMPOS PERIN R., PEZZOLI M., DANIEL R. T., STOOP R., SEGEV I., MARKRAM H., DE KOCK C. P.: Of mice and men: Dendritic architecture differentiates human from mouse neuronal networks. *iScience 28*, 7 (2025), 112928. doi:10.1016/j.isci.2025.112928. 2, 3

[LH21] LUO S., HU W.: Diffusion probabilistic models for 3d point cloud generation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021), pp. 2836–2844. doi:10.1109/CVPR46437.2021.00286. 3

[LNJJ*24] LEITE J., NHOATTO F., JACOB JR A., SANTANA R., LOBATO F.: Computational tools for neuronal morphometric analysis: a systematic search and review. *Neuroinformatics 22*, 3 (2024), 353–377. doi:10.1007/s12021-024-09674-6. 2

[LSG*25] LI T., SHI Y., GAO J., WANG J., YIN B.: Hkmcnn: Heat kernel mesh-based convolutional neural networks. *Knowledge-Based Systems 317* (2025), 113375. doi:10.1016/j.knosys.2025.113375. 3

[MH08] MAATEN L. V. D., HINTON G.: Visualizing data using t-sne. *Journal of machine learning research 9*, Nov (2008), 2579–2605. URL: https://dl.acm.org/doi/10.5555/3058878.3058894. 8

[MKKS22] MIRBAUER M., KRABEC M., KRIVANEK J., SIKUDOVÁ E.: Survey and evaluation of neural 3d shape classification approaches. *IEEE Transactions on Pattern Analysis and Machine Intelligence 44*, 11 (2022), 8635–8656. doi:10.1109/TPAMI.2021.3102676. 2

[MMR*15] MARKRAM H., MULLER E., RAMASWAMY S., REIMANN M. W., ABDELLAH M., SANCHEZ C. A., AILAMAKI A., ALONSO-NANCLARES L., ANTILLE N., ARSEVER S., ET AL.: Reconstruction and simulation of neocortical microcircuitry. *Cell 163*, 2 (2015), 456–492. doi:10.1016/j.cell.2015.09.029. 2

[MUM*21] MCDONALD T., USHER W., MORRICAL N., GYULASSY A., PETRUZZA S., FEDERER F., ANGELUCCI A., PASCUCCI V.: Improving the usability of virtual reality neuron tracing with topological elements. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (2021), 744–754. doi:10.1109/TVCG.2020.3030363. 2

[PD24] PIOCHOWIAK M., DACHSBACHER C.: Fast compressed segmentation volumes for scientific visualization. *IEEE Transactions on Visualization and Computer Graphics 30*, 1 (2024), 12–22. doi:10.1109/TVCG.2023.3326573. 2

[PZL*24] PANG B., ZHENG Z., LI Y., WANG G., WANG P.-S.: Neural laplacian operator for 3d point clouds. *ACM Trans. Graph. 43*, 6 (Nov. 2024). doi:10.1145/3687901. 3

[RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace–beltrami spectra as 'shape-dna' of surfaces and solids. *Computer-Aided Design 38*, 4 (2006), 342–366. Symposium on Solid and Physical Modeling 2005. doi:10.1016/j.cad.2005.10.011. 4

[SAB*25] SHAH U., AGUS M., BOGES D., ALDOUS H., CHIAPPINI V., ALZUBAIDI M., HADWIGER M., MAGISTRETTI P. J., HOUSEH M., CALÍ C.: Ai-guided immersive exploration of brain ultrastructure for collaborative analysis and education. *Computers Graphics 129* (2025), 104239. doi:https://doi.org/10.1016/j.cag.2025.104239. 2

[SACO22] SHARP N., ATTAIKI S., CRANE K., OVSJANIKOV M.: Diffusionnet: Discretization agnostic learning on surfaces. *ACM Trans. Graph. 41*, 3 (Mar. 2022). doi:10.1145/3507905. 1, 2, 3, 4, 5, 6, 9

[SMG21] SAINBURG T., MCINNES L., GENTNER T. Q.: Parametric umap embeddings for representation and semisupervised learning. *Neural Computation 33*, 11 (10 2021), 2881–2907. doi:10.1162/neco_a_01434. 8

[SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum* (2009), vol. 28, Wiley Online Library, pp. 1383–1392. doi:10.1111/j.1467-8659.2009.01515.x. 3, 4

[SS21] SMIRNOV D., SOLOMON J.: Hodgenet: learning spectral geometry on triangle meshes. *ACM Trans. Graph. 40*, 4 (July 2021). doi:10.1145/3450626.3459797. 3

[SWL*20] SUN Y., WANG Y., LIU Z., SIEGEL J. E., SARMA S. E.: Pointgrow: Autoregressively learned point cloud generation with self-attention. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)* (2020), pp. 61–70. doi:10.1109/WACV45572.2020.9093430. 3

[SZHQ25] SHENG P., ZHAO G., HAN T., QU L.: Self-supervised neuron morphology representation with graph transformer. *IEEE Transactions on Medical Imaging* (2025), 1–1. doi:10.1109/TMI.2025.3590484. 2

[TCG*22] TROIDL J., CALI C., GRÖLLER E., PFISTER H., HADWIGER M., BEYER J.: Barrio: Customizable spatial neighborhood analysis and comparison for nanoscale brain structures. *Computer Graphics Forum* (2022). doi:10.1111/cgf.14532. 2

[TMB*22] TURNER N. L., MACRINA T., BAE J. A., YANG R., WILSON A. M., SCHNEIDER-MIZELL C., LEE K., LU R., WU J., BODOR A. L., ET AL.: Reconstruction of neocortex: Organelles, compartments, cells, circuits, and activity. *Cell 185*, 6 (2022), 1082–1100. doi:10.1016/j.cell.2022.01.023. 2

[TWC*24] TROIDL J., WARCHOL S., CHOI J., MATELSKY J., DHANYASI N., WANG X., WESTER B., WEI D., LICHTMAN J. W., PFISTER H., BEYER J.: Vimo - visual analysis of neuronal connectivity motifs. *IEEE Transactions on Visualization and Computer Graphics 30*, 1 (2024), 748–758. doi:10.1109/TVCG.2023.3327388. 2

[VKA*23] VINODKUMAR P. K., KARABULUT D., AVOTS E., OZCINAR C., ANBARJAFARI G.: A survey on deep learning based segmentation, detection and classification for 3d point clouds. *Entropy 25*, 4 (2023). doi:10.3390/e25040635. 2

[VMA*23] VIANA J. F., MACHADO J. L., ABREU D. S., VEIGA A., BARSANTI S., TAVARES G., MARTINS M., SARDINHA V. M., GUERRA-GOMES S., DOMINGOS C., PAULETTI A., WAHIS J., LIU C., CALÌ C., HENNEBERGER C., HOLT M. G., OLIVEIRA J. F.: Astrocyte structural heterogeneity in the mouse hippocampus. *Glia 71*, 7 (2023), 1667–1682. doi:https://doi.org/10.1002/glia.24362. 9

[WL20] WANG Z., LU F.: Voxsegnet: Volumetric cnns for semantic part segmentation of 3d shapes. *IEEE Transactions on Visualization and Computer Graphics 26*, 9 (2020), 2919–2930. doi:10.1109/TVCG.2019.2896310. 3

[WXC*25] WU T., XU H., CHENG L., WU R., GUO F., CHEN X.: The nuclear envelope and nuclear pore complexes in neurodegenerative diseases. *Frontiers in Cell and Developmental Biology 13* (2025), 1550859. doi:10.3389/fcell.2025.1550859. 2

[YHH*19] YANG G., HUANG X., HAO Z., LIU M.-Y., BELONGIE S., HARIHARAN B.: Pointflow: 3d point cloud generation with continuous normalizing flows. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)* (2019), pp. 4540–4549. doi:10.1109/ICCV.2019.00464. 3

[YRW*21] YU X., RAO Y., WANG Z., LIU Z., LU J., ZHOU J.: Pointr: Diverse point cloud completion with geometry-aware transformers. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)* (2021), pp. 12478–12487. doi:10.1109/ICCV48922.2021.01227. 3

[ZLQH21] ZHANG D., LU X., QIN H., HE Y.: Pointfilter: Point cloud filtering via encoder-decoder modeling. *IEEE Transactions on Visualization and Computer Graphics 27*, 3 (2021), 2015–2027. doi:10.1109/TVCG.2020.3027069. 3

[ZX19] ZHANG W., XIAO C.: Pcan: 3d attention map learning using contextual information for point cloud based retrieval. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019), pp. 12428–12437. doi:10.1109/CVPR.2019.01272. 3