

Data-Driven Morphological Feature Perception of Single Neuron With Graph Neural Network

Tianfang Zhu, Gang Yao¹, Dongli Hu, Chuangchuang Xie¹, Pengcheng Li¹, Xiaoquan Yang, Hui Gong, Qingming Luo², and Anan Li¹, *Member, IEEE*

Abstract—Clarifying the morphological characteristics of neurons can promote the understanding of brain function. However, traditional morphometrics fail to capture the modeling of each point in reconstructed neurons, leading to limited ability to distinguish massive nerve fibers and restricted application scenarios. To address these challenges, we propose MorphoGNN, a single neuron morphological embedding based on a graph neural network in this study. MorphoGNN learns the point-level structure information of reconstructed nerve fibers by considering their nearest neighbors on each hidden layer. This enables MorphoGNN to capture the lower-dimensional representation of a single neuron through an end-to-end model. In order to meet the requirements of various tasks, both supervised and self-supervised training strategies are designed to learn the characteristics that fit artificial semantics or the morphological patterns of neurons, respectively. We quantitatively compare our embeddings with other features in neuron classification and retrieval tasks and demonstrate cutting-edge performance. Additionally, we introduce our embeddings to the task of reconstruction quality classification and neuron clustering, where they can help detect reconstruction errors and obtain similar subtyping results to existing work. Furthermore, our method can be handily combined with other modal features, such as microscopic image features and traditional morphometrics. Ablation and robustness tests are also conducted to analyze the impact of several network components and low-quality reconstructed neurons on the performance of our method. The code is available at <https://github.com/fun0515/MorphoGNN>.

Index Terms—Neuron morphology, graph neural network, point cloud, neuron reconstruction.

Manuscript received 22 February 2023; revised 12 April 2023 and 2 May 2023; accepted 5 May 2023. Date of publication 11 May 2023; date of current version 2 October 2023. This work was supported in part by the STI2030-Major Projects under Grant 2021ZD0201002 and in part by the National Natural Science Foundation of China under Grant T2122015. (Corresponding author: Anan Li.)

Tianfang Zhu, Gang Yao, Dongli Hu, and Chuangchuang Xie are with the Britton Wuhan National Laboratory for Optoelectronics, the MoE Key Laboratory for Biomedical Photonics, and the Chance Center for Biomedical Photonics, Huazhong University of Science and Technology, Wuhan 430074, China (e-mail: funfunfun@hust.edu.cn; yaogang@hust.edu.cn; hudongli@hust.edu.cn; ccxie@hust.edu.cn).

Pengcheng Li, Xiaoquan Yang, Hui Gong, and Anan Li are with the Wuhan National Laboratory for Optoelectronics, the MoE Key Laboratory for Biomedical Photonics, and the Britton Chance Center for Biomedical Photonics, Huazhong University of Science and Technology, Wuhan 430074, China, and also with the HUST-Suzhou Institute for Brainmatics, JITRI, Suzhou 215123, China (e-mail: pengchengli@mail.hust.edu.cn; xqyang@hust.edu.cn; huigong@mail.hust.edu.cn; aali@hust.edu.cn).

Qingming Luo is with the Key Laboratory of Biomedical Engineering of Hainan Province, School of Biomedical Engineering, Hainan University, Haikou 570228, China (e-mail: qluo@hainanu.edu.cn).

Digital Object Identifier 10.1109/TMI.2023.3275209

I. INTRODUCTION

THE rapid advances in brain-wide optical microscopy [1], [2], [3] have enabled researchers to obtain huge amounts of brain slice images, and a single mouse brain can generate several terabytes of data [4]. Digital tracing methods [5], [6] can then be used to reconstruct the whole structure of a single neuron from brain-wide three-dimensional images. As neuronal reconstructions are non-Euclidean data, they are typically represented in the form of SWC files [7]. These files contain a set of three-dimensional points connected in the brain space, which describe the tree-like morphological structure of the neurons. Although no two neurons have exactly the same morphology, neurons with similar structural characteristics often possess similar functions. Therefore, extracting the morphological and structural features of neurons plays a crucial role in the study of brain function [8], [9], [10].

Dozens of morphometrics [11], [12] have been proposed to extract neuronal morphological features, including the total neuron fiber length, number of sections, branch order, radial distance, and density maps. Classifiers can use these measurements to categorize neurons, but they are challenged by massive neuron data. Each measurement only describes a certain local morphological characteristic of neurons. In order to completely describe the structure of neurons, as many parameters as possible are often comprehensively used. Recently, researchers have focused on the more complex morphological features beyond these metrics. For example, the topological morphology descriptor [13] combines the topology of neuronal fibers with their spatial location, enabling it to distinguish neuron trees and randomly generated tree structures. However, designing morphological characteristics based on the experience of neuroscientists is not only cumbersome, but also incomplete in the use of morphological information. This deficiency leads to the limitation of identification ability. Thus, learnable methods have been proposed to capture neuron features at an advanced level. Cellular morphology neural networks [14] can convert reconstructed neurons into lower-dimensional vector representations. Neurons are described by a series of two-dimensional images from multiple viewpoints, and the representing features are extracted from the images by convolutional neural networks. This method uses image features to indirectly represent neuron morphology, which may introduce additional interference information from the visualization stage. MorphVAE [15] samples walks from neuron structures and learns the hidden representation at the branch-level through a variational seq2seq-autoencoder. This learning method directly processes the topological structures of neurons, but lacks the use of structural information between

points. Approaches that learn neuron morphology at the point-level remain to be studied. Moreover, there is also a lack of quantitative comparison of multiple morphological features under the same framework.

We have noticed that the reconstructed neuron data presents a similar structure to the point cloud, and several point cloud-based processing and analysis methods [16], [17] have been proposed in recent years. Multi-view convolutional neural networks [18] have demonstrated that a collection of two-dimensional views can accurately describe three-dimensional objects. Graph neural networks (GNN) [19], [20] capture the dependency relationship in the graph by information transfer between nodes, and they have also been introduced into the field of point cloud processing. Wang et al. [21] dynamically updates the graph structure between layers to learn the lower-dimensional features of the point set and applied it to classification and segmentation tasks. These point cloud methods provide learning techniques for three-dimensional spatial points, which lay the foundation for extracting point-level morphological features of neurons, thus fully exploiting the neural structure information.

In this study, we propose MorphoGNN to learn the morphologic representation of a single neuron. The main contributions are as follows:

- To our best knowledge, we propose a comprehensive framework that employs the GNN model to directly learn the morphologic information of a single neuron for the first time, covering the stages of pretreatment, analysis, utilization and evaluation. Through a data-driven approach, we capture the point-level geometric characters of neurons, leading to improved performance on a variety of morphological tasks.
- We also design both supervised and self-supervised training strategies to obtain the morphological features that suitable for multiple tasks. For diverse neural data, we incorporate dense connection, double-pooling operator, and a joint loss function into this network architecture. Their effects are analyzed through ablation experiments, and network robustness is also studied to clarify the impact of low-quality neurons on model performance.
- Our method is quantitatively assessed on neuron classification and retrieval tasks, demonstrating superior performance compared to traditional morphometrics and other baseline methods. Specifically, our method achieves an accuracy of 75.5% in judging the quality of reconstructed samples. By utilizing our extracted features, we are able to obtain neuron cluster results similar to those reported in previous studies. Experimental results also indicate that our morphological embedding is highly compatible and easily combinable with other features.

The basic principle, formula derivation, as well as network architecture of MorphoGNN are expounded in section III. Section IV introduces the experiment datasets and processing details, and presents the results of several morphological applications. Ablation study and network robustness are discussed in section IV-H. Finally, the limitation of our method and the possible follow-up work are discussed in section V.

II. RELATED WORK

A. Morphometrics

Brain-wide and high-resolution optical microscopes [1], [2] enable the acquisition of complete morphological images of neuronal axons and dendrites, providing a rich source of morphometrics for describing neuronal topology. Along with fundamental morphometrics such as total neuron fiber length, number of sections, and density maps [11], [12], recent morphological studies have concentrated on more refined and intricate features. For example, Gillet et al. [22], [23] encode the branches of axons and dendrites of nerve fibers, and calculate the similarity between trees by comparing the corresponding relationships of branches. Kanari et al. [13] and Sizemore et al. [24] describe neuronal topology based on persistent homology and fiber evolution in brain space, respectively. Le Gao et al. [25] propose a method for measuring the similarity between two neurons based on the distance that each point of one neuron maps to another in the other neuron. While this method cannot be used as a morphological description of individual neurons, it can be used to measure differences between projection patterns of paired neurons. Laturus et al. [12] construct a binary classifier based on logical regression and PCA preprocessing, systematically comparing the effects of density maps, persistent images, morphological statistics, and morphological distribution. Typically, these morphometrics focus on specific characteristics of neuron morphology, and researchers often combine them based on their experiences to study a particular task. The development of a method capable of automatically designing morphological features suitable for application could greatly assist neuroscience research. Additionally, the comprehensive and systematic comparison of morphological representations is also an important area for future research.

B. Learning Based Morphometrics

Learning based methods have made significant contributions to the biomedical fields, including brain region segmentation [26], [27], medical image registration [28], [29] and neuron reconstruction [30], [31]. However, these successful applications are based on regular data, such as medical images. The irregularity of neurons presents a challenge for applying deep learning techniques to morphological data, resulting in less work on morphological features based on learning methods. Inspired by [18], Schubert et al. [14] use a set of images to describe a section of nerve fibers. They analyzed the images through a two-dimensional convolutional neural network to extract morphological features indirectly. This learning-based method has an advantage over traditional morphometrics, as neuroscientists no longer need to consider feature design, but rather focus on the tasks they want to perform. The network automatically learns appropriate features according to the set targets. Schubert's method is tested for a series of automatic tasks such as glia detection on electron microscope data. However, the multi-view based method is vulnerable to interference during the visualization stage of neurons, which can affect the learning of morphological information. MorphVAE [15] operates on walks within the tree structure of a neuron and it can generate new morphologies by sampling new walks from the latent space. This approach processes reconstructed neurons at the branch-level. However, the fixed

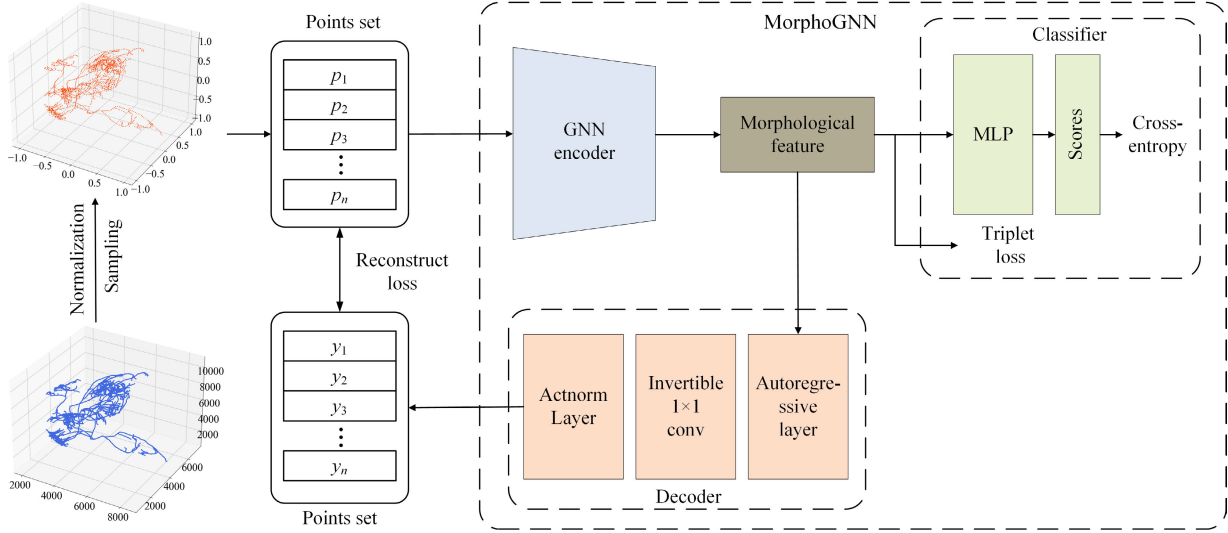


Fig. 1. Overview of the proposed MorphoGNN. This paper proposes a data-driven model for learning the morphological embedding of reconstructed neurons. Through the preprocessing steps, the nerve fiber is converted into a standard point cloud format. We employ a graph neural network to model the geometric relationship between points, and then capture the point-level morphological features of neurons. Two training strategies are provided to satisfy different morphological tasks. By fitting different labels, the model can supervised learn neuron characteristics suitable for different scenes. When labels are unavailable or need to describe their own morphological patterns, the model can be trained via self-supervision by constructing the training target to recovering neural shapes from hidden representations.

walking length and quantity may limit its flexibility in processing complex morphologies. Currently, research is needed to develop an approach that directly learns the morphological characteristics of a single neuron at the point-level.

III. METHOD

The overall framework of our proposed morphological embedding for a single neuron is shown in Fig. 1. After normalization and farthest point sampling (see in section IV-B), the morphology of each neuron can be represented as a point cloud $P = \{p_i \in \mathbb{R}^3 = \{x_i, y_i, z_i\} | i = 1, 2, \dots, n\}$, x_i , y_i , and z_i represent the three-dimensional coordinates of these nodes, respectively. Then a graph neural network is employed to generate morphological features of reconstructed neurons. To increase the applicability of our network, we offer both supervised and self-supervised learning strategies: (i) fitting existing classification labels or (ii) decoding neuron shapes from hidden features. This chapter details our methodology, including the principles behind learning point-level morphological features, the refined network structure, and the joint loss function employed.

A. GNN Encoder

Inspired by [21] and [16], we design a graph neural network that learns the geometric relationship between reconstructed points and then directly generates the representation embedding of the neuronal fibers, instead of manually designed morphometrics.

The problem focused in this paper can be described in (1):

$$h = f_n(P)$$

$$P = \{p_i \in \mathbb{R}^3 = \{x_i, y_i, z_i\} | i = 1, 2, \dots, n\}, \quad (1)$$

where f_n represent a embedding function. The function is supposed to capture the global features h of the reconstructed irregular neural fiber P .

Specifically, we update the graphs dynamically by the K-nearest neighbors(KNN) method and continuously learn node features for the higher dimensions. Node features at different hidden layers are linked to avoid vanishing-gradient. Then a double-pooling operator is utilized to capture the global features of neuronal fibers. The following are detailed explanations of key elements of the proposed method.

B. Local Graph Update

The points in cloud form a graph with the nearest K points, then new features of the central node are learnt from its neighbor nodes. As node features are updated, the graphs are dynamically recomputed between layers. A local graph G^l composed of vertices V and edges E at layer l can be defined as in (2):

$$G^l = (V_K^l, E_K^l), V_K^l \subset \mathbb{R}^C, E_K^l \subset V_K^l \times V_K^l, \quad (2)$$

where C and K represent the dimension of node features and the number of neighbors, respectively. We define the geometric relationship between center node and neighbor nodes as edges, which can be calculated by (3):

$$G^l = (V_{K_{l+1}}^l, E_{K_l}^l)$$

$$V_{K_{l+1}}^l = \{p_c, p_i | i = 1, 2, \dots, K_l\}$$

$$E_{K_l}^l = \{e_i | i = 1, 2, \dots, K_l\}$$

$$e_i = p_i - p_c, \quad (3)$$

where e_i is the directed edge from one of the neighbor nodes p_i to the updating node p_c .

C. Local Feature Update

The initial feature of each node for a single neuron is a three-dimensional vector (x_i, y_i, z_i) , and the local feature is updated to a higher dimension via the EdgeConv [21] layers.

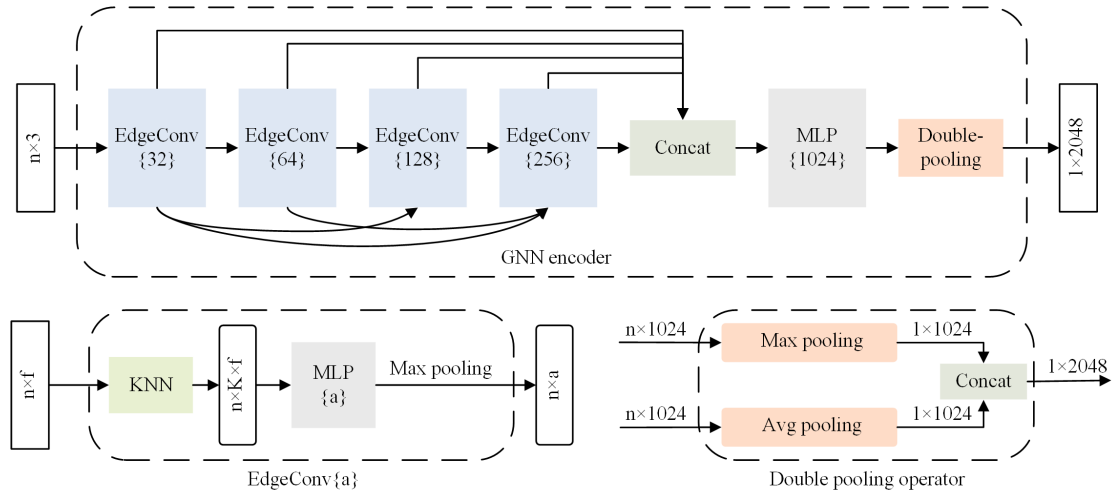


Fig. 2. Architecture of GNN encoder. Our encoder is a modular graph neural network that takes the three-dimensional coordinates of n points as input and outputs the global feature of neurons. EdgeConv $\{N\}$ denotes that after this EdgeConv layer, node features are expanded to N dimension. MLP represents for multi-layer perceptron. Point features at different levels are densely connected to reduce vanishing-gradient. Then the morphological embedding is combined of the outputs of a max-pooling and an average-pooling layer.

For a local graph, the local feature of the updating point p_c can be generated by (4):

$$s'_c = \max\{g(e_1), g(e_2), \dots, g(e_K)\}$$

$$g(e_i) = \text{LeakyReLU}(\Phi \cdot s_c + \Theta \cdot (s_i - s_c)), \quad (4)$$

where Φ and Θ are trainable parameters, while $g(e_i)$ represents the hidden vector of the edge e_i from the neighbor node p_i to p_c . The new local feature s'_c can be obtained through a max-pooling operation for all edge vectors. Since e_i is combined of the old local feature s_c , initially coordinates, and the morphologic information of the neighborhood $s_i - s_c$, s'_c can be considered as the morphological features of neurons in another representation.

D. Double-pooling Operator

The dimension C increases rapidly over several updating iterations on the node features S . To improve the local feature utilization, the global feature of neurons is computed by two kinds of pooling operators for node features, which could be described as shown in (5):

$$h = \max\{S\} || \text{avg}\{S\}, S \subset N \times C, \quad (5)$$

where h and N denotes the features and the number of points of a single neuron, respectively. And $||$ represents the concatenation operator. Max-pooling captures the prominent feature, while average-pooling preserves the overall feature strength. Their outputs are concatenated so the final shape of h is $1 \times 2C$.

E. Encoder Network Architecture

The network architecture of the encoder is depicted in Fig. 2. We begin by increasing the dimension of neuron node features from 3 to 256 through four EdgeConv layers. These features are then aggregate into the global features of nerve fibers using a double-pooling operator. To reduce vanishing-gradient and enhance feature utilization, node features at different levels are densely connected. The input of each EdgeConv layer, except for the first, is derived from

the outputs of each previous EdgeConv layer. To obtain rich global features, we pass the output of the last EdgeConv layer through a maximum-pooling layer and an average-pooling layer, respectively. Their outputs are then concatenated to form the final feature vector of nerve fibers.

F. Supervised Classifier

To fit different category labels, MorphoGNN can independently learn neuron features suitable for different tasks, such as morphological classification or quality classification, instead of manually designing morphometrics by the experience of neuroscientists. For classification, we aim to find a function f_c that maps the features into a probability distribution \mathcal{P}_d as defined in (6):

$$\mathcal{P}_d = f_c(h), h = f_n(p_1, p_2, \dots, p_n) \quad (6)$$

In some cases, such as quality classification in section IV-F, we need to jointly classify with the corresponding microscopic image features, which can be described in (7):

$$\mathcal{P}_d = f_c(h, h'), h = f_n(p_1, p_2, \dots, p_n), \quad (7)$$

where h' denotes image features learned by other methods.

In addition to the cross-entropy function commonly applied in classification tasks, we also use triplet loss [32] to deal with complex and diverse neurons. Considering the morphological embedding h_i of a single neuron, h_p and h_n are defined to represent the embeddings of the same and different categories in the training batch, respectively. So triplet loss L_t can be calculated by (8):

$$L_t = \max(d(h_i, h_p) - d(h_i, h_n) + \text{margin}, 0), \quad (8)$$

where d is a function to calculate the euclidean distance between two embeddings. And margin is a constant that always greater than 0, making $d(h_i, h_p)$ smaller and $d(h_i, h_n)$ larger. As described in (8), triplet loss shortens the distance of similar morphological embeddings and lengthens the distance of heterogeneous morphological embeddings, so as to identify neurons with minor inter-class differences. The total

TABLE I
INFORMATION OF THE 7-CLASS NEURONAL MORPHOLOGICAL DATASET

	Amacrine	Aspiny	Basket	Bipolar	Pyramidal	Spiny	Stellate
Sample size	232	98	130	355	239	247	92
The minimum number of points	1617	1118	1072	1163	1139	1468	1099
The maximum number of points	5719	5986	5354	5462	5701	5690	4603
The average number of points	3614	2873	2065	2117	2002	2920	1953

criterion L can be expressed as the weighted sum of two loss functions as below (9):

$$L = w_1 L_c + w_2 L_t, \quad (9)$$

where L_c represent the cross-entropy loss function, w_1 and w_2 are the weights of the two loss functions.

G. Self-Supervised Decoder

Certain neural tasks, such as unsupervised neuron clustering, necessitate unsupervised training. And in some cases, acquiring precise neuron labels can prove to be challenging. In such situations, MorphoGNN can adopt a self-supervised approach to learn the morphological feature of a single neuron. We achieve self-supervised learning by decoding the spatial structure of neurons from their hidden features h , and then computing the reconstruction log-likelihood [33] with input neurons P . The reconstruction loss L_{recon} can be described in (10):

$$L_{recon}(P; \phi, \psi) \triangleq \mathbb{E}_{E_\phi(h|P)}[\text{Log}D_\psi(P|h)], \quad (10)$$

where E_ϕ and D_ψ are the encoder and decoder of MorphoGNN. And ϕ and ψ represent the trainable parameters of encoder and decoder, respectively. The network architecture of the decoder follows the flow-based model [34], [35], [36], composed of the modules of an autoregressive layer, an invertible 1×1 convolution layer, and an actnorm layer.

IV. EXPERIMENT

This section outlines the data and training process for the proposed MorphoGNN and presents the model’s quantitative performance, as well as that of other baseline models, across various tasks. We also provide a detailed report on the results of the network ablation and robustness tests in this section.

A. Data Description

Three datasets are utilized to evaluate our model’s performance: a 7-class morphological dataset, a dataset comprising 6357 long-range mouse neurons, and a neuron reconstruction blocks dataset. These datasets are employed to assess morphological classification and retrieval, clustering, and reconstruction quality classification, respectively.

The 7-class dataset contains 1393 reconstructed neurons collected from NeuroMorpho.Org [37] and we denote this dataset as Neuro7. These neurons are labeled with their morphological category, so they can be used for supervised learning. However, the raw morphological data is non-standardized due to the various sources of the contributors. As shown in Table I, we selected seven representative types of mouse nerve cells, each of which results from multiple research groups.

Aspiny neurons come from three groups [38], [39], [40]. Spiny neurons are all from [38]. Reference [41] contributes all of the amacrine and bipolar neurons. Stellate neurons also consists of three groups [42], [43], [44], while basket neurons are from [45] and [46]. Finally, pyramidal neurons are totally taken from four groups [47], [48], [49], [50]. The number of points that make up a neuron structure varies from 1×10^3 to 6×10^3 in this dataset.

In addition to Neuro7 dataset, we also utilize an open-source dataset consisting of 6,357 long-range neurons mapped from the mouse prefrontal cortex, which have been previously identified into three large subtypes and sixty-four small groups based on their projection pattern [25]. In this paper, we attempt to classify these same neurons using our proposed morphological embedding and compare our results with existing methods.

Moreover, we introduce a reconstructed neuron dataset consisting of microscopic images collected from a collaborative platform, where neural segments were manually reconstructed by researchers from three-dimensional optical microscopic image blocks. The reconstructions were subsequently verified by other researchers to ensure their accuracy. This dataset consists of 3,000 paired samples, with an equal proportion of correct and incorrect samples. Each neuron segment is comprised of 1024 points, corresponding to a $64 \times 64 \times 64$ image block. These three datasets are randomly split into training and validation sets at a 7 : 3 ratio.

B. Morphological Data Preprocessing

Prior to network training, morphological data undergoes a two-step preprocessing procedure to transform them into regular point clouds.

Farthest point sampling is used to address the issue of inconsistent point numbers per sample and to accelerate training by enabling minibatches of differently sized neurons to be grouped together. We iteratively select the farthest point from the existing set of sampling points to obtain neurons of any size. Additionally, to ensure that graphs used for training contain the same number of points, we fill in some small reconstructed neurons with zero points $(0, 0, 0)$ until they reach the same size as the other neurons.

Normalization is another necessary step in preprocessing neuronal morphological data. The absolute location of neurons in the brain space can lead to a wide difference between the coordinates of the reconstructed points. Furthermore, there is no unified standard for the brain spatial location of neurons that have not been uniformly registered. The following formula (11) is used to normalize the neuron point coordinates:

$$v'_i = v_i - (v_{min} + v_{max})/2, v_i = (x_i, y_i, z_i) \\ V = \left\{ \frac{v'_i}{\max(v'_j)} \mid i = 1, 2, \dots, n \right\} \quad (11)$$




Fig. 3. Sampling and normalization for neurons. From left to right are the original neuron, sampling 2048 points and sampling 1024 points.

, where v_{max} and v_{min} are the maximum and minimum initial coordinate values v_i over a single neuron, respectively. And V represents the normalized coordinates set. we identify the smallest unit cube that can surround the original neuron and use it as the new origin for normalization. Fig. 3 shows that the normalized neuron retains its original shape, with its center serving as the spatial origin, and that the coordinate value of each point falls within the range $(-1, 1)$.

C. Implementation Details

Here we explain the experiments setup, especially the training process of our network. For supervised training, we employ the Adam optimizer [51] with a learning rate of 10^{-3} and momentum of 0.9 to train the entire network. The network is trained for 300 epochs, with a batch size of 16, and every 20 epochs, and the learning rate is decreased by half every 20 epochs. The cross-entropy and triplet loss are used jointly, with a weight ratio of 1 : 1. For self-supervised training, we set the initial learning rate of the Adam optimizer to 2×10^{-3} and reduce it by half after every 500 epochs. The network is trained for 3500 epochs, using a batch size of 32. All experiments are conducted on a single NVIDIA RTX 5000 GPU with 16GB of graphics memory, as well as four NVIDIA Tesla V100 graphics cards, each with 16GB of memory.

D. Morphological Classification

We first evaluate the performance of our proposed features and traditional morphometrics on the 7-class neuron dataset. Sixteen basic neuronal morphometrics [52] are used to form the feature vector of each neuron, including total length, number of tips, number of sections, number of segments, number of bifurcation points, maximum branch order, maximum neurite length, maximum radial distance, and maximum/median section tortuosity, as well as maximum/median/minimal section length and maximum/median/minimal remote bifurcation angle. These measurements are then used for neuron classification via a simple multi-layer perceptron network with the same settings as other methods. We trained MorphVAE via supervised learning on our dataset to achieve its best performance. Additionally, since reconstructed neurons are transformed into standard point clouds after preprocessing, we include three recent point cloud networks for comparison. Table II shows the overall accuracy(OA) and mean class accuracy(MA) of these methods.

Our network achieves the highest overall accuracy(85.58%) and mean class accuracy(79.45%), which are about 2.5% and 6.8% higher than those of PointNet++. The mean class accuracy of each method is significantly lower than the overall accuracy due to the less-obvious differences between classes

TABLE II

RESULTS OF MORPHOLOGICAL CLASSIFICATION ON NEURO7

Method	MA(%)	OA(%)
PointNet	50.48	63.46
PointNet++	72.69	83.08
DGCNN	71.56	81.73
Morphometrics	77.95	82.45
MorphVAE	63.08	73.73
Ours	79.45	85.58
Ours+Morphometrics	86.01	91.35




Fig. 4. UMAP visualization of different features.

and imbalanced class populations. Interestingly, the performance of morphometrics is close to, and even surpassed, some learning-based features, including PointNet++ and DGCNN. This could be attributed to the loss of some morphological information, such as the absolute size of neurons, in the preprocessing step, which limits the performance of learning-based methods. However, this drawback can be overcome by combining them with morphometrics. When our embedding and 16 morphological statistics are concatenated for classification, the overall accuracy and mean class accuracy are further improved by 5.8% and 6.6%, respectively. We use UMAP [53] to visualize the features extracted by these methods(see in Fig. 4), where each point represented the feature of a single neuron after dimensionality reduction. The features of the same class extracted by our method are close to each other and far away from other types. However, some morphologically similar categories, such as spiny and aspiny, are still not clearly distinguished.

E. Retrieval on Morphological Dataset

Current Large neuronal morphological databases [37], [54] typically support search based on meta information such as contributors and species. However, relying on such meta-information can result in a loss of valuable morphological details. Here we demonstrate the effectiveness of direct query based on embedding similarity between neurons.

These networks trained in section IV-D are used previously to extract features from each neuron to build the feature databases. Fig. 7(c) shows the time required to build the




Fig. 5. Samples of reconstruction quality classification.

database using different features. PointNet++ model with complex multi-scale learning takes the longest time to build a feature database, approximately 25 times longer than PointNet. The time consumption of traditional morphometrics is the second longest, approximately three times that of DGCNN. To only measure the performance of features, we retrieve similar neurons by simply calculating the cosine similarity between paired features. The accuracy of each retrieval operation is defined as the proportion of the same types among the ten most similar neurons retrieved. We searched each neuron in the Neuro7 dataset in turn, recorded the overall accuracy, and mean class accuracy. Fig. 7(a) shows the five most similar neuron fibers retrieved from the same neuron fiber by the six methods. The retrieval results of other methods are mixed with various kinds of neurons, while our method retrieved five neurons of the same class. As shown in Fig. 7(b), PointNet++ and MorphoGNN achieve the best retrieval performance. The retrieval performance of traditional morphometrics is between that of DGCNN and MorphVAE. Among them, our method achieves the highest overall accuracy(84.19%) and mean class accuracy(77.41%), respectively.

F. Reconstruction Quality Classification

Our research investigates the potential of combining our morphological embedding method with optical microscopic images in neuron reconstruction. Quality control has always been an important part of neuron reconstruction, traditionally achieved through multi-person collaboration where one team reconstructs nerve fibers and another team evaluates the accuracy of the reconstruction. While effective, this approach is labor-intensive and relies heavily on human resources. To reduce costs, we explore the possibility of automated methods that can assist in certain steps of the process. To this end, we leverage a collaborative reconstruction platform (see in section IV-A) that provides labeled data for training and testing our proposed method. Fig. 5 shows some data samples used in the reconstruction quality classification. One criterion to judge the quality is whether the reconstructed nerve fibers match the signals in microscopic images.

Fig. 6 displays the quality classification framework, which comprises two channels: MorphoGNN and 3D U-Net [55]. These channels model the morphology of neurons and their corresponding microscopic images, respectively. We concatenate the features of these two modal data for quality classification. Table III summarizes the results of the reconstruction quality classification, divided into two types: using

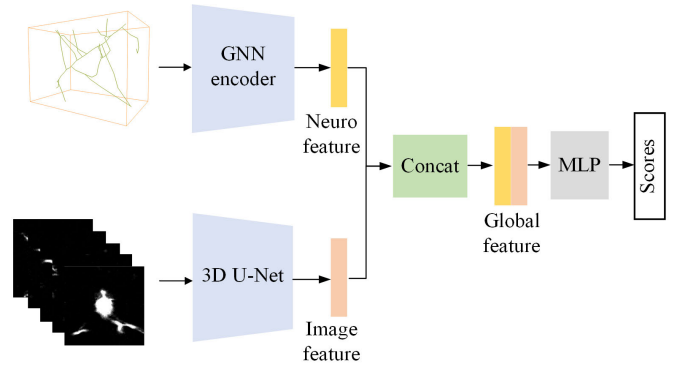


Fig. 6. Dual-channel network for reconstruction quality classification. We model neurons and their corresponding microscopic images respectively, and then concatenate their features for classification.

TABLE III
RESULTS OF RECONSTRUCTION QUALITY CLASSIFICATION

	Dual-channel		Single channel	
	Correct	Incorrect	Correct	Incorrect
Precision	75.23%	75.92%	60.95%	81.25%
Recall	76.38%	74.76%	90.27%	42.18%
F1-score	75.80%	75.34%	72.77%	55.53%

only neuronal channel and using both channels. When only the morphology of neurons is modeled, the precision of the network for correct samples is low(60.95%), while the recall rate is high(90.27%). This indicates that the network tends to judge all samples as correct. However, when both channels are combined, the judgment results of positive and negative samples are relatively balanced, resulting in an average precision of 75.57%. This method can automatically judge simple samples and identify uncertain ones that require manual labor for further assessment. Overall, our proposed method shows promise in reducing the labor required for neuron reconstruction and providing accurate judgments of reconstruction quality.

G. Neuron Clustering

Morphological subtyping is a top-down process. With the refinement of criteria, neurons with the similar anatomical characteristics can be clustered into the same subtype. Identifying neuron subtypes advances the understanding of brain structure and function, so it has always been a hot spot in neuroscience research [56], [57]. The difference between neuron clustering and neuron classification is that the former is an unsupervised task, where the subtypes of neurons are not known in advance. Current neuron clustering mainly relies on the manual features designed by neuroscientists that reflect the fields they are interested in. For instance, [25] mainly considers the information of brain spatial location and morphology to calculate the pairwise similarity of 6357 long-range neurons reconstructed in the mouse prefrontal cortex. These individual neurons are precisely identified into three large subtypes, including intratelencephalic neurons(IT), corticothalamic neurons(CT) and pyramidal tract neurons(PT), as well as 64 small groups further.

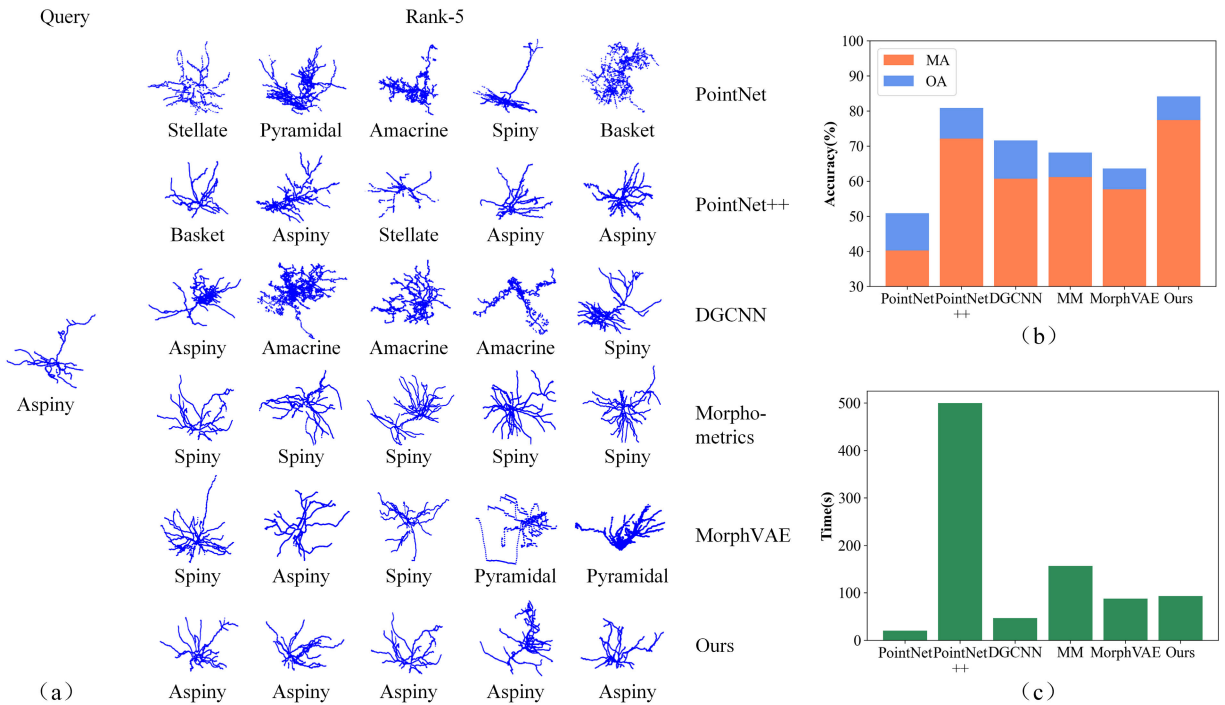


Fig. 7. Retrieval results based on different methods. (a) is an example of five most similar neurons retrieved based on different features for a randomly selected query neuron. (b) shows the retrieval accuracy of six methods, which is defined as the proportion of the same types among the ten most similar neurons retrieved. (c) compares the time-consuming of different methods to establish the feature database.

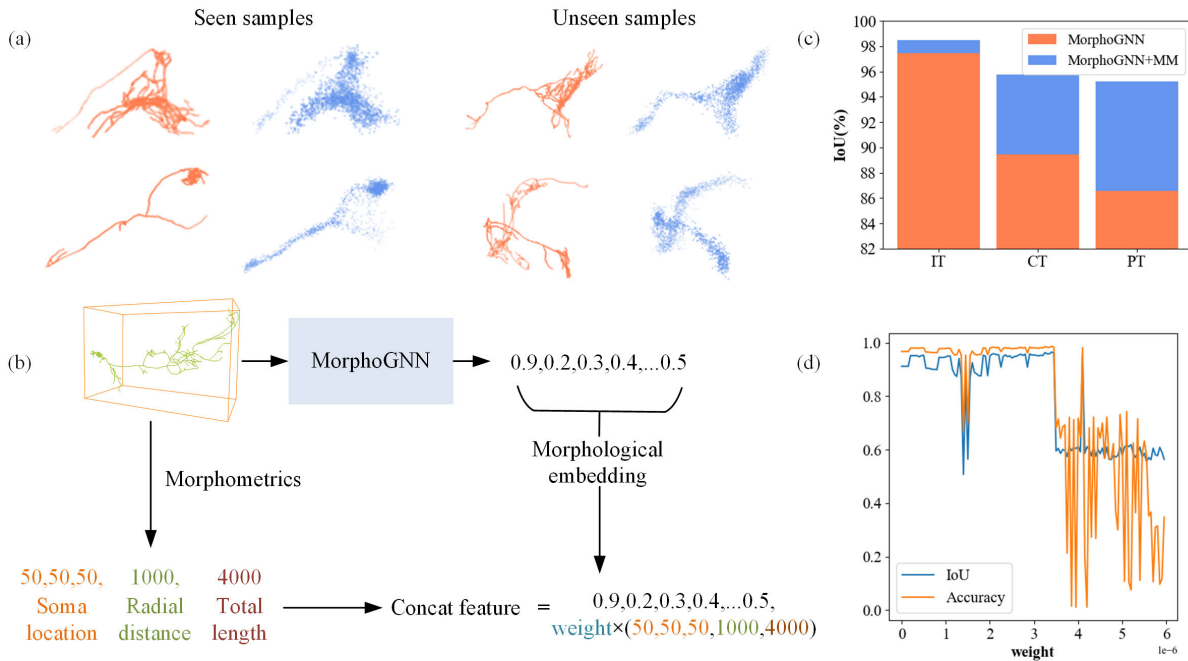


Fig. 8. Neuron clustering based on self-supervised learning features. (a) shows the reconstructed shape by the self-supervised model from the hidden features of neuron samples seen and not seen during the training stage. (b) displays the process of enhancing the description of neuronal information through combining our morphological embedding with traditional morphometrics. (c) compares the cluster similarity between our features and the existing work. The evaluation metric is intersection of union (IoU). (d) illustrates the influence of weight on the result of neuron clustering when our features combined with traditional morphometrics.

Here we compare the results of the three major subtypes of these prefrontal cortical neurons hierarchical clustered based on our morphological embeddings with [25], using the evaluation metrics of intersection of union (IoU) and accuracy. IoU describes the coincidence between subtypes. When the existing

cluster results are taken as absolutely correct references, accuracy can also be calculated. MorphoGNN is trained through self-supervised learning (see in section III-G). We construct the objective function by decoding the three-dimensional structure of neurons from hidden representations. Fig. 8(a) illustrates

TABLE IV

ABLATION TEST OF NETWORK COMPONENTS ON NEURO7. HERE ✓ INDICATES THIS COMPONENT IS APPLIED IN THE NETWORK

Dense connection	Avg pooling	Max pooling	Triplet loss	OA(%)	MA(%)
	✓	✓	✓	82.69	74.07
✓		✓	✓	71.26	57.78
✓	✓		✓	84.13	76.90
		✓	✓	81.73	73.21
	✓		✓	79.33	71.17
✓	✓	✓		81.49	74.14
✓	✓	✓	✓	85.58	79.45

the reconstruction performance of neurons seen and unseen in the training stage. Our hidden representations can restore an approximate structure of neurons, which is hard to achieve by traditional morphometrics. However, there is still a gap between reconstructed structures and real neurons, and designing special constraint terms to make the structure skeleton may improve this defect.

As shown in Fig. 8(c), the mean IoU and accuracy of two cluster results are only 91.15% and 96.70% when our morphological embeddings are used for hierarchical clustering. This is because the two methods focus on different areas. Our model focuses on the relative morphology of neurons, lacking the information on absolute size and brain spatial location, which are deeply concerned in [25]. Hence we introduce some additional information to strengthen our features. Specifically, we weighted concatenate several simple morphometrics, including soma location, radial distance, and total length, to complement this shortcoming(see in Fig. 8(b)). The influence of weight is shown in Fig. 8(d). After simple morphometrics are jointed, the mean IoU and accuracy increase to 96.49% and 98.62%, respectively(see in Fig. 8(c)). This obvious increase shows that our morphological embedding is an open architecture, which is easy to combine with other features. When more abundant additional information is added, the cluster results may be further improved. Overall, our study demonstrates the potential of using morphological embeddings for neuron clustering and the importance of incorporating additional information to improve the accuracy of the clustering results.

H. Ablation and Robustness Test

We further investigate the effect of several components specially designed in our network on the final performance. Table IV shows the classification performance of multiple versions of MorphoGNN on the Neuro7 dataset, including dense connection, double-pooling operation and triplet loss.

Dense connection greatly impacts on network performance, improving overall accuracy from 82.47% to 85.58% and mean class accuracy from 74.07% to 79.45%. The double-pooling operation also plays a notable role. Compared with it, no matter which single pooling operation is adopted, the mean class accuracy can be reduced by up to 2.8%. Moreover, maximum-pooling performs better than average-pooling when capturing global features. The triplet loss function brings an

TABLE V

RESULTS OF CLASSIFICATION WITH SELF-SUPERVISED TRAINING. ST AND SST REPRESENT SUPERVISED AND SELF-SUPERVISED TRAINING

	Overall accuracy		Parameters	
	ST	Add SST	ST	Add SST
PointNet	63.46%	64.46%	3.8×10^5	2.9×10^7
DGCNN	81.73%	83.09%	1.5×10^6	2.9×10^7
MorphoGNN	85.58%	84.38%	1.8×10^6	3.0×10^7

TABLE VI

ABLATION TEST OF LOSSES ON NEURO7

	Cross-entropy	Triplet loss	Recon loss	OA(%)	MA(%)
	✓			81.49	74.14
MorphoGNN	✓	✓		85.58	79.45
	✓		✓	85.58	78.19
	✓	✓	✓	84.38	76.47

additional 4.1% and 5.3% improvement in overall accuracy and mean class accuracy for the network, respectively.

In the biomedical domain, unsupervised learning has become increasingly important. Therefore, we further investigate the effects of self-supervised training proposed in this study. Specifically, we add self-supervised training to the morphological classification task in section IV-D and compare it with the original classification accuracy to determine whether self-supervised training could improve the performance of supervised classification tasks. Table V shows that the added self-supervised training slightly improves the overall accuracy of PointNet (1.00%) and DGCNN (1.36%), but at the cost of about 20 million more network parameters. However, the overall accuracy of MorphoGNN decreased by 1.21% after adding self-supervised training. To clarify this issue, we conduct further ablation experiments on the loss functions of MorphoGNN. Table VI illustrates that the reconstruction loss of self-supervised training conflicts with the triplet loss. While triplet loss and reconstruction loss can each improve the overall accuracy by 4.09%, when used together, the overall accuracy is only improved by 2.99%. These results indicate that self-supervised training does not always improve the performance of supervised tasks, and combining them requires careful research.

Due to inevitable defects in imaging and reconstruction, incomplete and inaccurate morphological data are often obtained. Discarding all these samples is costly, so algorithms are expected to be robust to low-quality reconstructed neurons. In this section, we also discuss the performance of those deep networks mentioned in this paper against two attacks: point dropping and point-wise Gaussian noise. The former randomly drops points in a single neuron to simulate incomplete reconstruction, while the latter adds point-wise noise to make the reconstructed shape deviate from reality. The neurons in the testset of Neuro7 are attacked on varying degrees, and the classification accuracy of these pretrained networks is recorded.

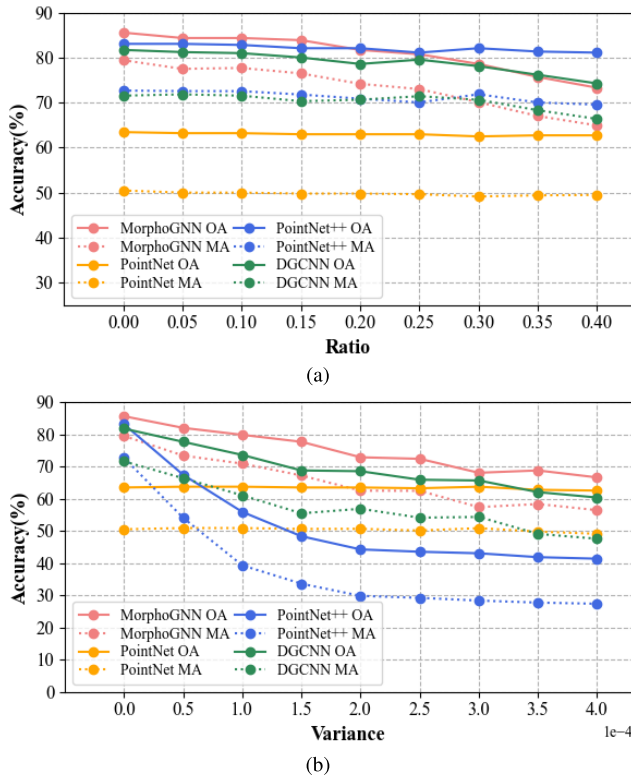


Fig. 9. Influence of two attacks on the classification performance of learning based methods. This experiment is done on the Neuro7 dataset. The abscissa of (a) is the ratio of discarded points, while that of (b) is the variance of Gaussian noise.

Fig. 9 shows the overall accuracy and mean class accuracy of deep networks against different degrees of point dropping and Gaussian noise attacks. The classification accuracy of all networks shows a downward trend, indicating that network performance is generally weakened when the morphology of neurons is destroyed. For instance, when MorphoGNN is attacked by Gaussian noise with a variance of 4×10^{-4} , its overall accuracy and mean class accuracy are decreased by around 19% and 23%, respectively. When 40% points are lost, these two metrics of MorphoGNN are reduced by 14% and 18%, respectively. The performance degradation of networks against point dropping is slower than that against random noise. When the variance of Gaussian noise increases to 2.5×10^{-4} , the overall accuracy and mean class accuracy of DGCNN rapidly drop to less than 80% of the initial network. However, the same network's performance is hardly decreased when the ratio of dropped points gradually increases to 35%. The same trend is shown on PointNet++. The better robustness to point discarding may be because some zero points are filled into neuron data during data preprocessing (see in section IV-B). Although the dropping operation is before filling zero points, these filled points can be regarded as a disturbance to the number of effective points added in the training stage, which makes networks less sensitive to point dropping than random noise. However, lightweight deep network such as PointNet is almost stable under different degrees of attacks. Overall, experimental results show that low-quality neurons significantly weaken network performance. Therefore, promoting network robustness is a necessary area for future work.

V. CONCLUSION

We present MorphoGNN, a novel method for capturing the morphological embeddings of reconstructed neuron fibers based on a proposed deep network, aiming to improve the utilization of existing approaches on the morphological information of neurons by a data-driven approach. By learning point-level geometric information, our proposed network can capture the feature of a single neuron and be quantitatively compared with other existing morphological features. We demonstrate that the network can be trained via both supervision and self-supervision. Moreover, the learned morphological embeddings can be combined with the features from microscopic images as well as traditional morphometrics, thereby adapting to multiple tasks. We demonstrate the effectiveness of our approach in neuron classification, retrieval, reconstruction quality classification, and neuron clustering on three thousand-level morphology datasets.

Limitations and future work: We construct a self-supervised objective function by reconstructing the shape of neurons from hidden representations, but our current model is not capable of accurately reconstructing the neuron structure due to the absence of a constraint function that reflects the prior knowledge of neuron structures. Our method does not consider that branches of neurons are usually slender and directional in order to transmit electrical signals, which may account for the imprecise reconstruction. Furthermore, the shape generated by our method is a point cloud rather than a topological tree structure, which necessitates further work to predict the connection between points. Additionally, our method does not incorporate information regarding the spatial location of the brain, resulting in the cluster results of mouse prefrontal cortex neurons being slightly different from those obtained in previous research. Although combining our approach with existing morphometrics can alleviate this deficiency, a significant gap still remains. Therefore, the integration of multimodal information, such as morphology, brain location, genomics, and proteomics, is necessary to fully comprehend the structure and function of neurons, and ultimately achieve a census of all brain cells.

REFERENCES

- [1] H. Gong et al., "High-throughput dual-colour precision imaging for brain-wide connectome with cytoarchitectonic landmarks at the cellular level," *Nature Commun.*, vol. 7, no. 1, p. 12142, Jul. 2016.
- [2] Q. Zhong et al., "High-definition imaging using line-illumination modulation microscopy," *Nature Methods*, vol. 18, no. 3, pp. 309–315, Mar. 2021.
- [3] X. Yang et al., "Observing single cells in whole organs with optical imaging," *J. Innov. Opt. Health Sci.*, vol. 16, no. 1, Jan. 2023, Art. no. 2330002.
- [4] A. Li et al., "Micro-optical sectioning tomography to obtain a high-resolution atlas of the mouse brain," *Science*, vol. 330, no. 6009, pp. 1404–1408, Dec. 2010.
- [5] H. Zhou et al., "GTree: An open-source tool for dense reconstruction of brain-wide neuronal population," *Neuroinformatics*, vol. 19, no. 2, pp. 305–317, Apr. 2021.
- [6] T. Quan et al., "NeuroGPS-tree: Automatic reconstruction of large-scale neuronal populations with dense neurites," *Nature Methods*, vol. 13, no. 1, pp. 51–54, Jan. 2016.
- [7] R. Parekh and G. A. Ascoli, "Neuronal morphology goes digital: A research hub for cellular and system neuroscience," *Neuron*, vol. 77, no. 6, pp. 1017–1038, Mar. 2013.
- [8] D. H. Hall and M. Treinin, "How does morphology relate to function in sensory arbors?" *Trends Neurosciences*, vol. 34, no. 9, pp. 443–451, Sep. 2011.

- [9] T. Sasaki, N. Matsuki, and Y. Ikegaya, "Effects of axonal topology on the somatic modulation of synaptic outputs," *J. Neurosci.*, vol. 32, no. 8, pp. 2868–2876, Feb. 2012.
- [10] Z. Šišková et al., "Dendritic structural degeneration is functionally linked to cellular hyperexcitability in a mouse model of alzheimer's disease," *Neuron*, vol. 84, pp. 1023–1033, 2014.
- [11] L. D. F. Costa, K. Zawadzki, M. Miazaki, M. P. Viana, and S. N. Taraskin, "Unveiling the neuromorphological space," *Frontiers Comput. Neurosci.*, vol. 4, p. 150, Dec. 2010.
- [12] S. Laturus, D. Kobak, and P. Berens, "A systematic evaluation of interneuron morphology representations for cell type discrimination," *Neuroinformatics*, vol. 18, no. 4, pp. 591–609, Oct. 2020.
- [13] L. Kanari et al., "A topological representation of branching neuronal morphologies," *Neuroinformatics*, vol. 16, no. 1, pp. 3–13, Jan. 2018.
- [14] P. J. Schubert, S. Dorckenwald, M. Januszewski, V. Jain, and J. Kornfeld, "Learning cellular morphology with neural networks," *Nature Commun.*, vol. 10, no. 1, Jun. 2019.
- [15] S. C. Laturus and P. Berens, "MorphVAE: Generating neural morphologies from 3D-walks using a variational autoencoder with spherical latent space," in *Proc. 38th Int. Conf. Mach. Learn.*, in Proceedings of Machine Learning Research, vol. 139, M. Meila and T. Zhang, Eds., Jul. 2021, pp. 6021–6031. [Online]. Available: <https://proceedings.mlr.press/v139/laturus21a.html>
- [16] R. Q. Charles, H. Su, M. Kaichun, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 77–85.
- [17] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep hierarchical feature learning on point sets in a metric space," in *Proc. NIPS*, 2017, pp. 1–10.
- [18] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller, "Multi-view convolutional neural networks for 3D shape recognition," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 945–953.
- [19] J. Zhou et al., "Graph neural networks: A review of methods and applications," 2018, *arXiv:1812.08434*.
- [20] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" 2018, *arXiv:1810.00826*.
- [21] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph CNN for learning on point clouds," *ACM Trans. Graph.*, vol. 38, no. 5, pp. 1–12, Oct. 2019.
- [22] T. A. Gillette and G. A. Ascoli, "Topological characterization of neuronal arbor morphology via sequence representation: I–motif analysis," *BMC Bioinf.*, vol. 16, no. 1, pp. 1–15, Dec. 2015.
- [23] T. A. Gillette, P. Hosseini, and G. A. Ascoli, "Topological characterization of neuronal arbor morphology via sequence representation: II–global alignment," *BMC Bioinf.*, vol. 16, no. 1, Dec. 2015.
- [24] A. E. Sizemore, J. E. Phillips-Cremins, R. Ghrist, and D. S. Bassett, "The importance of the whole: Topological data analysis for the network neuroscientist," *Netw. Neurosci.*, vol. 3, no. 3, pp. 656–673, Jan. 2019.
- [25] L. Gao et al., "Single-neuron projectome of mouse prefrontal cortex," *Nature Neurosci.*, vol. 25, no. 4, pp. 515–529, 2022.
- [26] L. Wang et al., "Benchmark on automatic six-month-old infant brain segmentation algorithms: The iSeg-2017 challenge," *IEEE Trans. Med. Imag.*, vol. 38, no. 9, pp. 2219–2230, Sep. 2019.
- [27] M. Ebner et al., "An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI," *NeuroImage*, vol. 206, Feb. 2020, Art. no. 116324.
- [28] G. Balakrishnan, A. Zhao, M. R. Sabuncu, A. V. Dalca, and J. Guttag, "An unsupervised learning model for deformable medical image registration," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 9252–9260.
- [29] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "VoxelMorph: A learning framework for deformable medical image registration," *IEEE Trans. Med. Imag.*, vol. 38, no. 8, pp. 1788–1800, Aug. 2019.
- [30] M. Januszewski et al., "High-precision automated reconstruction of neurons with flood-filling networks," *Nature Methods*, vol. 15, no. 8, pp. 605–610, Aug. 2018.
- [31] J. Zhao et al., "Progressive learning for neuronal population reconstruction from optical microscopy images," in *Proc. MICCAI*, 2019, pp. 750–759.
- [32] A. Hermans, L. Beyer, and B. Leibe, "In defense of the triplet loss for person re-identification," 2017, *arXiv:1703.07737*.
- [33] X. Chen et al., "Variational lossy autoencoder," 2016, *arXiv:1611.02731*.
- [34] G. Yang, X. Huang, Z. Hao, M. Liu, S. Belongie, and B. Hariharan, "PointFlow: 3D point cloud generation with continuous normalizing flows," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 4540–4549.
- [35] H. Kim, H. Lee, W. H. Kang, J. Y. Lee, and N. S. Kim, "SoftFlow: Probabilistic framework for normalizing flow on manifolds," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 16388–16397.
- [36] D. P. Kingma and P. Dhariwal, "Glow: Generative flow with invertible 1×1 convolutions," 2018, *arXiv:1807.03039*.
- [37] G. A. Ascoli, D. E. Donohue, and M. Halavi, "NeuroMorpho.Org: A central resource for neuronal morphologies," *J. Neurosci.*, vol. 27, no. 35, pp. 9247–9251, Aug. 2007.
- [38] L. Madisen et al., "Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance," *Neuron*, vol. 85, no. 5, pp. 942–958, Mar. 2015.
- [39] M. W. Bagnall, C. Hull, E. A. Bushong, M. H. Ellisman, and M. Scanziani, "Multiple clusters of release sites formed by individual thalamic afferents onto cortical interneurons ensure reliable transmission," *Neuron*, vol. 71, no. 1, pp. 180–194, Jul. 2011.
- [40] M. Helmstaedter, K. L. Briggman, S. C. Turaga, V. Jain, H. S. Seung, and W. Denk, "Connectomic reconstruction of the inner plexiform layer in the mouse retina," *Nature*, vol. 500, no. 7461, pp. 168–174, Aug. 2013.
- [41] T. J. Ellender et al., "Embryonic progenitor pools generate diversity in fine-scale excitatory cortical subnetworks," *Nature Commun.*, vol. 10, no. 1, p. 5224, Nov. 2019.
- [42] T. Guillamon-Vivancos et al., "Distinct neocortical progenitor lineages fine-tune neuronal diversity in a layer-specific manner," *Cerebral Cortex*, vol. 29, no. 3, pp. 1121–1138, Mar. 2019.
- [43] D. L. F. Garden, P. D. Dodson, C. O'Donnell, M. D. White, and M. F. Nolan, "Tuning of synaptic integration in the medial entorhinal cortex to the organization of grid cell firing fields," *Neuron*, vol. 60, no. 5, pp. 875–889, Dec. 2008.
- [44] T. Lalanne et al., "Synapse-specific expression of calcium-permeable AMPA receptors in neocortical layer 5," *J. Physiol.*, vol. 594, no. 4, pp. 837–861, Feb. 2016.
- [45] V. K. Vereczki et al., "Synaptic organization of perisomatic GABAergic inputs onto the principal cells of the mouse basolateral amygdala," *Frontiers Neuroanatomy*, vol. 10, p. 20, Mar. 2016.
- [46] L. Rovira-Esteban, Z. Péterfi, A. Vikór, Z. Máté, G. Szabó, and N. Hájos, "Morphological and physiological properties of CCK/CB1R-expressing interneurons in the basal amygdala," *Brain Struct. Function*, vol. 222, no. 8, pp. 3543–3565, Nov. 2017.
- [47] E. K. W. Brennan, S. K. Sudhakar, I. Jedrasiak-Cape, T. T. John, and O. J. Ahmed, "Hyperexcitable neurons enable precise and persistent information encoding in the superficial retrosplenial cortex," *Cell Rep.*, vol. 30, no. 5, pp. 1598–1612, Feb. 2020.
- [48] M. P. Fiske, M. Anstötz, L. J. Welty, and G. Maccaferri, "The intrinsic cell type-specific excitatory connectivity of the developing mouse subiculum is sufficient to generate synchronous epileptiform activity," *J. Physiol.*, vol. 598, no. 10, pp. 1965–1985, May 2020.
- [49] S. C. Borrie et al., "Loss of nogo receptor homolog NgR2 alters spine morphology of CA1 neurons and emotionality in adult mice," *Frontiers Behav. Neurosci.*, vol. 8, p. 175, May 2014.
- [50] Z. Bertels et al., "Neuronal complexity is attenuated in chronic migraine and restored by HDAC6 inhibition," *bioRxiv*, vol. 2020, pp. 1–48, Apr. 2020.
- [51] D. P. Kingma and J. L. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent.*, 2015, pp. 1–15.
- [52] J. Palacios et al., "BlueBrain/NeuroM, v3.2.2," Zenodo, 2022, doi: [10.5281/zenodo.7128173](https://doi.org/10.5281/zenodo.7128173).
- [53] L. McInnes, J. Healy, and J. Melville, "UMAP: Uniform manifold approximation and projection for dimension reduction," *J. Open Source Softw.*, vol. 3, no. 29, p. 861, 2018.
- [54] M. Costa, J. D. Manton, A. D. Ostrovsky, S. Prohaska, and G. S. Jefferis, "NBLAST: Rapid, sensitive comparison of neuronal structure and construction of neuron family databases," *Neuron*, vol. 91, no. 2, pp. 293–311, Jul. 2016, doi: [10.1016/j.neuron.2016.06.012](https://doi.org/10.1016/j.neuron.2016.06.012).
- [55] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: Learning dense volumetric segmentation from sparse annotation," 2016, *arXiv:1606.06650*.
- [56] M. N. Economo et al., "Distinct descending motor cortex pathways and their roles in movement," *Nature*, vol. 563, no. 7729, pp. 79–84, Nov. 2018.
- [57] N. W. Gouwens et al., "Integrated morphoelectric and transcriptomic classification of cortical GABAergic cells," *Cell*, vol. 183, no. 4, pp. 935–953, 2020.