CoNMoR: Contrastive Neuron Morphology Representation Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Morphology of neuron trees is a key indicator to delineate neuronal cell-types, analyze brain development process, and evaluate pathological changes in neurological diseases. Traditional analysis mostly relies on heuristic features and visual inspections. A quantitative, informative, and comprehensive representation of neuron morphology is largely absent but desired. To fill this gap, in this work, we adopt a Tree-LSTM network to encode neuron morphology and introduce a self-supervised learning framework named COntrastive Neuron MOrphology Representation learning (CoNMoR) to learn features without the need of human experts' annotations. We test CoNMoR on 2403 high-quality 3D neuron reconstructions of mouse brains from three different public resources. Our results show that CoNMoR is effective in both classifying major brain cell-types and identifying sub-types. To our best knowledge, CoNMoR is the very first to explore learning the representation of neuron tree morphology with contrastive learning. It has a great potential in shedding new light on quantitative neuron morphology analysis.

1 Introduction

The central nervous system is composed of hundreds to thousands of neuronal subtypes. Subtype is the critical determinant of neuronal function within local and long-range circuits and is identified by a combination of gene expression patterns, dendrite/axon morphology, synaptic connectivity, and electrophysiological properties. In particular, dendritic morphology is a central node of neuronal subtype identity, as dendritic branches receive the vast majority of the presynaptic input onto each neuron and directly shape the electrophysiological properties of the neuron that governs its firing output. Delineating neuron cell-types based on their dendritic morphology is one of the core problems of neuroscience for the past century [50].

With recent advances in imaging [9, 11, 46] and computational methods [20, 25, 41], 3D reconstruction of whole neuron trees with high-throughput is now accessible for mouse brain [26, 43]. However, due to the lack of effective measurement, most analysis of these 3D reconstructions are still largely based on biologists' visual inspection. Considering the enormous number of neurons in the brain and the 3D formation of neuron morphology, such visual inspection is subjective, biased, and of low throughput.

The past decade has witnessed the success of deep learning in learning representations of almost all kinds of data such as image, language, sound, and graph [2]. However, as a special type of graph or graphic data, most of the quantitative analysis of neuron trees is still based on simple heuristic features [32]. Only a few attempts have been made in learning deep representations of neuron morphology. In [51], the authors proposed to encode neuron morphology with tree recurrent neural network (TRNN). However, this method relies on manual annotations, limiting its application to supervised training only. In [21], the authors proposed a generative self-supervised learning

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

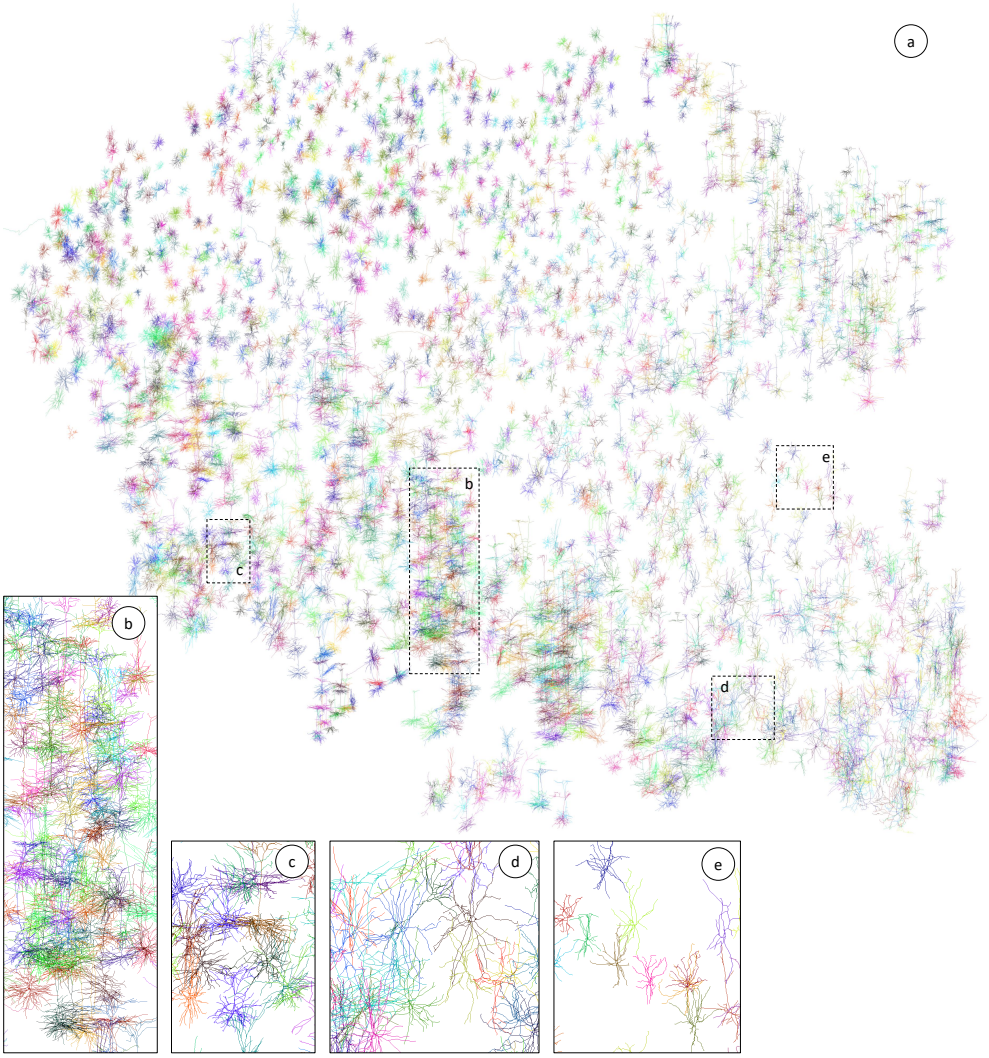

Figure 1: Distribution of neurons in the embedding space of CoNMoR visualized by t-SNE [38]. Thumbnails of dendrites are randomly colored. (b-e) Zoom-in view of regions in (a) showing clustering of isocortex layer 5 cells (b), layer 2/3 cells (c), basket cells (d), and Martinotti cells (e).

(SSL) method for neuron morphology representation learning, named MorphVAE, to reconstruct original neuron trees via a variational autoencoder and random walk sampling. Since this method cannot encode tree topology structure, its classification performance is limited and the method still relies on label information to generate reasonable predictions. In [31], the authors adopt multi-view based SSL to learn representations of neuron morphology. Though the work mainly focuses on the local components of neuron branch which is not available in whole neuron reconstruction data, its promising results advocate the possibility of learning morphology representation without labels.

In this paper, we propose a new SSL framework to directly learn neuron morphological features from 3D dendrite reconstructions without the need of any prior knowledge of a neuron such as cell type annotation. The framework is named *COntrastive Neuron MOrphology Representation learning* (CoNMoR). It first maps neuron trees input to a deep embedding space by our proposed neuron encoder, and then incorporates a SSL pipeline composed by neuron tree augmentations and a contrastive loss to learn representations. By testing the framework on 2403 neuron reconstructions from 3 different data resources, we are excited to find that CoNMoR can clearly delineate neuron morphology in the embedding space (Figure 1). Comparison with prior arts shows that CoNMoR can achieve premium results than supervised methods even *without annotations*.

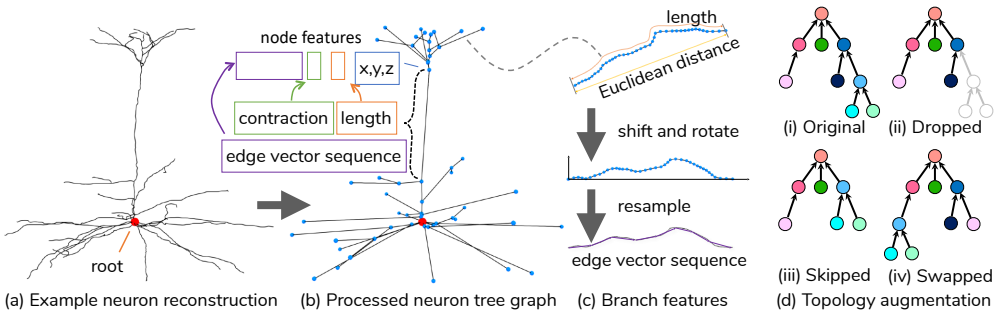

Figure 2: Illustration of (a) neuron reconstruction, (b) processed tree graph, (c) branch feature computation, (d) topology augmentations. (ii) to (iv) correspond to (9) to (11) in Table 1.

2 Background and Related Works

Neuron tree reconstruction and heuristic representation. Neuron reconstruction is usually represented by a tree structure in a "swc" format [3] as illustrated in Figure 2(a). The root of a tree is placed at the soma and the nodes (vertices) with 3D coordinate as well as other node properties are connected to the parent node. To quantitatively describe the neuron morphology, a set of heuristic features and related tools namely L-measure [32] has been proposed and widely used. These features can be summarized as 4 categories including (1) distance-based measurements (*e.g.*, branch length, distance to soma); (2) angle-based measurements (*e.g.*, bifurcation angle); (3) topology-based measurements (*e.g.*, number of bifurcations, branch order); and (4) size-based measurements (*e.g.*, branch radius, surface area). Statistics of these features such as mean or standard deviation are usually adopted to describe neuron morphology. Such representation is easy to compute and explain. However, they cannot differentiate complex shapes (examples in Appendix).

Cell-type and neuron morphology classification. As the primary function of dendritic branches is to compartmentalize functionally related synaptic inputs, dendritic morphology can be used to categorize neuronal subtypes based on total dendritic length, amount of branching across the dendritic trees, and localization of dendritic branch fields to different brain regions (such as layers of neocortex). Most existing cell-type classification algorithms rely on hand-crafted features [32] to predict or cluster cell types [39, 23]. In some recent works, researchers started to adopt tree-based model [51, 22] and sequence-based model [21] to classify neuron morphology. These methods take raw neuron data as input and predict cell-types in an end-to-end fashion. Despite competing performance than classic methods, training of these models requires manual annotations of cell types. Though [21] can be pre-trained without labels, the performance is limited as the adopted seq2seq model ignores critical tree structure of neurons. [31] adopted a triplet-loss to learn consistent embeddings between different views of the same neuron branch (Neuron2vec). The learned embeddings can identify subcellular compartments and the cell types of neuron reconstructions. However, this method mainly focuses on local shape and organelle. Such information requires nano resolution imaging techniques that are not applicable for whole neuron analysis.

Contrastive learning. Contrastive learning (CL) is a SSL method that aims at bringing positive sample pairs closer and spreading away the negative sample pairs in the feature embedding space. This method dates back to a decade ago [13] or earlier, and has recently drawn significant interest within and across different communities [44, 24, 17, 47, 16, 6, 5]. Some research efforts have been made in graph contrastive learning (GCL) [36, 28, 15, 48, 45]. Particularly, GraphCL [48] studies the graph data augmentations systematically for GCL, and InfoGCL [45] provides a theoretical analysis for it. However, to our best knowledge, no work has studied the tree graph, a special instance of graph, within the context of CL. As we discuss later, tree graph is too special to share the same setups as general GCL. We subsequently propose several tree-graph-oriented data augmentations that are helpful for tree CL. We hope this work could provide a useful data point, a missing piece of jigsaws, for tree graph representation learning. In addition, we believe we are the very first to explore CL in neuron tree morphology representation learning, which we expect to be a new powerful technique for analyzing neurons both quantitatively and qualitatively.

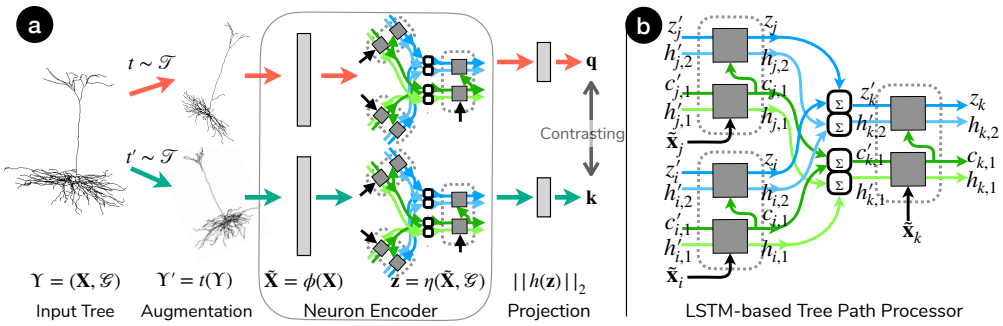

Figure 3: The overview of CoNMoR: (a) CoNMoR pipeline. (b) The zoom-in view of how the LSTM-based tree path processor η propagates information from child nodes to a parent node, taking a parent node k that has two child nodes i and j as an example. The symbol "Σ" denotes summation. The dashed rectangles mark our double-cell LSTM processor.

3 Methods

3.1 Tree graph representation of neuron

In this paper, we define a neuron tree as $\Upsilon \triangleq (\mathbf{X}, \mathcal{G})$, where $\mathcal{G}$ represents the tree topology and $\mathbf{X} = \{\mathbf{x}_k\}_{k=1}^{N}$ is an $N \times (3 + Q)$ matrix that is from N nodes with 3D coordinates and Q-dim additional features. Intuitively, the reconstruction of a neuron tree can be directly used as model input. However, most of the vertices in raw neuron reconstructions are 1-degree nodes along the branch which are less informative but overwhelm bifurcation nodes. Also, the number of such 1-degree nodes is biased to neuron reconstruction pipeline and raw data resolution. We define a complete sequence of 1-degree nodes between two bifurcation nodes or a bifurcation node and a tip node as a *branch*. To normalize neuron tree graphs across data source, we replace branches with edges directly connecting the root node, the bifurcation nodes, and the tip nodes in the graph (Figure 2(b)). To preserve the branch shape information, we calculate related features and concatenate them with the coordinate property of its downstream node (bifurcation node or tip node). This procedure also reduces the complexity of neuron tree graphs and benefits training efficiency.

For branch features, length and contraction have been computed. Branch length L is the accumulated distance between adjacent vertices along a branch: $L = \sum_i ||V_i - V_{i+1}||_2$. Branch contraction is the ratio between the Euclidean distance $D = ||V_{start} - V_{end}||_2$ and the length L of a branch: $C = D/L$. Length measures the size of a branch while contraction measures the curviness. Notably, both measurements are invariant to the order of edges and thus cannot differentiate certain patterns (details in Appendix). To allow a comprehensive encoding of branch pattern, we resample each branch with a fixed number of vertices (8 in this paper) and take the sequence of the edge vectors $\vec{E}_i = V_i - V_{i+1}$ as the its shape feature: $F_E \in \mathbb{R}^{24}$. In addition, the branch orientation is unified before extracting feature vectors such that the start and end points are on the x-axis and the principle direction is aligned with y-axis (Figure 2(c)). Together, there are 29 features for each neuron node: $\mathbf{x} = [x, y, z; L; C; F_E]$.

3.2 Neuron Encoder

We next propose a neuron encoder f to project the neuron tree graph into a feature space. Our model has two components: a trainable MLP (multi-layer perceptron) embedding layer (ϕ) that maps point coordinates and other node-related features into the embedding space, and a LSTM-based tree-path processor (η) that exploits tree topology information. Formally, we define the neuron encoder as $f \triangleq \eta \circ \phi$. Figure 3 (a) shows an overview of our neuron encoder.

Embedding Layer (ϕ). The embedding layer maps input features vector $\mathbf{x}$ into d-dim embedding vectors $\tilde{\mathbf{x}}$. It is implemented by a three-layer perceptron with `BatchNorm` and `ReLU` attached to the input and hidden layers. We refer to the outputs of this embedding layer as node embeddings. Formally, we compute $\tilde{\mathbf{x}}_k = \phi(\mathbf{x}_k)$, where $\mathbf{x}_k$ is the k-th node of a neuron, or jointly as $\tilde{\mathbf{X}} = \phi(\mathbf{X})$.

LSTM-based tree-path processor (η). The original LSTM [18] architecture applies for strictly sequential data. However, a neuron tree structure is more complicated than a simple sequence. Therefore, we adopt a child-sum Tree-LSTM architecture, similar to [37], to process such data. Recall the hidden state and the cell state at time t are usually denoted as $\mathbf{h}_t$ and $\mathbf{c}_t$ in LSTMs [18]. Here we denote them similarly, but indexed by tree node, as $\mathbf{h}_k$ and $\mathbf{c}_k$ at node k. To aggregate information from multiple child nodes, we simply sum over the input states. The transition equations can be formulated as followings:

$$\mathbf{h}'_{k,1} = \sum_{j\in \mathrm{C}(k)} \mathbf{h}_{j,1}, \quad \mathbf{c}'_{k,1} = \sum_{j\in \mathrm{C}(k)} \mathbf{c}_{j,1} \tag{1}$$

$$\mathbf{h}_{k,1}, \mathbf{c}_{k,1} = \mathrm{LSTM}_1(\tilde{\mathbf{x}}_k, (\mathbf{h}'_{k,1}, \mathbf{c}'_{k,1}))) \tag{2}$$

$$\mathbf{h}'_{k,2} = \sum_{j\in \mathrm{C}(k)} \mathbf{h}_{j,2}, \quad \mathbf{z}'_{k} = \sum_{j\in \mathrm{C}(k)} \mathbf{z}_{j} \tag{3}$$

$$\mathbf{h}_{k,2}, \mathbf{z}_k = \mathrm{LSTM}_2(\mathbf{c}_{k,1}, (\mathbf{h}'_{k,2}, \mathbf{z}'_k)) \tag{4}$$

where $\mathrm{C}(k)$ denotes the set of child nodes of node k, and $\mathbf{z}_k$ is the final cell state. We use two layers of LSTM cells (non-shared weights) to better capture the topology information (Eqns. 2, 4). Our tree-path processor is designed in a bottom-up manner where the information is aggregated from each individual leaf nodes and flowed to the root node. We then use the final cell state of the root node $\mathbf{z}$, referred to as *neuron representation*, for downstream tasks, *e.g.*, predicting class labels by an additional single linear layer classifier or projected by a projection head [5, 16] for contrastive learning. Formally, we compute the representation of a neuron tree by $\mathbf{z} = \eta(\tilde{\mathbf{X}}, \mathcal{G})$ (the propagation path is determined by $\mathcal{G}$). Figure 3 (b) demonstrates how information is propagated along the path.

3.3 Contrastive learning of tree graph

Recent progress in self-supervised contrastive learning, stemmed from visual representation learning [44, 16, 6, 5, 12], has been seen in many other fields, including graph contrastive learning [45, 48] — an overall under-explored direction. Attempts have been made to study self-supervised graph representation learning in terms of data augmentation [48], network architecture, and overall pipelines [45]. Although the neuron tree is a special case of graph structure, we argue that it is too special to share the same setups as traditional graphs in terms of contrastive learning [48]. Below we first introduce our CoNMoR pipeline, then elaborate on the subtle but important difference between a neuron tree and a general graph, subsequently introducing our data augmentations for tree structure.

Tree Contrastive Learning Figure 3 illustrates the overall pipeline of CoNMoR. The core idea of contrastive learning is to maximize the similarity between positive sample pairs and minimize it among negative sample pairs by optimizing a contrastive loss function [44, 5, 13]. We follow a successful exemplar, MoCo [16, 6], in visual contrastive learning to build our tree contrastive learning framework for self-supervised pre-training.

Conceptually, MoCo [16] abstracts contrastive learning as a dictionary look-up problem. For each encoded query $\mathbf{q}$, there is a set of encoded keys $\{\mathbf{k}_0, \mathbf{k}_1, \mathbf{k}_2, ...\}$ in a dictionary. The task for a model is to pull closer $\mathbf{q}$ and its matched positive key $\mathbf{k}^+$ in the dictionary while to spread $\mathbf{q}$ away from all other negative keys $\mathbf{k}^-$. When using the dot-product as similarity measurement, a form of contrastive loss function based on InfoNCE [44] becomes:

$$L_q = -\log \frac{\exp(\mathbf{q}\cdot\mathbf{k}^+/\tau)}{\exp(\mathbf{q}\cdot\mathbf{k}^+/\tau) + \sum_{\mathbf{k}^-} \exp(\mathbf{q}\cdot\mathbf{k}^-/\tau)} \tag{5}$$

where τ is a temperature hyper-parameter [44]. This loss is defined for one input instance, and the total loss is averaged over a batch.

To construct a pair of query and key, given a set of trees $\mathcal{D}$ and a set of possible augmentations $\mathcal{T}$, we first uniformly sample a tree $\Upsilon \sim \mathcal{D}$ and transform it into two augmented views $\Upsilon_q = t(\Upsilon)$ and $\Upsilon_k = t'(\Upsilon)$, referred to as the query view and the key view, by applying two sampled augmentations $t \sim \mathcal{T}, t' \sim \mathcal{T}$. Then, we compute their neuron representations using a neuron encoder defined in Sec. 3.2, *i.e.*, $\mathbf{z}_q = f(\Upsilon_q), \mathbf{z}_k = f'(\Upsilon_k)$. Following the improvement in SimCLR [6], we further map the neuron representations to another latent space using a non-linear projection head $h(\cdot)$ and normalize them by $\mathbf{q} = ||h(\mathbf{z}_q)||_2, \mathbf{k} = ||h'(\mathbf{z}_k)||_2$. Here, f' and h' stand for the momentum copies of f and h, named momentum encoder [16]. Therefore, $\mathbf{q}$ and $\mathbf{k}^+$ in Equation 5 are naturally generated from a single sample, and $\mathbf{k}^-$'s are from a queue [16] that stores features from previous batches.

Discussion: data augmentation on graphs, trees, and neuron morphology. Data augmentation is known to be essential for self-supervised contrastive learning, where its absence would affect the quality of learned representations significantly [5]. Current successful data augmentations for graphs [48] mainly include (1) "node dropping" that randomly drops graph nodes and their associated edges, (2) "edge perturbation" that randomly adds or drops edges, (3) "attribute masking" that acts like dropout [35] to node features, and (4) "subgraph sampling" that uses random walk to sample a partial graph. These augmentations can apply to general graph data without loss of generality. However, some of them are unsuitable for tree-structure graphs without special consideration. For example, the node dropping at a rate of 0.2 (default in [48]) would produce a significantly pruned tree since all the subtrees would be removed once their parent node was dropped. Besides, edge perturbation is unlikely to be feasible in a tree structure if no strategical design is considered. Two nodes in a neuron tree may be spatially close to each other while having long "distance" in terms of topology (*e.g.*, number of hops from one node to another) and vice versa. The perturbed edges would create cyclic graphs on the one hand, and break the consistency between space (3d coordinates) and structure (nodes connectivity) on the other hand. Therefore, we argue the special consideration is needed for the design of data augmentations on tree-structured data.

Data augmentation for neuron morphology representation learning. Regarding the characteristics of neuron data, we propose 11 available augmentations for neuron morphology representation learning, which are briefly summarized in Table 1 and some implementations are provided in Appendix. The rationale behind our design is that neuron representation should be invariant to perturbations — the golden rule applicable in many cases. For neuron data, such perturbations may occur in (i) node coordinates, (ii) overall morphology, (iii) attribute masking and (iv) topology. Thus we can augment neuron data from different perspectives as Table 1 shows. Jittering (4, 6) means adding random noise in a small range; deforming (7) means scaling each branch independently. We also include attribute masking (8) used in general graph augmentation. For (iv) topology, we propose three different augmentations. The first is (9) "truncated dropping subtrees" that recursively visits nodes from the root to leaves and drops a subtree with a small probability (*e.g.*, 0.05). The name "truncated" means at most m trees will be dropped. Similarly, we have (10) "skipping parent node" that at most m parent nodes will be replaced by one of their subtrees (respectively), and (11) "swapping sibling child" that picks a subtree from a sibling node and swaps that subtree with another subtree in the current node; all pickings are random. We empirically set $m = 10$ in this paper. Figure 2 (d) illustrates how topology augmentation acts like.

Table 1: Overview of data augmentations for neuron morphology representation learning.

Category	Target	Augmentations
(i) Point Transformation	Coordinates	1) Scaling, 2) Rotating, 3) Shifting, 4) Jittering, 5) Flipping
(ii) Morphology	Branch feats	6) Jittering branches, 7) Deforming branches
(iii) Attribute Masking	Branch feats	8) Masking attributes
(iv) Topology	Edges	9) Truncated dropping subtrees 10) Truncated skipping parent node 11) Truncated swapping sibling child

4 Experiments

4.1 Experiment setups

Data. We download data from 3 public resources: (1) BICCN fMOST data from Brain Image Library (BIL, `https://download.brainimagelibrary.org/biccn/zeng/luo/fMOST/`) [26], (2) Janelia MouseLight (JML, `http://mouselight.janelia.org/`) [43], and (3) Allen Cell Types (ACT, `https://celltypes.brain-map.org/`) [1]. In total, 3358 three-dimensional mouse brain neuron reconstructions are collected (refer to Table 2 for details). To facilitate analysis, we exclude cell types with less than 50 reconstructions, yielding 2403 neurons in total. Some of the samples include both dendrite and axon reconstructions. For simplicity, we only keep dendritic arbors in this paper and leave the analysis of axons for future works. All three datasets have registered reconstructions to the 3D Allen mouse brain Common Coordinate Framework (CCFv3 [40]). The

scales of the reconstructions are the same. Meanwhile, the absolute coordinates and orientations of registered neurons could also carry cell-type information that are related to brain anatomy instead of neuron morphology. To eliminate such potential information leakage, all dendrites are shifted and rotated such that the root is placed at the origin and the X/Y/Z coordinates are determined based on principle component analysis (PCA). Data are then processed and analyzed following the steps introduced in the Section 3.1. In the interest of evaluating CoNMoR's capability of classifying known cell types across data sources, we specifically use the cell types with rich coverage among 3 datasets only (grayed cell types in Table 2). We denote these sub-datasets as BIL-6, JML-4 and ACT respectively. For supervised training, 80% samples in each dataset are randomly picked as training set and 20% are reserved for testing. For CoNMoR pre-training, we use all samples except those from class "others", and evaluate performance in BIL-6, JML-4 and ACT datasets.

Table 2: Demographic table of reconstructions. Refer to Appendix for full names of brain regions.

Dataset	**Number of neuron reconstructions in each brain regions**											
	Total	**L2/3**	**L4**	**L5**	**L6**	**VPM**	**CP**	**VPL**	**SUB**	**PRE**	**MG**	**Others**
BIL [26]	1741	126	56	315	93	378	311	80	4	-	50	328
JML [43]	1107	64	-	179	114	12	-	5	70	60	-	602
ACT [1]	510	111	123	164	97	-	-	-	-	-	-	15
Total	3358	301	179	658	304	390	311	85	74	60	50	945

Implementations and evaluation. For implementation, we set the layer sizes of embedding MLP to 256 and 128, and the cell state size of LSTM-processor to 128. For CoNMoR, we use a queue length of 1024, a momentum of 0.99 to update momentum encoder and a temperature $\tau = 0.1$ for InfoNCE loss. For both supervised and self-supervised pre-training, we use a batch size of 128 and train for 100 epochs. For data augmentation, we use all the listed augmentations by default. More implementation details and augmentation details such as probability and hyper-parameters are provided in Appendix. For evaluation, we report the sample-wise accuracy. During pre-training, we evaluate the K Nearest-Neighbor (KNN) accuracy that uses training samples to vote for test samples. We set the number of neighbors to 20 for BIL-6 and ACT datasets, and 5 for JML-4 datasets since a minority class might only have 12 samples in total in JML-4.

4.2 CoNMoR delineates neuron morphology in embedding space

In this subsection, we demonstrate that CoNMoR can learn neuron tree representations that manifest different characteristics of neuron morphology *without any human labels*. To that end, we first follow procedures described in Section 3.3 to conduct CoNMoR pre-training and use the pre-trained neuron encoder to extract features from all three datasets (BIL, JML and ACT) and project them into a 2D space by t-SNE [38]. Figure 1 and 4 shows the thumbnails of all neurons placed at their corresponding 2D t-SNE coordinates.

Visual inspection of cell types delineation. By visual inspection, it is evidenced that CoNMoR can separate major cell-types such as pyramidal neuron (Figure 1(b-c)) and interneuron (Figure 1(d-e)) in the embedding space. Pyramidal neurons are the primary long-range output neurons of the brain [34] while interneurons mainly sculpt local network dynamics through a variety of circuit motifs. The asymmetric long projecting apical tuft is a unique feature that differentiates pyramidal neuron's morphology with interneuron's. CoNMoR can further delineate subtypes of pyramidal neuron and interneurons. In Figure 1 and 4(a), we can clearly see the separation of layer 2/3 (Figure 1(c)) and layer 5 (Figure 1(b)) pyramidal neurons. Layer 2/3 neurons can be distinguished by their shorter apical dendrites with extensive oblique branching [42] while layer 5 neurons feature extended apical dendrites with extensive horizontal branching in layer 1 and sparse oblique branching in layer 2/3 [14]. The two primary classes of inhibitory interneurons (parvalbumin-expressing basket cells and somatostatin-expressing Martinotti cells) are also clearly distinguished with CoNMoR. Parvalbumin-expressing basket cells have extensive branching field that extends radially around the soma (Figure 1(d)) [27]. In contrast, somatostatin-expressing Martinotti cells tend to have unipolar or bipolar dendritic branching that extends vertically above and below the soma (Figure 1(e)) [33]. It should be noted that such delineation is purely self-driven without any human labels or any prior knowledge of cell types.

Table 3: Comparisons to state-of-the-arts neuron cell-type classification. We report the sample-wise accuracy (%). The numbers of classes in BIL-6, JML-4 and ACT are 6, 4, and 4, respectively. Our method performs the best in all three datasets and can outperform current methods even without labels exposed during training (JML-4, ACT).

Methods	BIL-6	JML-4	ACT
MorphVAE [21]	66.80	40.00	41.05
TRNN [51]	80.86	54.29	43.16
CoNMoR frozen KNN	79.30	62.86	54.74
CoNMoR fine-tuining	**85.16**	**65.71**	**65.26**

Table 4: Ablation on augmentations.

	None	Point	Mask	Morpho	Drop	Skip	Swap
None	66.61	77.34	67.97	70.12	71.09	69.34	68.75
Point	77.34	78.12	79.69	77.54	71.88	76.37	79.5
Mask	67.97	79.69	68.17	71.88	70.51	71.49	67.97
Morpho	70.12	77.54	71.88	69.34	73.64	69.73	71.29
Drop	71.09	71.88	70.51	73.64	71.1	73.25	71.09
Skip	69.34	76.37	71.49	69.73	73.25	75.0	71.88
Swap	68.75	79.5	67.97	71.29	71.09	71.88	67.78

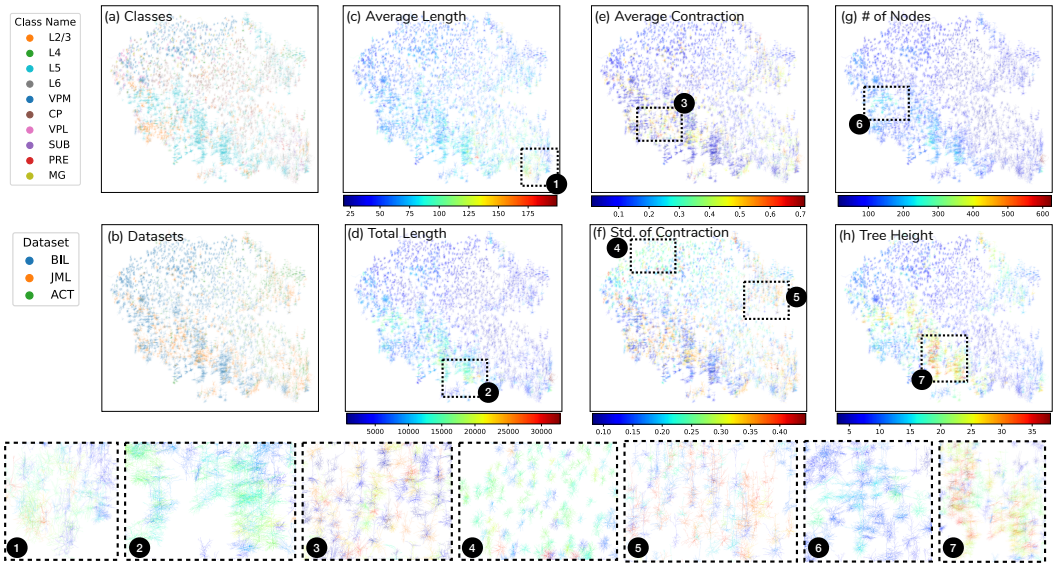

Figure 4: Thumbnails of neurons in the same embedding space of Figure 1 colored by different attributes. It is evidenced CoNMoR encodes cell types (a) and neuron morphology features (c)-(h), *without human labels*. Refer to Appendix for high-resolution figures.

Quantitative evaluation. Here we quantitatively evaluate CoNMoR by a pre-training & fine-tuning protocol [5] in classifying the known cell-types. For this setting, we pre-train a neuron encoder in joint dataset (merging three datasets) for 100 epochs and fine-tune it on each individual datasets. Table 3 report the results, with the best KNN evaluation accuracy during pre-training included. It can be observed that CoNMoR trained without label information can already achieve better performance than supervised MorphVAE [21] and TRNN [51]. Fine-tuning the pre-trained neuron encoder further improves cell-type classification. This is mainly because MorphVAE [21] is based on random walk to capture topology information, which ignores tree architecture. While TRNN [51] has less description power than our neuron encoder in capturing high-order interaction between features.

CoNMoR manifests neuron morphology distribution. Figure 4 shows the t-SNE plots of neuron representations. Each plot is colored by a different morphology feature, *i.e.*, average/total branch length (c, d), average / standard deviation of contraction (e, f), number of nodes in a neuron (g) and the height of a neuron tree (h). Neurons that share similar morphology features are well clustered in the embedding space (see (1) to (7)), demonstrating the good representations learnt by CoNMoR. These morphology attributes have not been explicitly provided as input features but it turns out that CoNMoR can mine such information by itself, *without human labels*. In addition, we can see it by cross-referencing Figure 4 (a) and (b) that neurons of a same class but from different datasets are clustered together in CoNMoR's representation space, showing CoNMoR's unbiasedness to data sources — a much desirable property for neuron morphology representation learning.

4.3 Ablation on augmentation methods and network architectures

In this subsection, we first assess the importance of data augmentations in contrastive neuron morphology learning as well as to understand how different augmentations contribute to it. Then we ablate on several design choices in our neuron encoder.

To systematically study the effects of different categories of data augmentations, we first group the proposed augmentations in Table 1 by their categories except the topology augmentation, the one we are more interested in. Then, we apply them individually or in pairs to CoNMoR framework when pre-training. For ablation purpose, in this experiment we train on the joint dataset for 50 epochs. Table 4 shows the KNN evaluation results under each individual category or a composition of data augmentations. Here we do not specify the order of augmentations, thus resulting a symmetric performance matrix. Unlike the observation in visual or graph representation contrastive learning that no single augmentation performs the best [5, 48], we find in neuron morphology representation learning certain kinds of augmentations are good enough for pre-training (*e.g.*, point transformation, skipping parent node). Besides, compared to Node Dropping [48], even our "truncated drop subtrees" applies under a significantly smaller probability (0.05 v.s. 0.2 in [48]) and a truncation constraint, the resulted performance (drop) is sometimes unsatisfactory, showing the importance of data augmentation designs with respect to tree structure data.

Then, we study the effects of different components in our neuron encoder. For this experiment, we train neuron encoders for 100 epochs on three datasets and report their best test performance. Table 5 (left) shows the results, where a steady increasing trend of model performance can be observed. We also include the KNN accuracy (Tab. 5 right) for 50-epoch pre-training on the joint dataset and evaluated in BIL-6 dataset just for ablation interest. Here we find it interesting that `BN` becomes more important in contrastive self-supervised learning than in supervised learning. No `BN` in the neuron encoder would give significantly worse results (first line). Further adding `BN` in the projection head is helpful for CoNMoR. This behavior has also been reported in non-contrastive SSL methods before, *e.g.*, BYOL [12, 29], where removing `BN` causes model collapse.

Table 5: **Ablation on model designs. (left)** Classification accuracy of supervised training for 100-epoch. **(right)** KNN accuracy of CoNMoR pre-training for 50 epochs and evaluated in BIL-6 dataset. BN matters more in CoNMoR. Gray rows indicate our default setting.

Settings			Classification Accuracy (%)		
BN	MLP	# LSTM Cells	BIL-6	JML-4	ACT
-	-	1	80.86	54.29	43.16
-	✓	1	83.98	57.14	50.53
✓	✓	1	84.77	58.57	**63.16**
✓	✓	2	**85.55**	**64.29**	62.11

Encoder BN	Projector BN	KNN accuracy (%) BIL-6
-	-	69.14
✓	-	75.00
✓	✓	**76.56**

5 Conclusion

In this work, we propose CoNMoR — a new SSL framework for tree structure representation learning. To our best knowledge, CoNMoR is the very first to explore learning the representation of neuron tree morphology with contrastive learning. Our results on 2403 neuron reconstructions show that CoNMoR has the unique potential to identify unbiased cell-types. Modern cell-type classification is heavily dependent on transcriptomic data [49, 8]. However, even neurons within the same mouse possessing identical genetic identity and birth-date demonstrate extremely variable dendritic morphology and synaptic distribution [19]. This variability underlies the need for high-throughput neuronal morphology classification techniques to probe for the existence of novel subtypes within classically established neuronal identities [4]. CoNMoR has a great potential to fill this need and to explore dendritic architecture not only in space, but also in time: especially to identify stages of development and disease progression [30]. These findings could aid in the development of treatments for neurological disorders characterized by progressive dendrite degeneration such as Alzheimer's disease and depression [10, 7]. Taken together, CoNMoR has immense potential to benefit human health by providing more rigorous output measurements for drug discovery experiments and possibly even diagnostic tools for human patients as noninvasive brain imaging continues to improve.

References

[1] Allen cell types database. allen institute for brain science. http://celltypes.brain-map.org/data.

[2] Yoshua Bengio, Yann Lecun, and Geoffrey Hinton. Deep learning for ai. *Communications of the ACM*, 64(7):58–65, 2021.

[3] Robert C Cannon, DA Turner, GK Pyapali, and HV Wheal. An on-line archive of reconstructed hippocampal neurons. *Journal of neuroscience methods*, 84(1-2):49–54, 1998.

[4] Mark S Cembrowski and Nelson Spruston. Heterogeneity within classical cell types is the rule: lessons from hippocampal pyramidal neurons. *Nature Reviews Neuroscience*, 20(4):193–204, 2019.

[5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.

[6] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.

[7] David Cotter, Daniel Mackay, Sabine Landau, Robert Kerwin, and Ian Everall. Reduced glial cell density and neuronal size in the anterior cingulate cortex in major depressive disorder. *Archives of general psychiatry*, 58(6):545–553, 2001.

[8] Javier DeFelipe, Pedro L López-Cruz, Ruth Benavides-Piccione, Concha Bielza, Pedro Larrañaga, Stewart Anderson, Andreas Burkhalter, Bruno Cauli, Alfonso Fairén, Dirk Feldmeyer, et al. New insights into the classification and nomenclature of cortical gabaergic interneurons. *Nature Reviews Neuroscience*, 14(3):202–216, 2013.

[9] Michael N Economo, Nathan G Clack, Luke D Lavis, Charles R Gerfen, Karel Svoboda, Eugene W Myers, and Jayaram Chandrashekar. A platform for brain-wide imaging and reconstruction of individual neurons. *Elife*, 5:e10566, 2016.

[10] Eric Falke, Jonathan Nissanov, Thomas W Mitchell, David A Bennett, John Q Trojanowski, and Steven E Arnold. Subicular dendritic arborization in alzheimer's disease correlates with neurofibrillary tangle density. *The American journal of pathology*, 163(4):1615–1621, 2003.

[11] Hui Gong, Dongli Xu, Jing Yuan, Xiangning Li, Congdi Guo, Jie Peng, Yuxin Li, Lindsay A Schwarz, Anan Li, Bihe Hu, et al. High-throughput dual-colour precision imaging for brain-wide connectome with cytoarchitectonic landmarks at the cellular level. *Nature communications*, 7(1):1–12, 2016.

[12] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33:21271–21284, 2020.

[13] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE, 2006.

[14] Kenneth D Harris and Thomas D Mrsic-Flogel. Cortical connectivity and sensory coding. *Nature*, 503(7474):51–58, 2013.

[15] Kaveh Hassani and Amir Hosein Khasahmadi. Contrastive multi-view representation learning on graphs. In *International Conference on Machine Learning*, pages 4116–4126. PMLR, 2020.

[16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020.

[17] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, 2018.

[18] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[19] Daniel Maxim Iascone, Yujie Li, Uygar Sümbül, Michael Doron, Hanbo Chen, Valentine Andreu, Finola Goudy, Heike Blockus, Larry F Abbott, Idan Segev, et al. Whole-neuron synaptic mapping reveals spatially precise excitatory/inhibitory balance limiting dendritic and somatic spiking. *Neuron*, 106(4):566–578, 2020.

[20] Lida Kanari, Paweł Dłotko, Martina Scolamiero, Ran Levi, Julian Shillcock, Kathryn Hess, and Henry Markram. A topological representation of branching neuronal morphologies. *Neuroinformatics*, 16(1):3–13, 2018.

[21] Sophie C. Laturnus and Philipp Berens. Morphvae: Generating neural morphologies from 3d-walks using a variational autoencoder with spherical latent space. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 6021–6031. PMLR, 18–24 Jul 2021.

[22] Congping Lin, Yuanfei Huang, Tingwei Quan, and Yiwei Zhang. Modelling brain-wide neuronal morphology via rooted cayley trees. *Scientific reports*, 8(1):1–10, 2018.

[23] Bojan Mihaljević, Pedro Larrañaga, Ruth Benavides-Piccione, Sean Hill, Javier DeFelipe, and Concha Bielza. Towards a supervised classification of neocortical interneuron morphologies. *BMC bioinformatics*, 19(1):1–22, 2018.

[24] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

[25] Hanchuan Peng, Jianyong Tang, Hang Xiao, Alessandro Bria, Jianlong Zhou, Victoria Butler, Zhi Zhou, Paloma T Gonzalez-Bellido, Seung W Oh, Jichao Chen, et al. Virtual finger boosts three-dimensional imaging and microsurgery as well as terabyte volume image visualization and analysis. *Nature communications*, 5(1):1–13, 2014.

[26] Hanchuan Peng, Peng Xie, Lijuan Liu, Xiuli Kuang, Yimin Wang, Lei Qu, Hui Gong, Shengdian Jiang, Anan Li, Zongcai Ruan, et al. Morphological diversity of single neurons in molecularly defined cell types. *Nature*, 598(7879):174–181, 2021.

[27] Frédéric Pouille and Massimo Scanziani. Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science*, 293(5532):1159–1163, 2001.

[28] Jiezhong Qiu, Qibin Chen, Yuxiao Dong, Jing Zhang, Hongxia Yang, Ming Ding, Kuansan Wang, and Jie Tang. Gcc: Graph contrastive coding for graph neural network pre-training. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1150–1160, 2020.

[29] Pierre H Richemond, Jean-Bastien Grill, Florent Altché, Corentin Tallec, Florian Strub, Andrew Brock, Samuel Smith, Soham De, Razvan Pascanu, Bilal Piot, et al. Byol works even without batch statistics. *arXiv preprint arXiv:2010.10241*, 2020.

[30] Sandrine Romand, Yun Wang, Maria Toledo-Rodriguez, and Henry Markram. Morphological development of thick-tufted layer v pyramidal cells in the rat somatosensory cortex. *Frontiers in neuroanatomy*, 5:5, 2011.

[31] Philipp J Schubert, Sven Dorkenwald, Michał Januszewski, Viren Jain, and Joergen Kornfeld. Learning cellular morphology with neural networks. *Nature communications*, 10(1):1–12, 2019.

[32] Ruggero Scorcioni, Sridevi Polavaram, and Giorgio A Ascoli. L-measure: a web-accessible tool for the analysis, comparison and search of digital reconstructions of neuronal morphologies. *Nature protocols*, 3(5):866–876, 2008.

[33] Gilad Silberberg and Henry Markram. Disynaptic inhibition between neocortical pyramidal cells mediated by martinotti cells. *Neuron*, 53(5):735–746, 2007.

[34] Nelson Spruston. Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience*, 9(3):206–221, 2008.

[35] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.

[36] Fan-Yun Sun, Jordan Hoffmann, Vikas Verma, and Jian Tang. Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. *arXiv preprint arXiv:1908.01000*, 2019.

[37] Kai Sheng Tai, Richard Socher, and Christopher D Manning. Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*, 2015.

[38] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.

[39] Xavier Vasques, Laurent Vanel, Guillaume Villette, and Laura Cif. Morphological neuron classification using machine learning. *Frontiers in neuroanatomy*, 10:102, 2016.

[40] Quanxin Wang, Song-Lin Ding, Yang Li, Josh Royall, David Feng, Phil Lesnar, Nile Graddis, Maitham Naeemi, Benjamin Facer, Anh Ho, et al. The allen mouse brain common coordinate framework: a 3d reference atlas. *Cell*, 181(4):936–953, 2020.

[41] Yimin Wang, Qi Li, Lijuan Liu, Zhi Zhou, Zongcai Ruan, Lingsheng Kong, Yaoyao Li, Yun Wang, Ning Zhong, Renjie Chai, et al. Teravr empowers precise reconstruction of complete 3-d neuronal morphology in the whole brain. *Nature communications*, 10(1):1–9, 2019.

[42] Verena C Wimmer, Randy M Bruno, Christiaan PJ De Kock, Thomas Kuner, and Bert Sakmann. Dimensions of a projection column and architecture of vpm and pom axons in rat vibrissal cortex. *Cerebral cortex*, 20(10):2265–2276, 2010.

[43] Johan Winnubst, Erhan Bas, Tiago A Ferreira, Zhuhao Wu, Michael N Economo, Patrick Edson, Ben J Arthur, Christopher Bruns, Konrad Rokicki, David Schauder, et al. Reconstruction of 1,000 projection neurons reveals new cell types and organization of long-range connectivity in the mouse brain. *Cell*, 179(1):268–281, 2019.

[44] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018.

[45] Dongkuan Xu, Wei Cheng, Dongsheng Luo, Haifeng Chen, and Xiang Zhang. Infogcl: Information-aware graph contrastive learning. *Advances in Neural Information Processing Systems*, 34, 2021.

[46] Fang Xu, Yan Shen, Lufeng Ding, Chao-Yu Yang, Heng Tan, Hao Wang, Qingyuan Zhu, Rui Xu, Fengyi Wu, Yanyang Xiao, et al. High-throughput mapping of a whole rhesus monkey brain at micrometer resolution. *Nature Biotechnology*, 39(12):1521–1528, 2021.

[47] Mang Ye, Xu Zhang, Pong C Yuen, and Shih-Fu Chang. Unsupervised embedding learning via invariant and spreading instance feature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6210–6219, 2019.

[48] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. *Advances in Neural Information Processing Systems*, 33:5812–5823, 2020.

[49] Rafael Yuste, Michael Hawrylycz, Nadia Aalling, Argel Aguilar-Valles, Detlev Arendt, Ruben Armañanzas, Giorgio A Ascoli, Concha Bielza, Vahid Bokharaie, Tobias Borgtoft Bergmann, et al. A community-based transcriptomics classification and nomenclature of neocortical cell types. *Nature neuroscience*, 23(12):1456–1468, 2020.

[50] Hongkui Zeng and Joshua R Sanes. Neuronal cell-type classification: challenges, opportunities and the path forward. *Nature Reviews Neuroscience*, 18(9):530–546, 2017.

[51] Tielin Zhang, Yi Zeng, Yue Zhang, Xinhe Zhang, Mengting Shi, Likai Tang, Duzhen Zhang, and Bo Xu. Neuron type classification in rat brain based on integrative convolutional and tree-based recurrent neural networks. *Scientific reports*, 11(1):1–14, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [No]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]