

IN VIVO CELL-TYPE AND BRAIN REGION CLASSIFICATION VIA MULTIMODAL CONTRASTIVE LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Current electrophysiological approaches can track the activity of many neurons, yet it is usually unknown which cell-types or brain areas are being recorded without further molecular or histological analysis. Developing accurate and scalable algorithms for identifying the cell-type and brain region of recorded neurons is thus crucial for improving our understanding of neural computation. In this work, we develop a multimodal contrastive learning approach for neural data that can be fine-tuned for different downstream tasks, including inference of cell-type and brain location. We utilize this approach to jointly embed the activity autocorrelations and extracellular waveforms of individual neurons. We demonstrate that our embedding approach, **Neuronal Embeddings via Multimodal** contrastive learning (**NEMO**), paired with supervised fine-tuning, achieves state-of-the-art cell-type classification for an opto-tagged visual cortex dataset and brain region classification for the public International Brain Laboratory brain-wide map dataset. Our method represents a promising step towards accurate cell-type and brain region classification from electrophysiological recordings.

1 INTRODUCTION

High-density electrode arrays now allow for simultaneous extracellular recording from hundreds to thousands of neurons across interconnected brain regions (Jun et al., 2017; Steinmetz et al., 2021; Ye et al., 2023b; Trautmann et al., 2023). While much progress has been made on developing algorithms for tracking neural activity (Buccino et al., 2020; Magland et al., 2020; Boussard et al., 2023; Pachitariu et al., 2024), identifying the cell-types and brain regions being recorded is still an open problem. Brain region classification is particularly challenging given the diversity of cell types in each region (Yao et al., 2023) and how the locations of regions differ between animals. The development of an accurate method for extracting cell-type and brain region labels will provide insights into brain development, microcircuit function, and brain disorders (Zeng & Sanes, 2017).

Traditional approaches for electrophysiological cell-type classification utilize simple features of the extracellular action potential (EAP) such as its width or peak-to-trough amplitude (Mountcastle et al., 1969; Matthews & Lee, 1991; Nowak et al., 2003; Barthó et al., 2004; Vigneswaran et al., 2011) or features of neural activity, such as the inter-spike interval distribution (Latuske et al., 2015; Jouty et al., 2018). These simple features are interpretable and easy to visualize but lack discriminative power and robustness across different datasets (Weir et al., 2015; Gouwens et al., 2019). Current automated featurization methods for EAPs (Lee et al., 2021; Vishnubhotla et al., 2024) and neural activity (Schneider et al., 2023a) improve upon manual features but are limited to a single modality.

There has been a recent push to develop multimodal methods that can integrate information from both recorded EAPs and spiking activity. PhysMAP (Lee et al., 2024) is a UMAP-based (McInnes et al., 2018a) approach that can predict cell-types using multiple physiological modalities through a weighted nearest neighbor graph. Another recently introduced method utilizes variational autoencoders (VAEs) to embed each physiological modality separately and then combines these embeddings before classification (Beau et al., 2024). Although both methods show promising results, PhysMAP cannot be fine-tuned for downstream tasks as it is nondifferentiable, and the VAE-based method captures features that are important for reconstruction, not discrimination, impairing downstream performance (Guo et al., 2017). Neither approach has been used to classify brain regions.

054 In this work, we introduce a multimodal contrastive learning method for neurophysiological data,
055 Neuronal Embeddings via Multimodal Contrastive Learning (NEMO), which utilizes large amounts
056 of unlabeled paired data for pre-training and can be fine-tuned for different downstream tasks includ-
057 ing cell-type and brain region classification. We utilize a recently developed contrastive learning
058 framework (Radford et al., 2021) to jointly embed individual neurons’ activity autocorrelations and
059 average extracellular waveforms. The key assumption of our method is that information shared be-
060 tween different modalities captures intrinsic properties of a neuron that are predictive of its cell-type
061 and corresponding brain region. Jointly embedding different modalities into a shared latent space,
062 will capture this shared information while discarding modality-specific noise (Huang et al., 2024).
063 We evaluate NEMO on cell-type classification using an optotagged Neuropixels Ultra (NP Ultra)
064 recording from mouse visual cortex (Ye et al., 2023b) and on brain region classification using the In-
065 ternational Brain Laboratory (IBL) brain-wide map dataset (IBL et al., 2023). Across both datasets
066 and tasks, NEMO significantly outperforms current state-of-the-art unsupervised (PhysMAP and
067 VAEs) and supervised methods. These results highlight NEMO’s ability to extract informative rep-
068 resentations despite high variability in the data, and demonstrate that NEMO is a significant advance
069 towards accurate cell-type and brain region classification from electrophysiological recordings.

070 2 RELATED WORK

071 2.1 CONTRASTIVE LEARNING FOR NEURONAL DATASETS

072
073 The goal of contrastive learning is to find an embedding space where similar examples are close
074 together while dissimilar ones are well-separated (Le-Khac et al., 2020). Contrastive learning has
075 found success across a number of domains including language (Reimers & Gurevych, 2019), vision
076 (Chen et al., 2020), audio (Saeed et al., 2021), and multimodal learning (Radford et al., 2021; Tian
077 et al., 2020). There has been a surge in the development of contrastive learning methods for neural
078 data. Contrastive learning has been applied to neuronal morphological data (Chen et al., 2022), con-
079 nectomics data (Dorkenwald et al., 2023) and preprocessed spiking activity (Azabou et al., 2021;
080 Urzay et al., 2023; Schneider et al., 2023b; Antoniadis et al., 2023). In each of these applications,
081 associated downstream tasks such as 3D neuron reconstruction, cellular sub-compartment classifi-
082 cation, or behavior prediction have shown improvement using this contrastive paradigm. One con-
083 trastive method, CEED, has been applied to raw extracellular recordings to perform spike sorting
084 and cell-type classification. In contrast to NEMO, CEED is unimodal (it ignores neural activity) and
085 was never applied to optotagged data or brain region classification (Vishnubhotla et al., 2024).
086
087

088 2.2 CELL-TYPE CLASSIFICATION

089
090 The goal of cell-type classification is to assign neurons to distinct classes based on their morphol-
091 ogy, function, electrophysiological properties, and molecular markers (Masland, 2004). Current
092 transcriptomic (Tasic et al., 2018; Yao et al., 2021; 2023) and optical methods (Cardin et al., 2010;
093 Kravitz et al., 2013; Lee et al., 2020) are effective but require extensive sample preparation or spe-
094 cialized equipment, limiting their scalability and applicability for in-vivo studies (Lee et al., 2024).
095 Recently, calcium imaging has been utilized in conjunction with molecular cell-typing to identify
096 cell-types (Bugeon et al., 2022; Mi et al., 2023). This approach suffers from low temporal resolu-
097 tion and requires significant post hoc effort to collect the molecular imaging data and align it to the
098 calcium data (and therefore this approach can not be used in closed loop in vivo experiments).

099
100 A promising alternative is to use the electrophysiological properties of recorded neurons as they
101 capture some of the variability of the transcriptomic profile (Bomkamp et al., 2019). This approach
102 obviates the need for genetic manipulations to identify such types. Simple electrophysiological fea-
103 tures extracted from a neuron’s EAPs and spiking activity are typically utilized to identify putative
104 cell-types (Gouwens et al., 2019). Automated methods including EAP-specific methods (Lee et al.,
105 2021) and spiking activity-based methods (Schneider et al., 2023a) are an improvement in com-
106 parison to manual features. Most recently, multi-modal cell-type classification methods including
107 PhysMAP (Lee et al., 2024) and a VAE-based algorithm (Beau et al., 2024) have been introduced
which make use of multiple physiological modalities such as the EAP, activity, or peri-stimulus time
histogram (PSTH) of a recorded neuron.

108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161

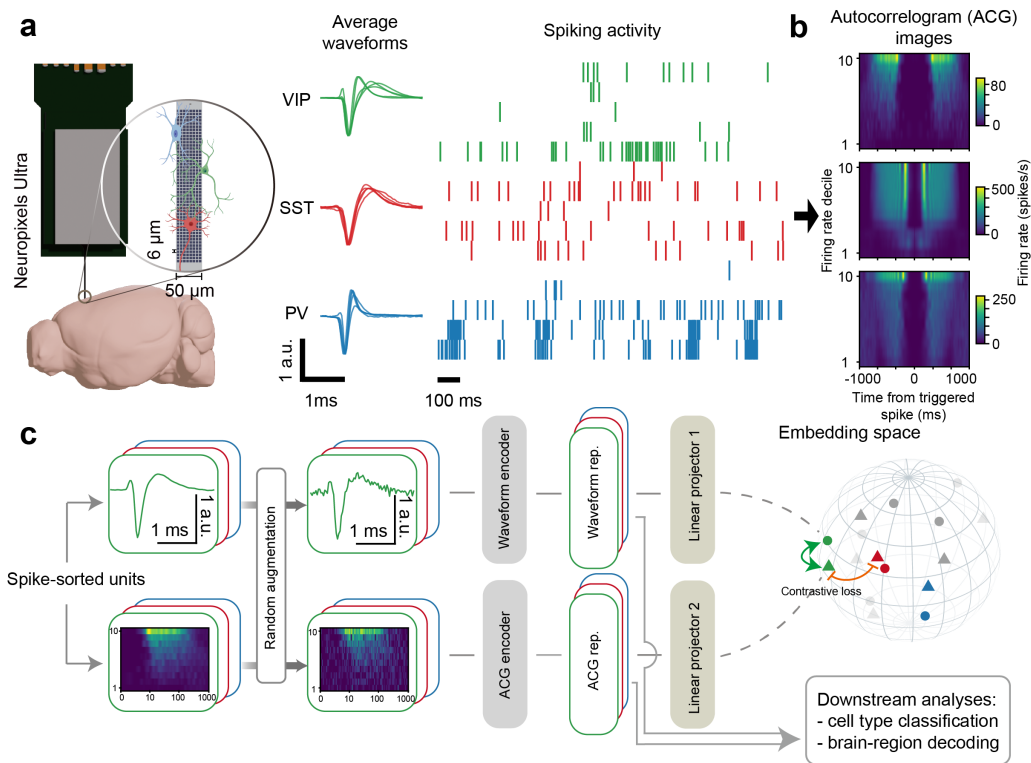


Figure 1: **Schematic illustration of NEMO.** (a) Neuropixels Ultra (Ye et al., 2023b) recordings capture activity from many different cell-types which have distinct extracellular action potentials (EAPs) and spiking activity. We present waveform and spiking activity snippets from five example neurons for each cell-type. (b) We transform the spiking activity of each neuron into a compact autocorrelogram (ACG) image from (Beau et al., 2024) that accounts for variations in the firing rate (see Section 4.1) (c) NEMO utilizes a CLIP-based learning strategy where an EAP encoder and a ACG image encoder are trained to learn an embedding which pushes together randomly augmented EAPs and ACG image from the same neuron while keeping different neurons separate. The learned representations can be utilized for downstream tasks (with or without fine-tuning) such as cell-type and brain-region classification. Schematic created using BioRender.com.

2.3 BRAIN REGION CLASSIFICATION

Brain region classification consists of predicting the location of a neuron or electrode based on the recorded physiological features (Steinmetz et al., 2018). Rather than predicting a 3D location, the task is to classify the brain region a neuron or electrode occupies, which can be estimated using post-insertion localization via histology (Sunken et al., 2012). Brain region classification is an important task for understanding fundamental differences in physiology between brain areas as well as for targeting regions that are hard to hit via insertion. Most importantly, brain region classification can provide a real-time estimate of the probe’s location in the brain during experiments, significantly increasing the success rate of hitting the target location. Additionally, insertions in primates and human subjects often lack histological information, instead relying on the experimental heuristics that lack standardization between laboratories. As this task is relatively new, only simple features of the EAPs have been utilized for classification (Jia et al., 2019; Tolossa et al., 2024). To our knowledge, no automated feature extraction methods have been utilized for brain region classification.

3 DATASETS

For cell-type classification, we utilize an opto-tagged dataset gathered with Neuropixels Ultra probes (NP Ultra; Ye et al., 2023b). For brain region classification, we use a dataset gathered by the IBL: a brain-wide map of neural activity of mice performing a decision-making task (IBL et al., 2023).

NP Ultra opto-tagged mouse data. This dataset consists of NP Ultra recordings of spontaneous activity from the visual cortex of mice. Opto-tagging is utilized to label 477 ground-truth neurons with three distinct cell-types. The ground-truth neurons are composed of 243 parvalbumin (PV), 116 somatostatin (SST), and 118 vasoactive intestinal peptide cells (VIP). There are also 8699 unlabelled neurons that we can utilize for pretraining.

IBL Brain-wide Map. This dataset consists of Neuropixels recordings from animals performing a decision-making task. Each neuron is annotated with the brain region where it is located. We utilize 675 insertions from 444 animals (37017 units). Each brain is parcellated with 10 brain atlas annotations, dividing the atlas into 10 broad areas: isocortex, olfactory areas (OLF), cortical subplate (CTXsp), cerebral nuclei (CNU), thalamus (TH), hypothalamus (HY), midbrain (MB), hindbrain (HB), cerebellum (CB) and hippocampal formation (HPF). We divide this dataset into a training, validation, and testing set by insertion such that we can evaluate each model on heldout insertions.

4 NEMO

We introduce Neuronal Embeddings via Multimodal contrastive learning (NEMO) which learns a multimodal embedding of neurophysiological data. To extract representations from multiple modalities, we utilize Contrastive Language-Image Pretraining (CLIP; Radford et al., 2021). CLIP uses a contrastive objective to learn a joint representation of images and captions. For NEMO, we utilize the same objective but with modality-specific data augmentations and encoders (see Figure 1c).

4.1 PREPROCESSING

We construct a paired dataset of spiking activity and EAPs for all recorded neurons. Using the open-source Python package NeuroPyxels (Beau et al., 2021), we computed an autocorrelogram (ACG) image for each neuron by smoothing the spiking activity with a 250-ms width boxcar filter, dividing the firing rate distribution into 10 deciles, and then building ACGs for each decile (see Figure 1b). This ACG image is a useful representation because the activity autocorrelations of a neuron can change as a function of its firing rate. By computing ACGs for each firing rate decile, the ACG image will account for firing rate dependent variations in the autocorrelations, allowing for comparisons between different areas of the brain, behavioral contexts, and animals (Beau et al., 2024). For the EAPs, we construct a ‘template’ waveform which is the mean of approximately 500 recorded waveforms for that neuron. For all experiments in the main text, we restrict the template to one channel with maximal amplitude. For multi-channel template results, see Supplement F.

4.2 DATA AUGMENTATIONS

Previous work on contrastive learning for spiking activity utilizes data augmentations including sparse multiplicative noise (pepper noise), Gaussian noise, and temporal jitter (Azabou et al., 2021). As it is computationally expensive to construct ACG images for each batch during training, we instead design augmentations directly for the ACG images rather than the original spiking data. Our augmentations include temporal Gaussian smoothing, temporal jitter, amplitude scaling, additive Gaussian noise, and multiplicative pepper noise (see Supplemental B for more details and Supplementary Figure 14 for an ablation). For our templates, we use additive Gaussian noise as our only augmentation. For a description of multi-channel template augmentations, see Supplement F.

4.3 ENCODERS

We employ separate encoders for each electrophysiological modality. For the ACG image encoder, we use a version of the convolutional architecture introduced in (Beau et al., 2024) with 2 layers and Gaussian Error Linear Units (GeLU) (Hendrycks & Gimpel, 2016). For the waveform encoder, we use a 2 layer multilayer perceptron (MLP) with GeLU units. The representation sizes are 200 dimensional and 300 dimensional for the ACG image encoder and the waveform encoder, respectively. We set the projection size to be 512. For details about hyperparameters, see Supplement B.

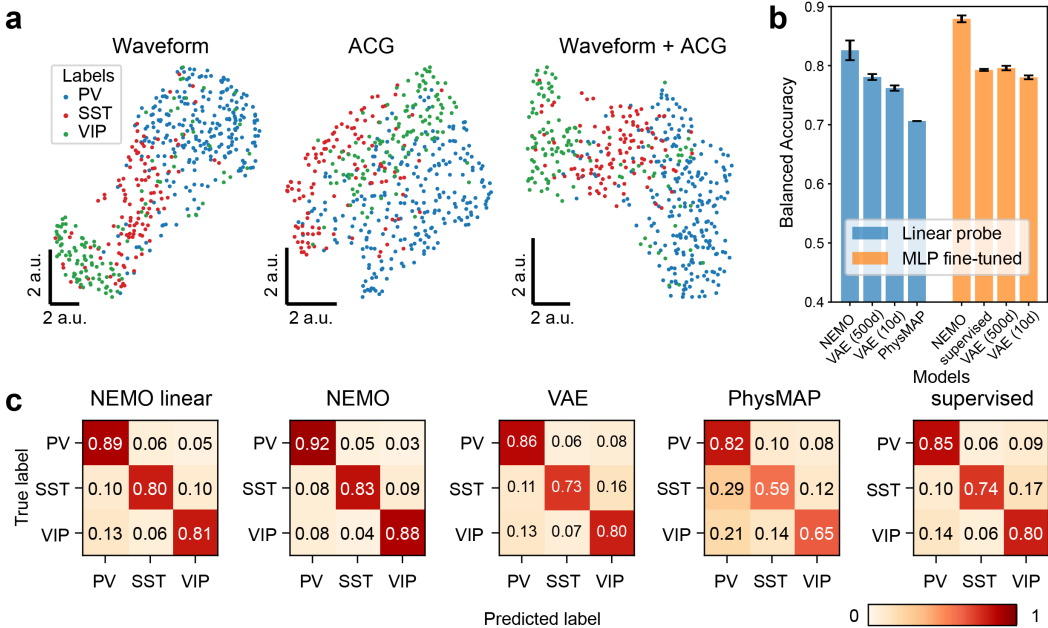


Figure 2: **Comparing NEMO to baseline models on the NP Ultra opto-tagged dataset.** (a) UMAP visualization of the pretrained NEMO representations of unseen opto-tagged visual cortex units, colored by different cell-types. Neurons of the same class form clusters, particularly when combined modalities are used. (b) Balanced accuracy and (c) Confusion matrices for the NP Ultra classification results, normalized by ground truth label and averaged across 5 random seeds. NEMO outperforms the other embedding methods by a significant margin across all cell-types. Surprisingly, NEMO outperforms all other methods even when just using a linear classifier on the embeddings.

4.4 CONTRASTIVE OBJECTIVE

We utilize the contrastive objective defined in CLIP. Let z_{acg} and z_{wf} be the L2 normalized projections of each modality. For a batch B , the objective is as follows,

$$\mathcal{L} = -\frac{1}{2|B|} \sum_{i=1}^{|B|} \left[\log \frac{\exp(z_{acg_i} \cdot z_{wf_i} / \tau)}{\sum_{j=1}^{|B|} \exp(z_{acg_i} \cdot z_{wf_j} / \tau)} + \log \frac{\exp(z_{acg_i} \cdot z_{wf_i} / \tau)}{\sum_{j=1}^{|B|} \exp(z_{acg_j} \cdot z_{wf_i} / \tau)} \right] \quad (1)$$

where τ is a temperature parameter which we fix during training. The objective function encourages the model to correctly match z_{acg_i} with its corresponding z_{wf_i} , and vice versa, over all other possible pairs in the batch. This loss can easily be extended to more than two modalities including PSTHs.

4.5 SINGLE-NEURON AND MULTI-NEURON BRAIN REGION CLASSIFICATION

Brain region classification is a new problem for electrophysiological datasets that requires novel classification schemes. We develop two classification schemes for our evaluation: a single-neuron and multi-neuron classifier. For our single-neuron classifier, we predict the brain area for each neuron independently using its embedding. For our multi-neuron classifier, we predict the brain region for each 20 micron bin along the depth of the probe by ensembling the predictions of nearby neurons within a 60-micron radius (i.e., averaging the logits of the single-neuron model) as shown in Figure 3a (ii). When more than five neurons fall within this range, only the nearest five are selected.

5 EXPERIMENTAL SETUP

5.1 BASELINES

For our baselines, we compare against current state-of-the-art multimodal cell-type embedding methods and a fully supervised method. For the multimodal cell-type embedding methods, we

Table 1: **Cell-type classification for the NP-Ultra dataset.** The accuracy and F1-scores are reported for three conditions: (i) a linear layer and (ii) MLP on top of the frozen pretrained representations (for VAE and NEMO), and (iii) after MLP finetuning.

Model	Linear		MLP		MLP fine-tuned	
	Acc	F1	Acc	F1	Acc	F1
Supervised	N/A	N/A	N/A	N/A	0.79 ± 0.00	0.79 ± 0.00
PhysMAP	0.71 ± 0.00^1	0.69 ± 0.00^1	N/A	N/A	N/A	N/A
VAE (10d latent)	0.76 ± 0.00	0.76 ± 0.00	0.75 ± 0.01	0.75 ± 0.01	0.78 ± 0.00	0.78 ± 0.00
VAE (500d rep)	0.78 ± 0.00	0.78 ± 0.00	0.79 ± 0.00	0.79 ± 0.00	0.80 ± 0.00	0.79 ± 0.00
NEMO (500d rep)	0.83 ± 0.01	0.83 ± 0.01	0.84 ± 0.01	0.84 ± 0.01	0.88 ± 0.01	0.88 ± 0.01

compare to PhysMAP (Lee et al., 2024) and a VAE-based method (Beau et al., 2024). For fair comparison, we utilize the same encoder architectures for NEMO and the VAE-based method. We include two versions of the VAE baseline: (1) the latent space (10D) is used to predict cell-type or brain region (from Beau et al. (2024)), or (2) the layer before the latent space (500D) is used to predict cell-type or brain region. For the VAE, we use default hyperparameters from Beau et al. (2024) (see Supplement L for a hyperparameter sensitivity analysis). For the supervised baseline, we again use the same encoder architectures as NEMO. For training NEMO, we use an early stopping strategy which utilizes validation data. For the VAE-based method, we use the training scheme introduced in (Beau et al., 2024). We fix the hyperparameters for all methods across both datasets. For more details about baselines, training, and hyperparameters, see Supplements B and E.

5.2 EVALUATION

For both NEMO and the VAE-based method, the representations from the ACG image and EAPs are concatenated together before classification or fine-tuning. We apply three classification schemes for evaluation including (1) freezing the model and training a linear classifier on the final layer, (2) freezing the model and training a MLP-based classifier on the final layer, (3) fine-tuning both the original model and a MLP-based classifier on the final layer. To ensure balanced training data, we implement dataset resampling prior to fitting the linear classifier. The only method that cannot be fine-tuned is PhysMAP, as it is UMAP-based and therefore non-differentiable. For PhysMAP comparisons, we utilize the weighted graph approach with alignment mapping provided in (Lee et al., 2024) for all comparisons. The macro-averaged F1 score, calculated as the unweighted mean of F1 scores for each class, and balanced accuracy, which measures average accuracy per class, are selected as the primary metrics for our classification tasks. This choice addresses data imbalance, a common issue in opto-tagged datasets and brain region classification. For additional details about baseline hyperparameters, see Supplement B.

Table 2: **single-neuron brain region classification for the IBL dataset.** The accuracy and F1-scores are reported for three conditions: (i) a linear layer and (ii) MLP on top of the frozen pretrained representations (for VAE and NEMO), and (iii) after MLP finetuning. We only ran PhysMAP one time as it is deterministic.

Model	Linear		MLP		MLP fine-tuned	
	Acc	F1	Acc	F1	Acc	F1
Supervised	N/A	N/A	N/A	N/A	0.45 ± 0.01	0.42 ± 0.01
PhysMAP	0.31^1	0.27^1	N/A	N/A	N/A	N/A
VAE	0.33 ± 0.00	0.29 ± 0.00	0.40 ± 0.01	0.37 ± 0.00	0.43 ± 0.00	0.40 ± 0.00
NEMO	0.42 ± 0.00	0.40 ± 0.00	0.45 ± 0.01	0.42 ± 0.00	0.47 ± 0.01	0.44 ± 0.01

5.3 EXPERIMENTS

NP Ultra opto-tagged dataset. For the NP Ultra dataset, we pretrain NEMO and the VAE-based method on 8699 unlabelled neurons. This pretraining strategy is important for reducing overfitting to the small quantity of labeled cell-types. To evaluate each model after pretraining, we perform the three evaluation schemes introduced in Section 5.2: freezing + linear classifier, freezing + MLP

¹We utilize PhysMAP’s anchor alignment technique to evaluate its performance (not a linear classifier).

Table 3: **multi-neuron brain region classification for the IBL dataset.** The accuracy and F1-scores for brain region classification are reported for three conditions: (i) a linear layer and (ii) MLP on top of the frozen pretrained representations (for VAE and NEMO), and (iii) after MLP finetuning.

Model	Linear		MLP		MLP fine-tuned	
	Acc	F1	Acc	F1	Acc	F1
Supervised	N/A	N/A	N/A	N/A	0.50 ± 0.00	0.48 ± 0.01
VAE	0.36 ± 0.00	0.32 ± 0.00	0.45 ± 0.00	0.42 ± 0.00	0.48 ± 0.00	0.46 ± 0.00
NEMO	0.48 ± 0.00	0.45 ± 0.00	0.50 ± 0.00	0.48 ± 0.00	0.51 ± 0.00	0.50 ± 0.00

classifier, and full end-to-end finetuning. For PhysMAP, we utilize the anchor alignment technique introduced by (Lee et al., 2024). For all methods and evaluation schemes, we perform a 5-fold cross-validation with 10 repeats to evaluate each model.

IBL Brain-wide Map. For the IBL dataset, we randomly divide all insertions (i.e., Neuropixels recordings) into a 70-10-20 split to create a training, validation, and test set for each method. We then pretrain NEMO and the VAE-based method on all neurons in the training split. We then perform the three evaluation schemes introduced in Section 5. For PhysMAP, we utilize the anchor alignment technique. We train both a single-neuron and multi-neuron classifier using the representations learned by NEMO and the VAE-based method. For PhysMAP, we only evaluate the single-neuron classifier. We compute the average and standard deviation of the metrics using five random seeds.

6 RESULTS

6.1 CLASSIFICATION

NP Ultra cell-type classifier. The results for the NP Ultra opto-tagged dataset are shown in Table 1 and Figure 2. For all three evaluation schemes, NEMO achieves the highest macro-averaged F1 score and balanced accuracy by a significant margin. Surprisingly, the VAE-based method, even with MLP-based fine-tuning, is outperformed by a frozen NEMO model with a linear decoder (by $\sim 5\%$). With MLP fine-tuning, NEMO reaches .88 F1 score and balanced accuracy which is an $\sim 11\%$ improvement over the baselines. For the VAE-based method, we found that utilizing the 500D representations *before* the latent space performed better than using the 10D latent space and also outperformed a VAE trained with a 500D latent space so we included this as a baseline. Except for NEMO, all baseline models struggle to differentiate VIP and SST cells as they have the fewest labels and the VIP cells display bimodality in their waveform distribution. These results demonstrate that NEMO is an accurate method for cell-type classification in visual cortical microcircuits.

IBL single-neuron region classifier. We then aim to investigate how much relevant information NEMO extracts from each neuron about its anatomical location, i.e., brain region. We investigate this by training classifiers that use single neuron features to identify anatomical regions for the IBL dataset (see Table 2 and Figure 3b for results). We again find that NEMO outperforms all other methods using both linear and MLP-based classification schemes. While NEMO outperforms the supervised baseline, we find that the VAE-based method performs worse than a fully supervised model trained on the same raw data. Without end-to-end fine-tuning, NEMO with an MLP classification head is already on par with the supervised MLP. NEMO’s success with both the linear and MLP classifier with frozen encoder weights indicates that NEMO is able to extract a region-discriminative representation of neurons without additional fine-tuning. This representation can be further improved by fine-tuning NEMO and the classifier end-to-end. The confusion matrix for PhysMAP is shown in Supplementary Figure 8.

IBL multi-neuron region classifier. We further investigate whether combining information from multiple neurons at each location can improve brain region classification. We use the nearest-neurons ensembling approach as described in 4.5 and shown in Figure 3a. Averaging the logits of predictions from single neurons improves classification performance over the single-neuron model. NEMO still has the best region classification performance (see Table 3 and Figure 3c for results).

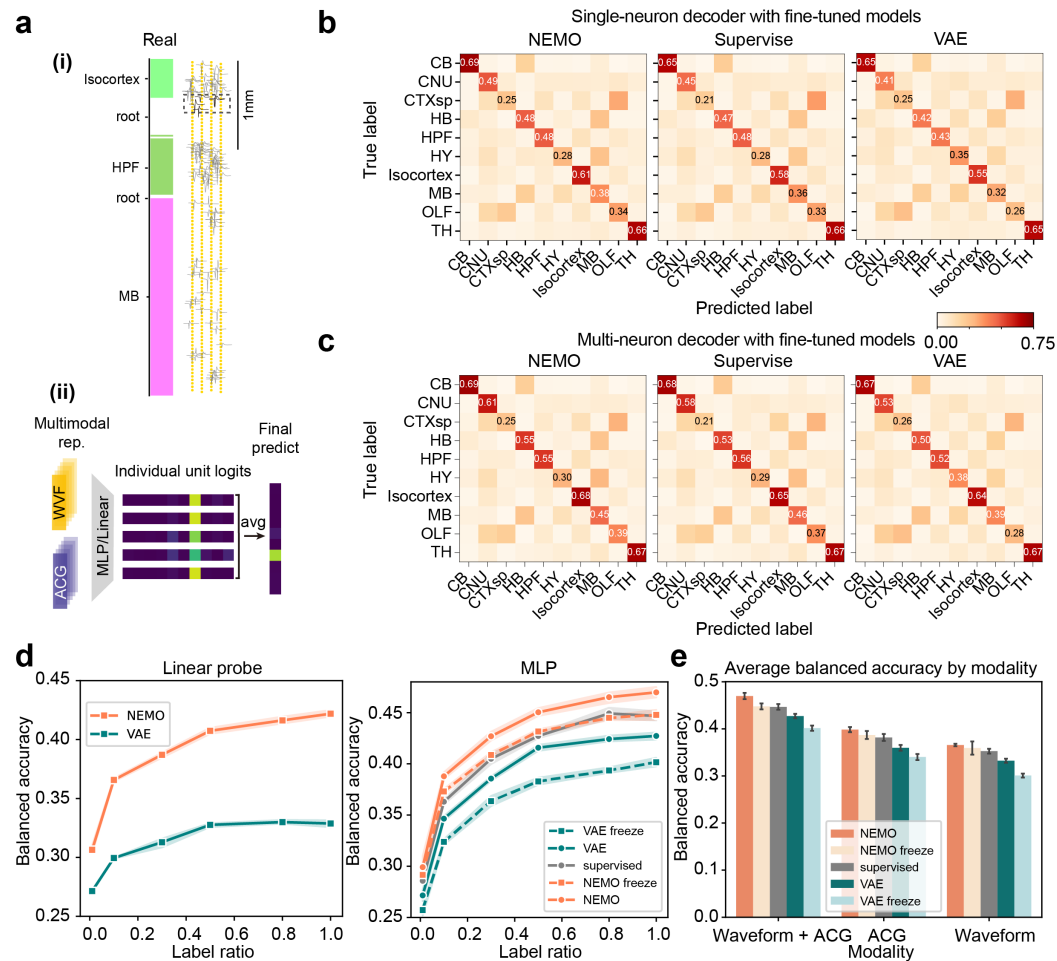


Figure 3: **Results for NEMO on the IBL brain region classification task.** (a) Schematic for multi-neuron classifier. (i) At each depth, the neurons within $60 \mu\text{m}$ were used to classify the anatomical region. Only the nearest 5 neurons were selected if there were more than 5 neurons within that range. (ii) For logits averaging, single-neuron classifier logits are predicted using a linear model/MLP trained on the representations of our two physiological modalities. The final prediction is based on the average of the individual logits. (b) Confusion matrices for the single-neuron region classifier with fine-tuned NEMO pretrained encoders, the fully supervised model, and with VAE pretrained encoders, averaged across 5 runs. (c) Confusion matrices for the multi-neuron region classifier, averaged across 5 runs. (d) Single neuron balanced accuracy with linear classifier and the MLP head for each model trained/fine-tuned with different label ratios. (e) Single-neuron MLP-classification balanced accuracy for each modality separately and for the combined representation.

6.2 CLUSTERING

We next examine the clusterability of NEMO representations for the IBL Brain-wide map. We followed the clustering strategy used in Lee et al. (2021) by running Louvain clustering on a UMAP graph constructed from the representations extracted by NEMO from the IBL training neurons. We adjusted two main settings: the neighborhood size in UMAP and the resolution in Louvain clustering. We selected these parameters by maximizing the modularity index, which had the effect of minimizing the number of resulting clusters (Figure 4c). The clustering results relative to the region labels are presented in Figures 4a and b. The UMAP visualization of the NEMO representations, colored by region label, demonstrates that the regions are separable in the representation space. Notably, there is a distinct separation of thalamic neurons from other regions, along with an isolated cluster of cerebellar neurons. Neurons from other regions are also well organized by region labels within the NEMO representation space, allowing for their clustering into several distinct clusters. Additionally, overlaying the neurons colored by their cluster IDs onto their anatomical

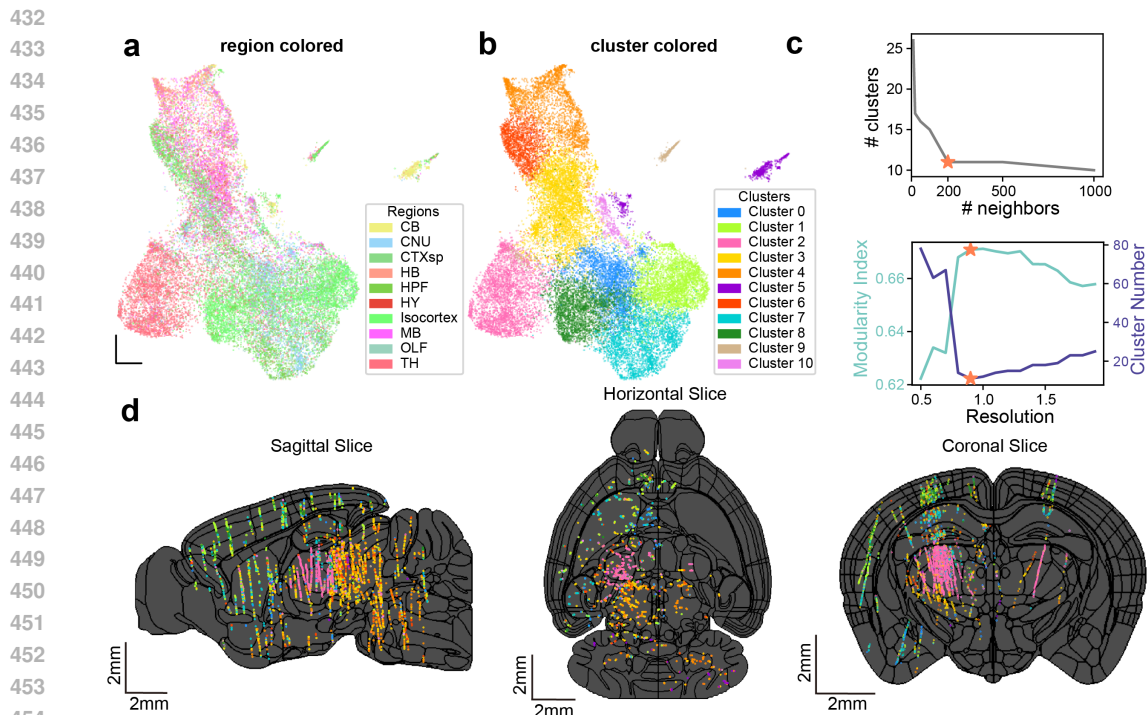


Figure 4: **IBL neuron clustering using NEMO pretraining.** (a) A UMAP visualization of the representations that NEMO extracts from the training data colored by anatomical brain region. (b) The same UMAP as shown in (a) but instead colored by cluster labels using a graph-based approach (Louvain clustering). (c) We tuned the neighborhood size in UMAP and the resolution for the clustering. These parameters were selected by maximizing the modularity index which minimized the number of clusters. (d) 2D brain slices across three brain views with the location of individual neurons colored using the cluster IDs shown in (b). The black lines show the region boundaries of the Allen mouse atlas (Wang et al., 2020). The cluster distribution found using NEMO is closely correlated with the anatomical regions and is consistent across insertions from different labs.

locations (Figure 4) reveals a cluster distribution closely correlated with anatomical regions which is consistent across insertions from different labs (Supplementary Figure 10). We find that clustering NEMO’s representations leads to a more region-selective clustering than when we use the raw features directly (Supplementary Figures 11 and 12). These results suggest that NEMO is able to extract features that capture the electrophysiological diversity across regions.

6.3 ABLATIONS

Label ratio sweep. We hypothesize that NEMO requires less labeled data for fine-tuning to achieve comparable performance to other models for brain region classification due to the contrastive pre-training. To test this, we conducted a label ratio sweep with our single-neuron region classifiers. We trained the linear classifier and the MLP classifier under two conditions: with frozen weights and with full end-to-end fine-tuning. We use 1%, 10%, 30%, 50%, 80%, and 100% of the labeled data for this experiment. The accuracy results are depicted in Figure 3d (for F1, see Supplementary Figure 5). The NEMO fine-tuned model outperforms all other methods for every label ratio. For the linear model, using just 10% of the training labels, NEMO achieves superior performance compared to using the full training set for the VAE. Similarly, NEMO paired with the MLP-classifier shows impressive results; with only 50% of the labeled data, the NEMO outperforms the supervised MLP. With 30% of the labels, NEMO exceeds the performance of the VAE model trained with all labels.

Single modality classifier. We next explore the additional region-relevant information gained by combining both modalities instead of just one, and whether NEMO enhances information extraction from each modality by aligning the embeddings of the two. We compared the classification performance of the MLP classifier with encoder weights frozen and end-to-end fine-tuned, across all models using: 1) waveforms only 2) ACGs only 3) combining waveforms and ACGs. Brain region

486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

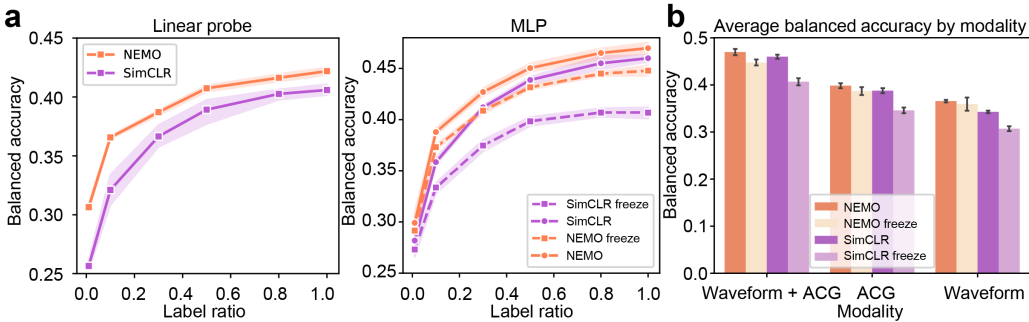


Figure 5: **Ablating joint vs. independent learning for NEMO.** To evaluate the importance of learning a shared representation between modalities, we train a version of NEMO on the IBL brain classification task where each modality is independently embedded using SimCLR. We refer to this method as SimCLR. (a) Across all label ratios and classifiers, we find that NEMO trained with CLIP outperforms the SimCLR version. (b) NEMO trained with CLIP also extracts more informative representations for each modality (especially waveforms) than when training with SimCLR.

classification balanced accuracies are shown in Figure 3e (for F1, see Supplementary Figure 5). We found that bimodal models, in general, outperform unimodal models, indicating that combining both modalities provides extra information on the anatomical location of neurons. We also find that the NEMO fine-tuned model delivers the best performance. This demonstrates that NEMO improves the information extracted from a single modality by leveraging information from the other modality.

Joint vs. independent learning for NEMO. To ablate the importance of learning a shared representation of each modality, we train a version of NEMO where we independently learn an embedding for each modality using a unimodal contrastive method, SimCLR (Chen et al., 2020). The results of the IBL brain region classification task are shown in Figure 5a where NEMO trained with CLIP outperforms NEMO trained with SimCLR for all label ratios and classification methods. NEMO trained with CLIP is also able to extract more informative representations from each modality (especially for waveforms) as shown in Figure 5b. These results demonstrate that learning a shared representation of the two modalities is important for the downstream performance of NEMO. Supplementary Table 9 shows that joint training with CLIP leads to a modest improvement over SimCLR for cell-type classification of the NP Ultra dataset.

7 DISCUSSION

In this work, we proposed NEMO, a pretraining framework for electrophysiological data that utilizes multi-modal contrastive learning. We demonstrate that NEMO is able to extract informative representations for cell-type and brain region classification with minimal fine-tuning. This capability is particularly valuable in neuroscience research where ground truth data, such as opto-tagged cells, are costly and labor-intensive to acquire (or in some cases even impossible to acquire, for example in human datasets). Our work has some limitations. Firstly, we primarily focus on shared information between two modalities, assuming this is most informative for identifying cell identity or anatomical location. While our analysis using a single modality classifier indicates that our model captures modality-exclusive information, it would be beneficial to distinguish between the shared and private information within each modality (Liu et al., 2024; Liang et al., 2024). Additionally, NEMO utilizes the activity of each neuron independently. Recent work has demonstrated that there is additional information in neuron correlations that can help distinguish different cell-types (Mi et al., 2023).

NEMO opens up several promising avenues for future research. Our framework can be adapted for studies of peripheral nervous systems, such as the retina (Wu et al., 2023). NEMO can also be combined with RNA sequencing to find features that are shared between RNA and electrophysiological data (Li et al., 2023). It will also be possible to correlate the cell-types discovered using NEMO with animal behavior to characterize their functional properties. Finally, we imagine that the representations found by NEMO can be integrated with current large-scale pretraining approaches for neural population activity to provide additional cell-type information which could improve generalizability to unseen sessions or animals (Azabou et al., 2023; Ye et al., 2023a; Zhang et al., 2024).

REFERENCES

- 540
541
542 Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna:
543 A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM*
544 *SIGKDD international conference on knowledge discovery & data mining*, pp. 2623–2631, 2019.
- 545
546 Antonis Antoniadis, Yiyi Yu, Joseph Canzano, William Wang, and Spencer LaVere Smith. Neu-
547 roformer: Multimodal and multitask generative pretraining for brain data. *arXiv preprint*
548 *arXiv:2311.00136*, 2023.
- 549 Thomas Aynaud. python-louvain x.y: Louvain algorithm for community detection.
550 <https://github.com/taynaud/python-louvain>, 2020.
- 551
552 Mehdi Azabou, Mohammad Gheshlaghi Azar, Ran Liu, Chi-Heng Lin, Erik C Johnson, Kiran
553 Bhaskaran-Nair, Max Dabagia, Bernardo Avila-Pires, Lindsey Kitchell, Keith B Hengen, et al.
554 Mine your own view: Self-supervised learning through across-sample prediction. *arXiv preprint*
555 *arXiv:2102.10106*, 2021.
- 556 Mehdi Azabou, Vinam Arora, Venkataramana Ganesh, Ximeng Mao, Santosh Nachimuthu,
557 Michael J Mendelson, Blake Richards, Matthew G Perich, Guillaume Lajoie, and Eva L Dyer.
558 A unified, scalable framework for neural population decoding. *arXiv preprint arXiv:2310.16046*,
559 2023.
- 560 Peter Barthó, Hajime Hirase, Lenaïc Monconduit, Michael Zugaro, Kenneth D Harris, and Gyorgy
561 Buzsáki. Characterization of neocortical principal cells and interneurons by network interactions
562 and extracellular features. *Journal of neurophysiology*, 92(1):600–608, 2004.
- 563
564 Maxime Beau, Federico D’Agostino, Ago Lajko, Gabriela Martínez, and Dimitar Kostadinov.
565 Neuropixels: loading, processing and plotting neuropixels data in python. *Zenodo* [https://doi.](https://doi.org/10.5281/ZENODO)
566 [org/10.5281/ZENODO](https://doi.org/10.5281/ZENODO), 5509733, 2021.
- 567 Maxime Beau, David J Herzfeld, Francisco Naveros, Marie E Hemelt, Federico D’Agostino, Mar-
568 lies Oostland, Alvaro Sánchez-López, Young Yoon Chung, Michael Maibach, Hannah N Stabb,
569 et al. A deep-learning strategy to identify cell types across species from high-density extracellular
570 recordings. *bioRxiv*, pp. 2024–01, 2024.
- 571 BioRender.com. Biorender. <https://www.biorender.com/>. Accessed: 2024-05-22.
- 572
573 Claire Bomkamp, Shreejoy J Tripathy, Carolina Bengtsson Gonzales, Jens Hjerling-Leffler,
574 Ann Marie Craig, and Paul Pavlidis. Transcriptomic correlates of electrophysiological and mor-
575 phological diversity within and across excitatory and inhibitory neuron classes. *PLoS computa-*
576 *tional biology*, 15(6):e1007113, 2019.
- 577
578 Julien Boussard, Charlie Windolf, Cole Hurwitz, Hyun Dong Lee, Han Yu, Olivier Winter, and Liam
579 Paninski. Dartsort: A modular drift tracking spike sorter for high-density multi-electrode probes.
580 *bioRxiv*, pp. 2023–08, 2023.
- 581
582 Alessio P Buccino, Cole L Hurwitz, Samuel Garcia, Jeremy Magland, Joshua H Siegle, Roger
583 Hurwitz, and Matthias H Hennig. Spikeinterface, a unified framework for spike sorting. *Elife*, 9:
e61834, 2020.
- 584
585 Stephane Bugeon, Joshua Duffield, Mario Dipoppa, Anne Ritoux, Isabelle Prankerd, Dimitris Ni-
586 coloutsopoulos, David Orme, Maxwell Shinn, Han Peng, Hamish Forrest, et al. A transcriptomic
587 axis predicts state modulation of cortical interneurons. *Nature*, 607(7918):330–338, 2022.
- 588
589 Jessica A Cardin, Marie Carlén, Konstantinos Meletis, Ulf Knoblich, Feng Zhang, Karl Deisseroth,
590 Li-Huei Tsai, and Christopher I Moore. Targeted optogenetic stimulation and recording of neu-
591 rons in vivo using cell-type-specific expression of channelrhodopsin-2. *Nature protocols*, 5(2):
247–254, 2010.
- 592
593 Hanbo Chen, Jiawei Yang, Daniel Iascone, Lijuan Liu, Lei He, Hanchuan Peng, and Jianhua Yao.
Treemoco: Contrastive neuron morphology representation learning. *Advances in Neural Infor-*
mation Processing Systems, 35:25060–25073, 2022.

- 594 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for
595 contrastive learning of visual representations. In *International conference on machine learning*,
596 pp. 1597–1607. PMLR, 2020.
- 597
- 598 Sven Dorkenwald, Peter H Li, Michał Januszewski, Daniel R Berger, Jeremy Maitin-Shepard,
599 Agnes L Bodor, Forrest Collman, Casey M Schneider-Mizell, Nuno Maçarico da Costa, Jeff W
600 Lichtman, et al. Multi-layered maps of neuropil with segmentation-guided contrastive learning.
601 *Nature Methods*, 20(12):2011–2020, 2023.
- 602 Nathan W Gouwens, Staci A Sorensen, Jim Berg, Changkyu Lee, Tim Jarsky, Jonathan Ting, Su-
603 san M Sunkin, David Feng, Costas A Anastassiou, Eliza Barkan, et al. Classification of electro-
604 physiological and morphological neuron types in the mouse visual cortex. *Nature neuroscience*,
605 22(7):1182–1195, 2019.
- 606
- 607 Xifeng Guo, Long Gao, Xinwang Liu, and Jianping Yin. Improved deep embedded clustering with
608 local structure preservation. In *Ijcai*, volume 17, pp. 1753–1759, 2017.
- 609
- 610 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint*
611 *arXiv:1606.08415*, 2016.
- 612
- 613 Wei Huang, Andi Han, Yongqiang Chen, Yuan Cao, Zhiqiang Xu, and Taiji Suzuki. On
614 the comparison between multi-modal and single-modal contrastive learning. *arXiv preprint*
615 *arXiv:2411.02837*, 2024.
- 616
- 617 IBL, Kush Banga, Julius Benson, Niccolò Bonacchi, Sebastian A Bruijns, Rob Campbell, Gaëlle A
618 Chapuis, Anne K Churchland, M Felicia Davatolhagh, Hyun Dong Lee, et al. Reproducibility of
in-vivo electrophysiological measurements in mice. *bioRxiv*, pp. 2022–05, 2022.
- 619
- 620 IBL, Brandon Benson, Julius Benson, Daniel Birman, Niccolo Bonacchi, Matteo Carandini, Joana A
621 Catarino, Gaëlle A Chapuis, Anne K Churchland, Yang Dan, et al. A brain-wide map of neural
622 activity during complex behaviour. *bioRxiv*, pp. 2023–07, 2023.
- 623
- 624 Xiaoxuan Jia, Joshua H Siegle, Corbett Bennett, Samuel D Gale, Daniel J Denman, Christof Koch,
625 and Shawn R Olsen. High-density extracellular probes reveal dendritic backpropagation and
facilitate neuron classification. *Journal of neurophysiology*, 121(5):1831–1847, 2019.
- 626
- 627 Jonathan Jouty, Gerrit Hilgen, Evelyne Sernagor, and Matthias H Hennig. Non-parametric physi-
628 ological classification of retinal ganglion cells in the mouse retina. *Frontiers in Cellular Neuro-*
629 *science*, 12:481, 2018.
- 630
- 631 James J Jun, Nicholas A Steinmetz, Joshua H Siegle, Daniel J Denman, Marius Bauza, Brian Barbar-
632 its, Albert K Lee, Costas A Anastassiou, Alexandru Andrei, Çağatay Aydın, et al. Fully integrated
silicon probes for high-density recording of neural activity. *Nature*, 551(7679):232–236, 2017.
- 633
- 634 Alexxai V Kravitz, Scott F Owen, and Anatol C Kreitzer. Optogenetic identification of striatal
635 projection neuron subtypes during in vivo recordings. *Brain research*, 1511:21–32, 2013.
- 636
- 637 Patrick Latuske, Oana Toader, and Kevin Allen. Interspike intervals reveal functionally distinct cell
638 populations in the medial entorhinal cortex. *Journal of Neuroscience*, 35(31):10963–10976, 2015.
- 639
- 640 Phuc H Le-Khac, Graham Healy, and Alan F Smeaton. Contrastive representation learning: A
framework and review. *Ieee Access*, 8:193907–193934, 2020.
- 641
- 642 Candice Lee, Andreeanne Lavoie, Jiashu Liu, Simon X Chen, and Bao-hua Liu. Light up the brain:
643 the application of optogenetics in cell-type specific dissection of mouse brain circuits. *Frontiers*
644 *in neural circuits*, 14:18, 2020.
- 645
- 646 Eric Kenji Lee, Hymavathy Balasubramanian, Alexandra Tsolias, Stephanie Udochukwu Anakwe,
647 Maria Medalla, Krishna V Shenoy, and Chandramouli Chandrasekaran. Non-linear dimensional-
ity reduction on extracellular waveforms reveals cell type diversity in premotor cortex. *Elife*, 10:
e67490, 2021.

- 648 Eric Kenji Lee, Asim Gul, Gregory Heller, Anna Lakunina, Santiago Jaramillo, Pawel Przytycki,
649 and Chandramouli Chandrasekaran. Physmap-interpretable in vivo neuronal cell type identifica-
650 tion using multi-modal analysis of electrophysiological data. *bioRxiv*, pp. 2024–02, 2024.
651
- 652 Qiang Li, Zuwan Lin, Ren Liu, Xin Tang, Jiahao Huang, Yichun He, Xin Sui, Weiwen Tian, Hao
653 Shen, Haowen Zhou, et al. Multimodal charting of molecular and functional cell states via in situ
654 electro-sequencing. *Cell*, 186(9):2002–2017, 2023.
- 655 Paul Pu Liang, Zihao Deng, Martin Q Ma, James Y Zou, Louis-Philippe Morency, and Ruslan
656 Salakhutdinov. Factorized contrastive learning: Going beyond multi-view redundancy. *Advances*
657 *in Neural Information Processing Systems*, 36, 2024.
658
- 659 Shengzhong Liu, Tomoyoshi Kimura, Dongxin Liu, Ruijie Wang, Jinyang Li, Suhas Diggavi, Mani
660 Srivastava, and Tarek Abdelzaher. Focal: Contrastive learning for multimodal time-series sensing
661 signals in factorized orthogonal latent space. *Advances in Neural Information Processing Systems*,
662 36, 2024.
- 663 Jeremy Magland, James J Jun, Elizabeth Lovero, Alexander J Morley, Cole Lincoln Hurwitz,
664 Alessio Paolo Buccino, Samuel Garcia, and Alex H Barnett. Spikeforest, reproducible web-facing
665 ground-truth validation of automated neural spike sorters. *Elife*, 9:e55167, 2020.
666
- 667 Richard H Masland. Neuronal cell types. *Current Biology*, 14(13):R497–R500, 2004.
668
- 669 RT Matthews and WL Lee. A comparison of extracellular and intracellular recordings from medial
670 septum/diagonal band neurons in vitro. *Neuroscience*, 42(2):451–462, 1991.
- 671 Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and
672 projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018a.
- 673
- 674 Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. Umap: Uniform manifold
675 approximation and projection. *The Journal of Open Source Software*, 3(29):861, 2018b.
- 676
- 677 Lu Mi, Trung Le, Tianxing He, Eli Shlizerman, and Uygur Sümbül. Learning time-invariant rep-
678 resentations for individual neurons from population dynamics. *Advances in Neural Information*
679 *Processing Systems*, 36:46007–46026, 2023.
- 680
- 681 Vernon B Mountcastle, William H Talbot, Hideo Sakata, and J Hyvärinen. Cortical neuronal mecha-
682 nisms in flutter-vibration studied in unanesthetized monkeys. neuronal periodicity and frequency
683 discrimination. *Journal of neurophysiology*, 32(3):452–484, 1969.
- 684
- 685 Lionel G Nowak, Rony Azouz, Maria V Sanchez-Vives, Charles M Gray, and David A McCormick.
686 Electrophysiological classes of cat primary visual cortical neurons in vivo as revealed by quanti-
687 tative analyses. *Journal of neurophysiology*, 89(3):1541–1566, 2003.
- 688
- 689 Marius Pachitariu, Shashwat Sridhar, Jacob Pennington, and Carsen Stringer. Spike sorting with
690 kilosort4. *Nature Methods*, pp. 1–8, 2024.
- 691
- 692 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
693 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
694 models from natural language supervision. In *International conference on machine learning*, pp.
695 8748–8763. PMLR, 2021.
- 696
- 697 Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-
698 networks. *arXiv preprint arXiv:1908.10084*, 2019.
- 699
- 700 Aaqib Saeed, David Grangier, and Neil Zeghidour. Contrastive learning of general-purpose audio
701 representations. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and*
Signal Processing (ICASSP), pp. 3875–3879. IEEE, 2021.
- 702
- 703 Aidan Schneider, Mehdi Azabou, Louis McDougall-Vigier, David F Parks, Sahara Ensley, Kiran
704 Bhaskaran-Nair, Tomasz Nowakowski, Eva L Dyer, and Keith B Hengen. Transcriptomic cell
705 type structures in vivo neuronal activity across multiple timescales. *Cell reports*, 42(4), 2023a.

- 702 Steffen Schneider, Jin Hwa Lee, and Mackenzie Weygandt Mathis. Learnable latent embeddings for
703 joint behavioural and neural analysis. *Nature*, 617(7960):360–368, 2023b.
704
- 705 Nicholas A Steinmetz, Christof Koch, Kenneth D Harris, and Matteo Carandini. Challenges and
706 opportunities for large-scale electrophysiology with neuropixels probes. *Current opinion in neu-*
707 *robiology*, 50:92–100, 2018.
- 708 Nicholas A Steinmetz, Cagatay Aydin, Anna Lebedeva, Michael Okun, Marius Pachitariu, Marius
709 Bauza, Maxime Beau, Jai Bhagat, Claudia Böhm, Martijn Broux, et al. Neuropixels 2.0: A minia-
710 turized high-density probe for stable, long-term brain recordings. *Science*, 372(6539):eabf4588,
711 2021.
712
- 713 Susan M Sunkin, Lydia Ng, Chris Lau, Tim Dolbeare, Terri L Gilbert, Carol L Thompson, Michael
714 Hawrylycz, and Chinh Dang. Allen brain atlas: an integrated spatio-temporal portal for exploring
715 the central nervous system. *Nucleic acids research*, 41(D1):D996–D1008, 2012.
716
- 717 Bosiljka Tasic, Zizhen Yao, Lucas T Graybuck, Kimberly A Smith, Thuc Nghi Nguyen, Darren
718 Bertagnolli, Jeff Goldy, Emma Garren, Michael N Economo, Sarada Viswanathan, et al. Shared
719 and distinct transcriptomic cell types across neocortical areas. *Nature*, 563(7729):72–78, 2018.
- 720 Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Computer*
721 *Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings,*
722 *Part XI 16*, pp. 776–794. Springer, 2020.
723
- 724 Gemechu B Tolossa, Aidan M Schneider, Eva L Dyer, and Keith B Hengen. A conserved code for
725 anatomy: Neurons throughout the brain embed robust signatures of their anatomical location into
726 spike trains. *bioRxiv*, pp. 2024–07, 2024.
- 727 Eric M Trautmann, Janis K Hesse, Gabriel M Stine, Ruobing Xia, Shude Zhu, Daniel J O’Shea,
728 Bill Karsh, Jennifer Colonell, Frank F Lanfranchi, Saurabh Vyas, et al. Large-scale high-density
729 brain-wide neural recording in nonhuman primates. *bioRxiv*, pp. 2023–02, 2023.
730
- 731 Carolina Urzay, Nauman Ahad, Mehdi Azabou, Aidan Schneider, Geethika Atamkuri, Keith B Hen-
732 gen, and Eva L Dyer. Detecting change points in neural population activity with contrastive metric
733 learning. In *2023 11th International IEEE/EMBS Conference on Neural Engineering (NER)*, pp.
734 1–4. IEEE, 2023.
- 735 Ganesh Vigneswaran, Alexander Kraskov, and Roger N Lemon. Large identified pyramidal cells in
736 macaque motor and premotor cortex exhibit “thin spikes”: implications for cell type classification.
737 *Journal of Neuroscience*, 31(40):14235–14242, 2011.
738
- 739 Ankit Vishnubhotla, Charlotte Loh, Akash Srivastava, Liam Paninski, and Cole Hurwitz. Towards
740 robust and generalizable representations of extracellular data using contrastive learning. *Advances*
741 *in Neural Information Processing Systems*, 36, 2024.
- 742 Quanxin Wang, Song-Lin Ding, Yang Li, Josh Royall, David Feng, Phil Lesnar, Nile Graddis,
743 Maitham Naeemi, Benjamin Facer, Anh Ho, et al. The allen mouse brain common coordinate
744 framework: a 3d reference atlas. *Cell*, 181(4):936–953, 2020.
745
- 746 Keiko Weir, Oriane Blanquie, Werner Kilb, Heiko J Luhmann, and Anne Sinning. Comparison of
747 spike parameters from optically identified gabaergic and glutamatergic neurons in sparse cortical
748 cultures. *Frontiers in cellular neuroscience*, 8:460, 2015.
- 749 Eric G Wu, Andra M Rudzite, Martin O Bohlen, Peter H Li, Alexandra Kling, Sam Cooler, Colleen
750 Rhoades, Nora Brackbill, Alex R Gogliettino, Nishal P Shah, et al. Decomposition of retinal
751 ganglion cell electrical images for cell type and functional inference. *bioRxiv*, 2023.
752
- 753 Zizhen Yao, Cindy TJ Van Velthoven, Thuc Nghi Nguyen, Jeff Goldy, Adriana E Seden-Cortes,
754 Fahimeh Baftizadeh, Darren Bertagnolli, Tamara Casper, Megan Chiang, Kirsten Crichton, et al.
755 A taxonomy of transcriptomic cell types across the isocortex and hippocampal formation. *Cell*,
184(12):3222–3241, 2021.

756 Zizhen Yao, Cindy TJ van Velthoven, Michael Kunst, Meng Zhang, Delissa McMillen, Changkyu
757 Lee, Won Jung, Jeff Goldy, Aliya Abdelhak, Matthew Aitken, et al. A high-resolution transcrip-
758 tomic and spatial atlas of cell types in the whole mouse brain. *Nature*, 624(7991):317–332, 2023.
759

760 Joel Ye, Jennifer Collinger, Leila Wehbe, and Robert Gaunt. Neural data transformer 2: multi-
761 context pretraining for neural spiking activity. *bioRxiv*, pp. 2023–09, 2023a.


762 Zhiwen Ye, Andrew M Shelton, Jordan R Shaker, Julien Bussard, Jennifer Colonell, Daniel Bir-
763 man, Sahar Manavi, Susu Chen, Charlie Windolf, Cole Hurwitz, et al. Ultra-high density elec-
764 trodes improve detection, yield, and cell type identification in neuronal recordings. *bioRxiv*,
765 2023b.

766 Hongkui Zeng and Joshua R Sanes. Neuronal cell-type classification: challenges, opportunities and
767 the path forward. *Nature Reviews Neuroscience*, 18(9):530–546, 2017.
768

769 Yizi Zhang, Tianxiao He, Julien Bussard, Charles Windolf, Olivier Winter, Eric Trautmann, Noam
770 Roth, Hailey Barrell, Mark Churchland, Nicholas A Steinmetz, et al. Bypassing spike sorting:
771 Density-based decoding using spike localization from dense multielectrode probes. *Advances in*
772 *Neural Information Processing Systems*, 36, 2024.
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

A IBL BRAIN-WIDE-MAP DATASET

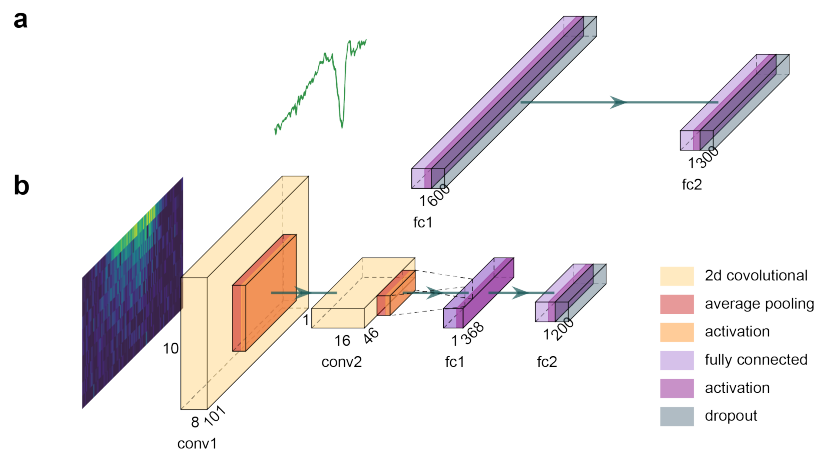
The full IBL Brain-Wide-Map dataset contains 699 insertions. Among those, only 675 successfully went through the pre-processing, with 37017 units that pass the IBL spike-sorting quality control. The region distribution of those units are shown in Supplementary Figure 1.



Supplementary Figure 1: Region distribution of the IBL Brain-Wide-Map dataset. The dataset is very imbalanced and has a small number of units in hypothalamus and cortical subplates.

B BASELINES AND HYPERPARAMETERS

B.1 MODEL ARCHITECTURE



Supplementary Figure 2: Encoder architecture illustration: a) A two-layer multilayer perceptron for waveform encoding, and b) A convolutional neural network for ACG image encoding.

B.2 MODEL PARAMETERS

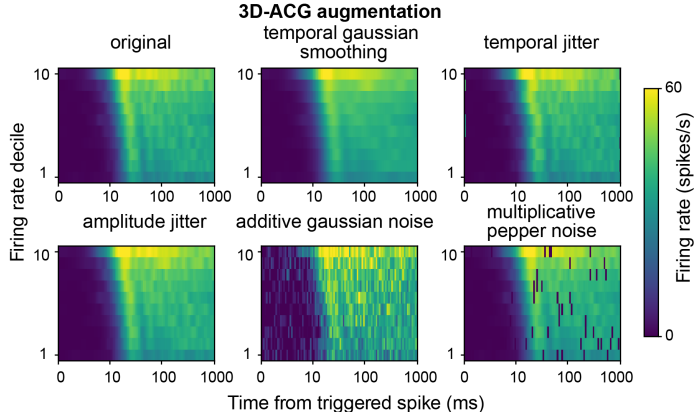
We apply the model augmentations in table 1 on the training data during NEMO. p is the probability an augmentation gets applied to an instance. We use the Adam optimizer with learning rate 0.0005 and Cosine Annealing scheduler with $T_0 = 20$. Other model hyperparameters are in table 2.

Supplementary Table 1: Model Augmentations

Augmentation	Description	Type	p
Gaussian noise	Gaussian noise with mean 0 and std $0.1 \times \text{std of WVF}$	WVF	0.3
Temporal Gaussian smoothing	Smooths an ACG using a Gaussian filter along the temporal axis with $\sigma = 2$	ACG	0.5
Temporal jittering	Jitters the temporal axis of an ACG by a random integer between -3 and 3 inclusive	ACG	0.5
Amplitude scaling	Scales the amplitude of an ACG by a random number between 0.9 and 1.1	ACG	0.5
Additive Gaussian noise	Adds Gaussian noise with mean 0 and std $0.1 \times \text{maximum of ACG}$	ACG	0.5
Multiplicative pepper noise	Each value in the ACG has a 5% of being set to 0	ACG	0.5

B.3 MLP FINE-TUNING

For cell-type classification, we fix the hyperparameters of the MLP classifier for all methods to those used in Beau et al. (2024). This is because we have very few labels for evaluation and it is challenging to hold out any for hyperparameter tuning. For the brain region classification experiments, we tune the optimizer, learning rate, model size and dropout using the validation set of insertions.



Supplementary Figure 3: Illustration of the augmentations on spiking activity ACGs used by NEMO.

Supplementary Table 2: Model Hyperparameters

Parameter	Description	Value
Epochs	Numbers of epochs to run	6000
Batch Size	Number of items processed in a single operation	1024
Learning Rate	Learning rate of the model	0.0005
Log every n steps	Save model and cross validate results every n epochs	100
dim_embed	Latent dimension size	512
Temperature	τ in the formula for contrastive loss	0.5
ACG dropout	Dropout probability in ACG encoder	0.2
Waveform dropout	Dropout probability in waveform encoder	0.1

B.4 SIMCLR-BASED NEMO

For the SimCLR-based NEMO, we randomly apply the same set of augmentation methods to each data point, generating two correlated augmented views. We then train each encoder and a linear projection layer to maximize agreement using contrastive learning. Compared to other methods, we use stronger augmentations for SimCLR. This is because SimCLR focuses on the differences between augmentations from the same data point, whereas CLIP-based learning strategies compare augmentations across different modalities. Stronger augmentations improved SimCLR’s performance.

Both methods use contrastive learning to maximize similarity between augmented views of the same neuron, while keeping views of different neurons distinct. The main difference between the SimCLR-based method and the CLIP-based method is that CLIP is multimodal and SimCLR is unimodal. For the SimCLR-based method, we use the contrastive objective from (Chen et al., 2020). Let $z_{\text{view}1}$ and $z_{\text{view}2}$ be the normalized projections of two augmented views from the same modality. For a batch B , the objective is as follows:

$$\mathcal{L} = -\frac{1}{2|B|} \sum_{i=1}^{|B|} \left[\log \frac{\exp(z_{\text{view}1_i} \cdot z_{\text{view}2_i} / \tau)}{\sum_{j=1}^{|B|} \mathbb{1}_{\{j \neq i\}} \exp(z_{\text{view}1_i} \cdot z_{\text{view}2_j} / \tau) + \sum_{j=1}^{|B|} \mathbb{1}_{\{j \neq i\}} \exp(z_{\text{view}1_i} \cdot z_{\text{view}1_j} / \tau)} + \log \frac{\exp(z_{\text{view}2_i} \cdot z_{\text{view}1_i} / \tau)}{\sum_{j=1}^{|B|} \mathbb{1}_{\{j \neq i\}} \exp(z_{\text{view}2_i} \cdot z_{\text{view}1_j} / \tau) + \sum_{j=1}^{|B|} \mathbb{1}_{\{j \neq i\}} \exp(z_{\text{view}2_i} \cdot z_{\text{view}2_j} / \tau)} \right] \quad (2)$$

For SimCLR, we use the same types of augmentation as CLIP with increased strength: for the EAPs, we applied Gaussian noise with standard deviation 1 with probability 0.5, rather 0.1 standard with a probability of than a 0.3. For the ACG images, we keep the augmentation methods consistent across all models.

In our SimCLR-based method, we use the pretrained encoder to obtain separate embeddings for each modality, which are then combined to train downstream tasks such as cell-type and brain region classification.

C AUGMENTATION ABLATION STUDY

D DIFFERENT PROJECTION TECHNIQUES FOR SIMCLR

We explored different projection techniques for SimCLR. One approach was to use a SimCLR model with a smaller projection size than the representation, reducing the projection size from 512 to 128. Another approach was to use a 2-layer MLP with a 512-dimensional hidden layer and a 128-dimensional projection size instead of a linear projection layer. However, our original linear projector still achieved the highest validation metrics, as shown in Supplementary Table 3. Therefore, we continued to use the original linear projector for all experiments.

Supplementary Table 3: Linear brain region classification for IBL dataset

Model	Acc	F1
original	0.388	0.338
128-dim linear projection	0.371	0.326
mlp-based projection	0.377	0.328

E DATASET SPLIT AND VALIDATION STRATEGY

Given the limited quantity of ground-truth labels for the NP Ultra dataset, we evaluate the performance of each model using a 5-fold cross-validation of the labeled cells with 10 repeats. Since we do not have a separate validation set while training NEMO, we utilize a nested cross-validation approach to choose an evaluation checkpoint for NEMO. In other words, for each cross-validation

fold, we perform a nested 5-fold cross-validation on the *training* folds with a linear classifier to choose the best checkpoint for NEMO. We then use this NEMO checkpoint to perform the evaluation on the heldout fold of the original 5-fold cross-validation. This checkpoint choosing procedure is done only on the training folds and does not use any information from the testing fold. Without this procedure, we are still able to achieve high performance.

For the IBL dataset, we use a standard 70-10-20 split for the training, validation, and test sets, respectively. During training, we monitor the performance of the model using a linear classifier trained on the training set and validated on the validation set. We then compute our evaluation metrics on the test set using the checkpoint with the highest validation F1.

F MULTICHANNEL

Using multi-channel templates has the potential to be more informative for identifying cell-types and classifying brain regions. In the following subsections, we demonstrate that using multi-channel templates does improve these downstream tasks, however, more work is needed to generalize this paradigm to different probe geometries.

F.1 IBL MULTICHANNEL

For the IBL multichannel experiments, we use 25 channel templates centered on the peak channel if possible (this is hard for edge spikes). We add two additional template augmentations: amplitude jitter (Vishnubhotla et al., 2024) and electrode dropout. Amplitude jitter rescales a channel’s amplitude by a uniform value between 0.9 and 1.1 with $p = 0.3$. Electrode dropout zeros out a channel with $p = 0.1$. If all channels are zeroed out, we leave the peak channel. Our model architecture is the same as NEMO with the only difference being the input size of the waveform encoder.

We tune hyperparameters for both the linear and MLP classification models. For the linear model, we tune the inverse of the regularization strength between $1e - 5$ and $1e4$ using the python module optuna (Akiba et al., 2019). For the MLP model, we do a grid search over the dropout probability(0.1, 0.2, 0.3, 0.4), learning rate($1e - 4$, $1e - 5$, $1e - 6$), number of layers(1, 2, 3) and layer size(128, 256, 512).

Supplementary Table 4: Multi-channel, single-neuron brain region decoding for the IBL dataset

Model	Linear		MLP		MLP fine-tuned	
	Acc	F1	Acc	F1	Acc	F1
NEMO	0.42 ± 0.00	0.40 ± 0.00	0.45 ± 0.01	0.42 ± 0.00	0.47 ± 0.01	0.44 ± 0.01
NEMO (25-channel)	0.45 ± 0.01	0.42 ± 0.00	0.46 ± 0.00	0.43 ± 0.00	0.48 ± 0.00	0.45 ± 0.00

Supplementary Table 5: multi-neuron, multi-channel brain region classification for the IBL dataset

Model	Linear		MLP		MLP fine-tuned	
	Acc	F1	Acc	F1	Acc	F1
NEMO	0.48 ± 0.00	0.45 ± 0.00	0.50 ± 0.00	0.48 ± 0.00	0.51 ± 0.00	0.50 ± 0.00
NEMO (25-channel)	0.50 ± 0.01	0.47 ± 0.00	0.51 ± 0.01	0.48 ± 0.01	0.52 ± 0.00	0.51 ± 0.01

F.2 NP-ULTRA MULTICHANNEL

For NP-Ultra multichannel, we tested using 9 or 25 channels centered on the peak channel if possible. We use the same augmentations and encoding model architecture as IBL multichannel and we do not finetune the hyperparameters in the MLP or linear model.

G PICKING PARAMETERS FOR IBL UNIT REPRESENTATION CLUSTERING

We used the Python-implemented UMAP package (McInnes et al., 2018b) and the Python-Louvain package (Aynaud, 2020) for our clustering analyses. For our clustering analysis, we aim to find the

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

Supplementary Table 6: Multichannel cell-type classification for the NP-Ultra dataset

Model	Linear		MLP (5-fold)		MLP fine-tuned(5-fold)	
	Acc	F1	Acc	F1	Acc	F1
NEMO (500d rep)	0.833 ± 0.008	0.833 ± 0.007	0.839 ± 0.007	0.839 ± 0.008	0.878 ± 0.005	0.878 ± 0.005
NEMO (25-channel)	0.853 ± 0.006	0.85 ± 0.006	0.865 ± 0.007	0.862 ± 0.007	0.875 ± 0.008	0.875 ± 0.007
NEMO (9-channel)	0.854 ± 0.014	0.851 ± 0.015	0.869 ± 0.003	0.865 ± 0.004	0.886 ± 0.003	0.883 ± 0.004

most informative clustering with the smallest number of clusters. There are two parameters to tune: 1) the size of local neighborhood used for the UMAP graph manifold approximation resolution that controls the community size in Louvain clustering ($n_{neighbor}$), and 2) the resolution γ that determines the size of the communities. We tuned $n_{neighbor}$ while keeping the resolution to be 1.0 and tracked the final number of clusters. We picked the ‘elbow’ that has the smallest $n_{neighbor}$ (200). The resolution was chosen by maximizing the modularity index Q of Louvain clustering with $n_{neighbor} = 200$. The modularity index of a graph is defined as:

$$Q = \sum_{c=1}^n \left[\frac{L_c}{m} - \gamma \left(\frac{k_c}{2m} \right)^2 \right]$$

where each c represents a community in the graph, m is the number of edges, L_c is the number of intra-community links for community c , k_c is the sum of degrees of the nodes in community c , and γ is the resolution parameter.

H PARAMETERS FOR CLASSIFIERS AND FINE-TUNING METHODS

Supplementary Table 7: Linear probe best hyperparameters

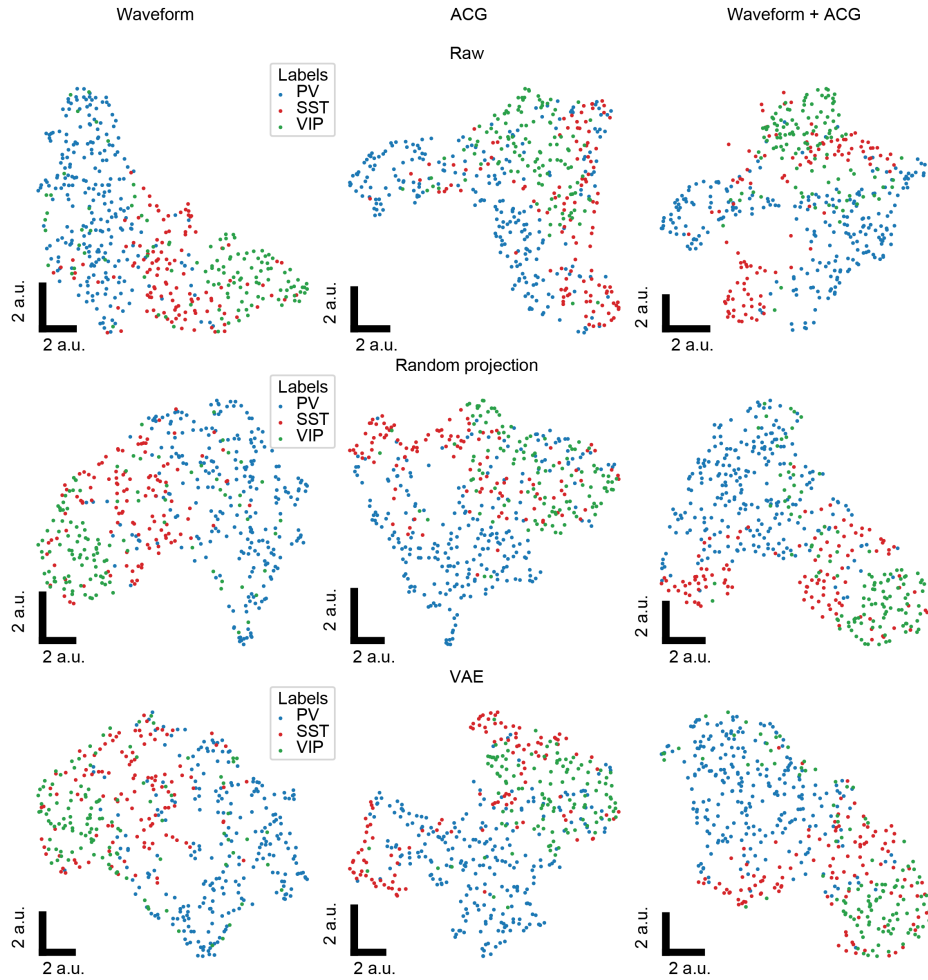
Hyperparameter	Value
max_{iter}	1000
tol	1e-5
NP-Ultra Celltype c	0.02
IBL NEMO joint c	0.02
IBL VAE c	0.001
IBL supervise c	0.001
IBL NEMO independent c	2.5

Supplementary Table 8: IBL MLP hyperparameters

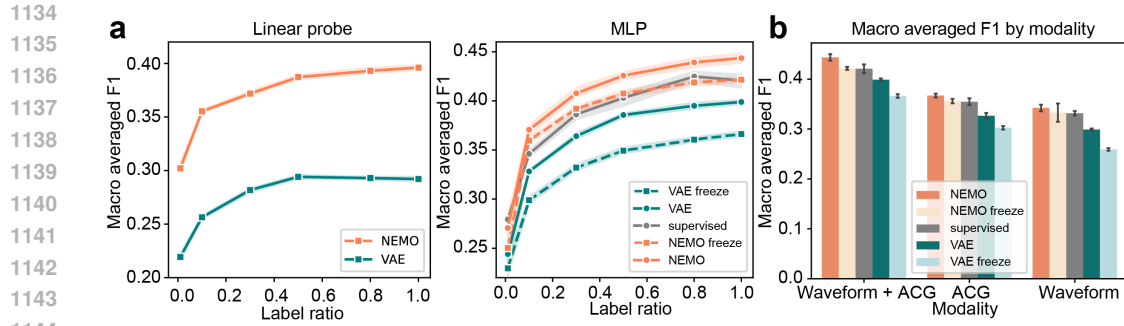
Hyperparameter	Value
n_{layer}	1
layer0 size	256
Dropout rate	0.2
Supervised scheduler	CosineAnnealingWarmRestarts
Supervised T_0	20
Supervised T_{mult}	1
Other models’ scheduler	StepLR
StepLR $step_{size}$	200
StepLR γ	0.8
Supervised lr	1×10^{-4}
Other models’ lr	1×10^{-5}

H.1 UMAP EMBEDDINGS WITH RAW FEATURES, RANDOM PROJECTION AND VAE

We visualize the UMAP embeddings of the raw features, randomly initialized encoder projection and the VAE projection in Supplementary figure 4. We find that NEMO representations are visually more structured.



Supplementary Figure 4: UMAP visualizations with raw features, random projections, and from the VAE. The random projections are representations which are passed through randomly initialized encoders with the same architecture as NEMO.



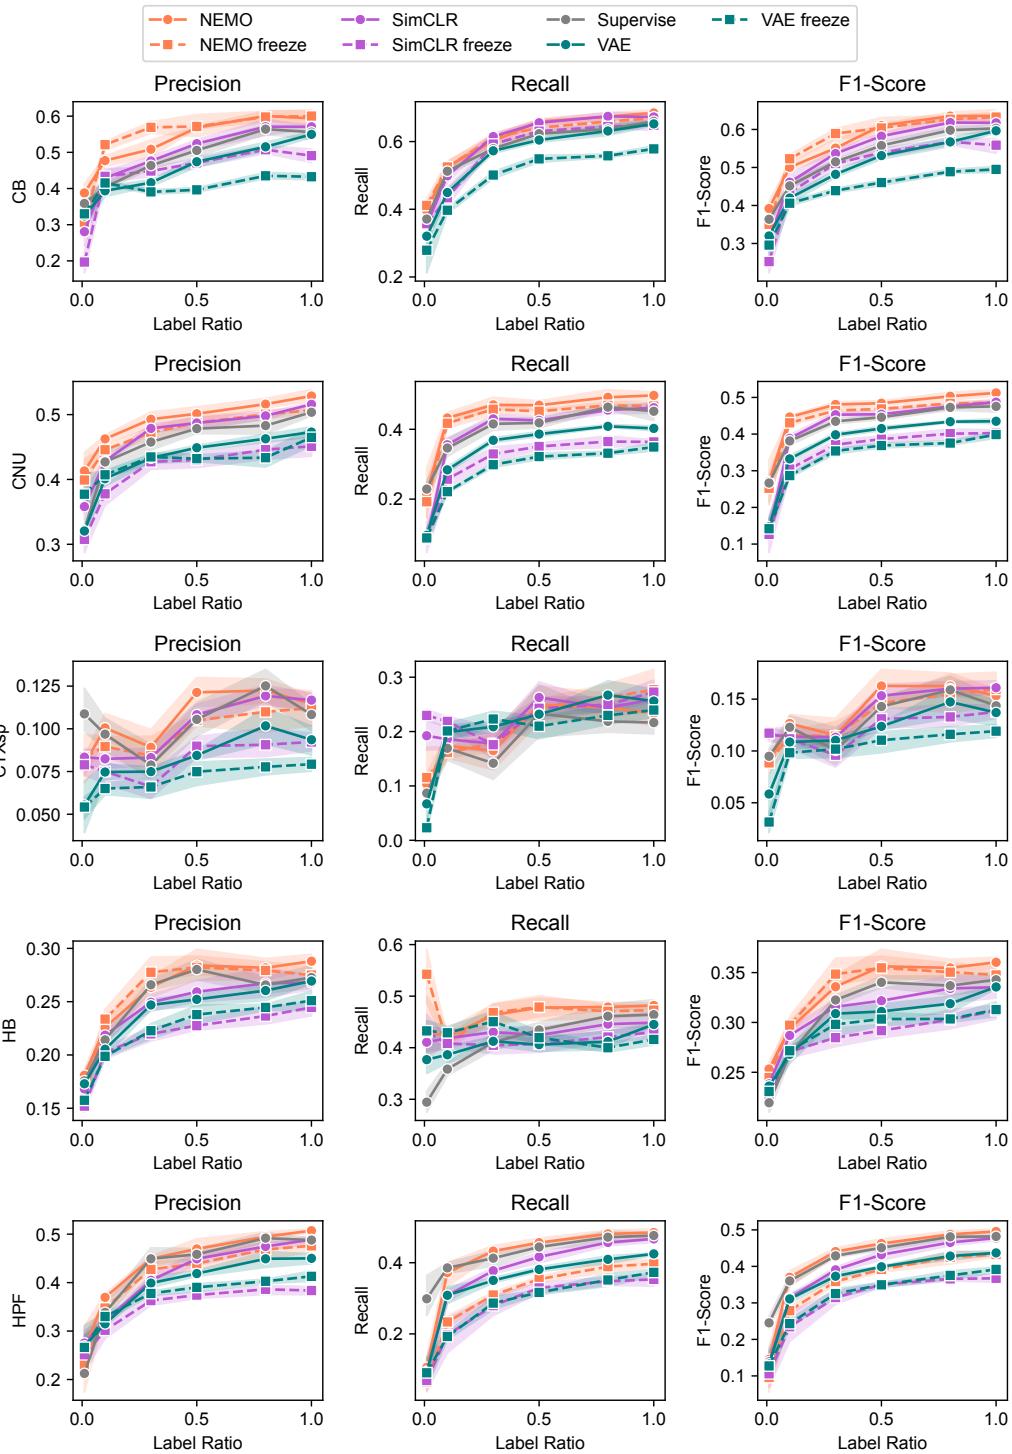
Supplementary Figure 5: Left and middle, macro-averaged F1 scores for linear and MLP-based single-neuron classification of brain region by label ratio, as in Fig. 3(d), replacing accuracy with F1 score. Right, single-neuron MLP-classification balanced F1 scores for uni- and bimodal models, as in Fig. 3(e), replacing accuracy with F1 score.

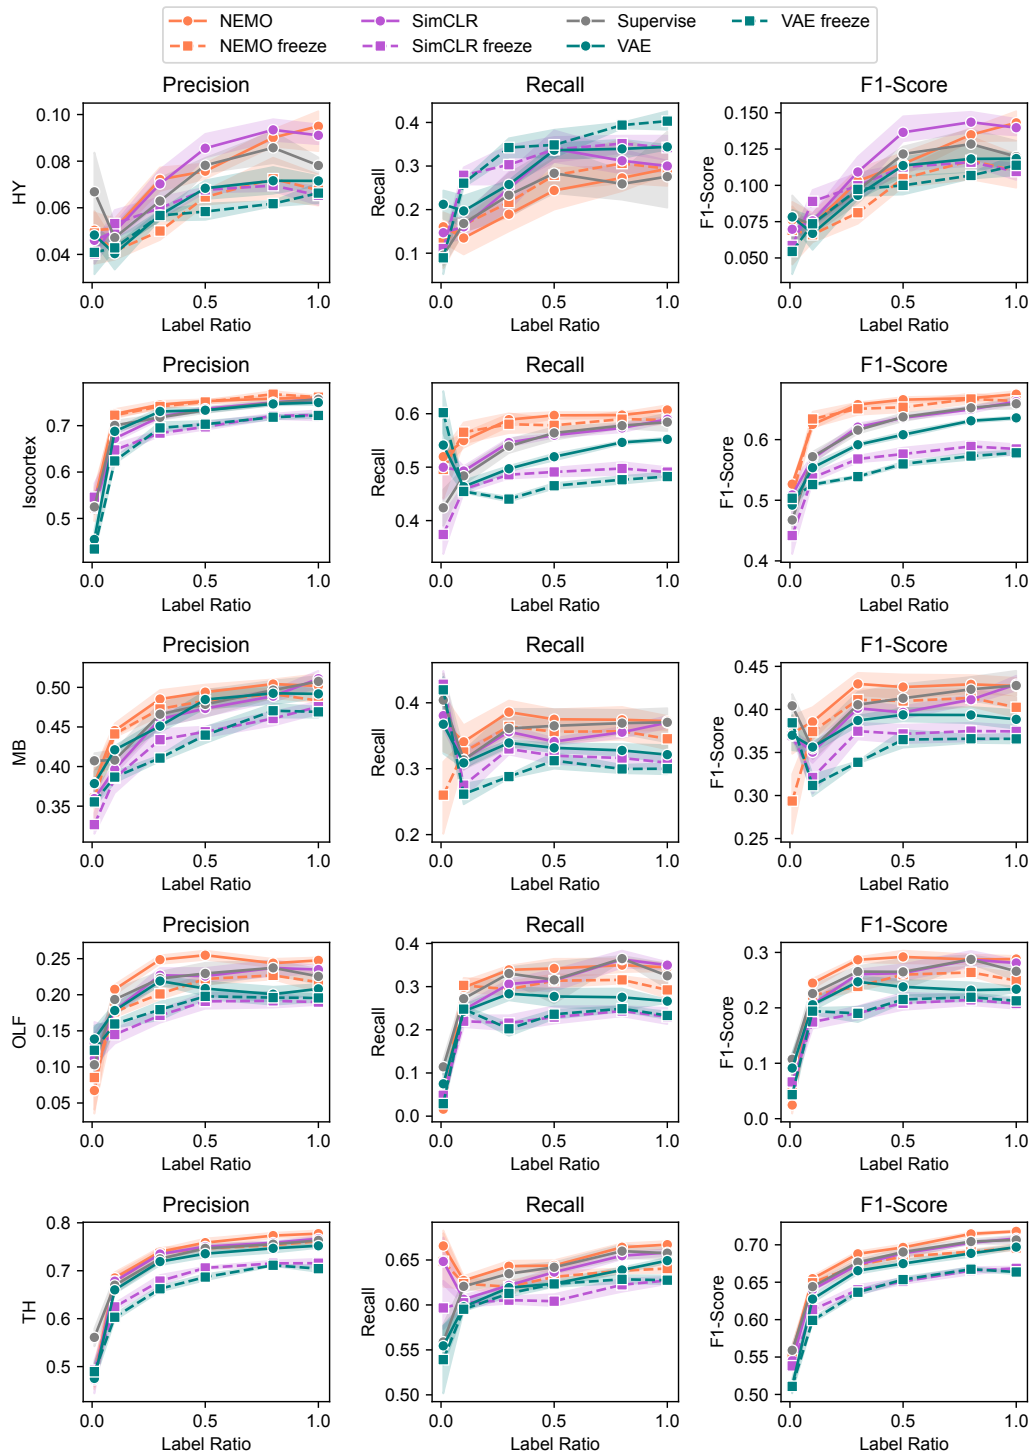
H.2 IBL REGION CLASSIFICATION WITH LABEL RATIO SWEEP

In Supplementary Figure 5, we show the macro-averaged F1 scores for single neuron classification of brain region by label ratio as complementary to Figure 3.

In Supplementary Figure 6, we study the effect of varying the ratio of labeled data used to train the brain region classifier. Fully supervised methods used the same labeled examples, but were not pretrained. Means and standard-deviation bands are computed over five random initializations for each label ratio.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241





Supplementary Figure 6: Detailed comparison of NEMO and SimCLR for region classification performance on IBL data with varying label ratio. We show the precision, recall and F1 score for each class. NEMO shows superior performance in most of the classes, except hypothalamus (HY), which has a small sample size compared to other classes.


I INDEPENDENT LEARNING ABLATION FOR NEMO

We compared the cell-type classification performance between NEMO with independent and joint learning. The results are shown in Supplementary table 9.

Supplementary Table 9: Cell-type classification for the NP-Ultra dataset for independent vs. joint learning NEMO.

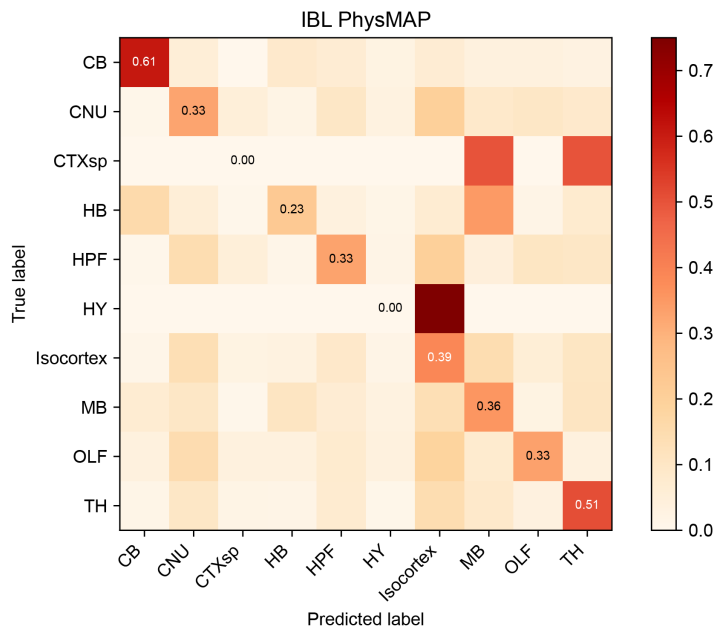
Model	Linear		MLP		Finetuned MLP	
	Acc	F1	Acc	F1	Acc	F1
independent NEMO	0.84 ± 0.00	0.83 ± 0.00	0.83 ± 0.00	0.82 ± 0.01	0.87 ± 0.01	0.87 ± 0.01
joint NEMO	0.83 ± 0.01	0.83 ± 0.01	0.84 ± 0.01	0.84 ± 0.01	0.88 ± 0.01	0.88 ± 0.01

Supplementary Figure 7 shows the macro-averaged F1 scores for single neuron classification of brain region by label ratio with the IBL dataset, complementary to Figure 5.



Supplementary Figure 7: Left and middle, macro-averaged F1 scores for linear and MLP-based single-neuron classification of brain region by label ratio for independent learning ablation, as in Fig. 5(a), replacing accuracy with F1 score. Right, single-neuron MLP-classification balanced F1 scores for uni- and bimodal models, as in Fig. 5(b), replacing accuracy with F1 score.

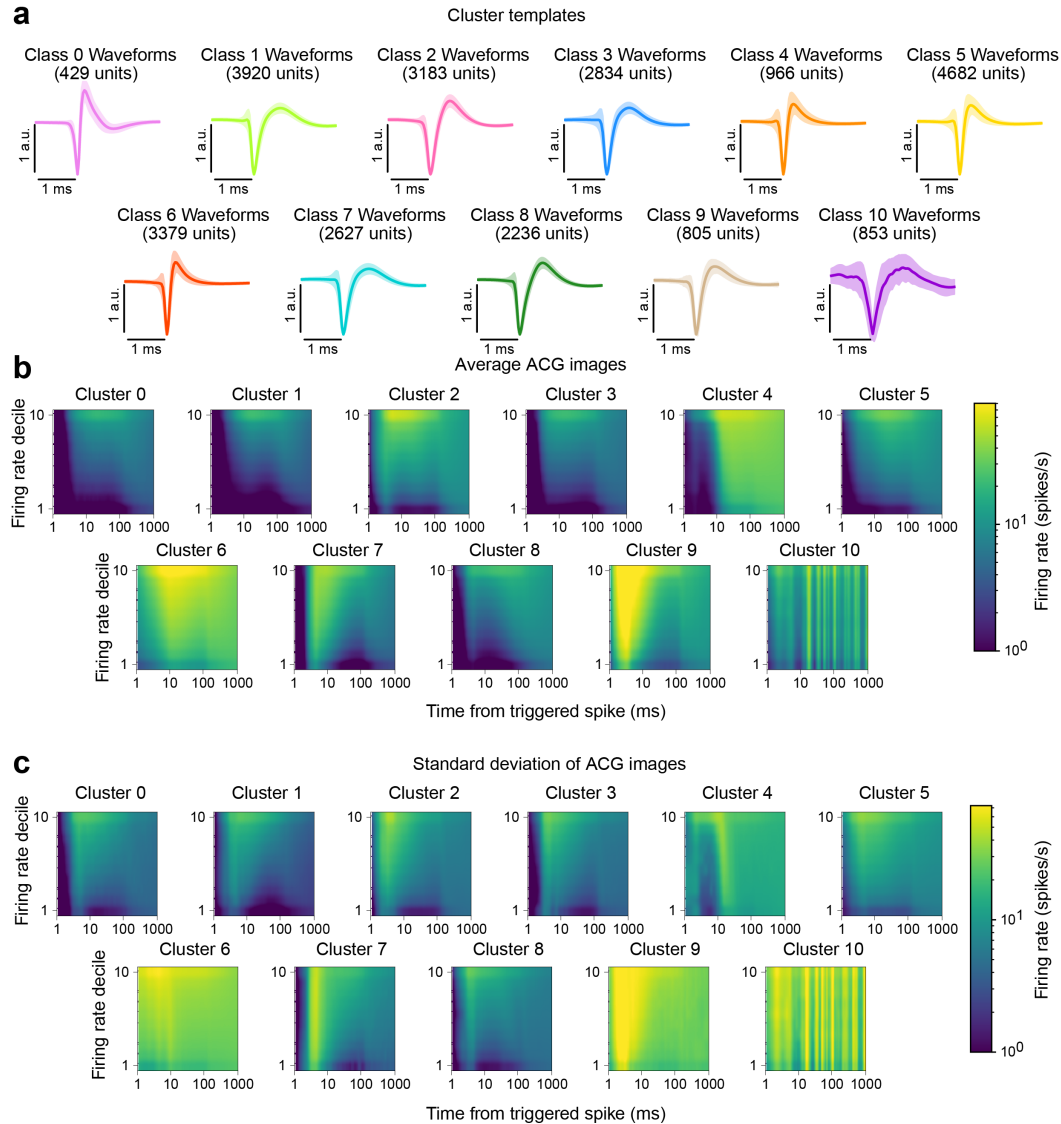
1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403



Supplementary Figure 8: Classification results of PhysMAP on the IBL brain region classification task. Due to label imbalance, PhysMAP is unable to predict CTXsp and HY, which leads to a low balanced accuracy and F1 score.

J IBL UNIT CLUSTERING RESULTS WITH NEMO

Supplementary Figure 9 shows averages and standard deviations for each cluster’s template waveforms and ACG images. Supplementary Figure 10 shows the distribution of cluster labels over brain regions, separated by individual insertions and by the recording lab.

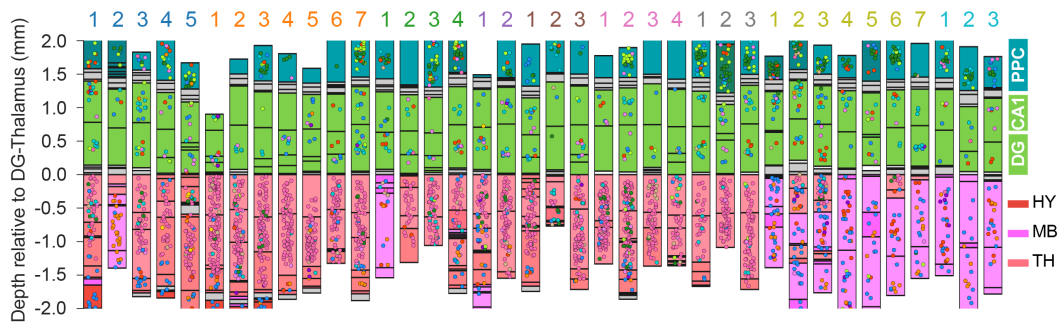


Supplementary Figure 9: Average and standard deviation of the template waveforms and ACGs for the units in each cluster, as shown in Figure 4. (a) The waveforms are consistent within clusters and distinct across clusters. (b) The ACG images are also distinct across clusters. These results suggest that NEMO is able to find distinct clusterings of neurons across the whole-brain.

K IBL UNIT CLUSTERING RESULTS WITH RAW FEATURE INPUT

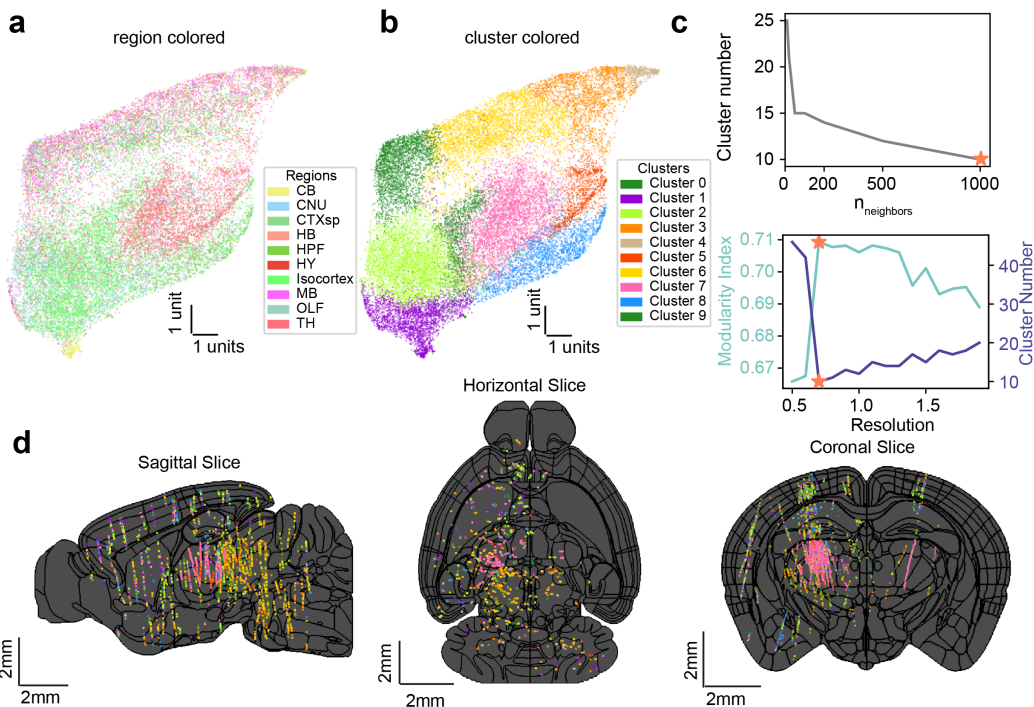
Supplementary Figure 11 shows clustering results for the IBL dataset. The hyper-parameters were selected with similar criteria as in Section 6. Since the number of clusters does not show a similar ‘elbow,’ but keeps decreasing as $n_{neighbor}$ increases, we picked $n_{neighbor} = 1000$ and used a resolution γ that maximizes the the modularity and minimizes the number of clusters. These clusters are

1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511



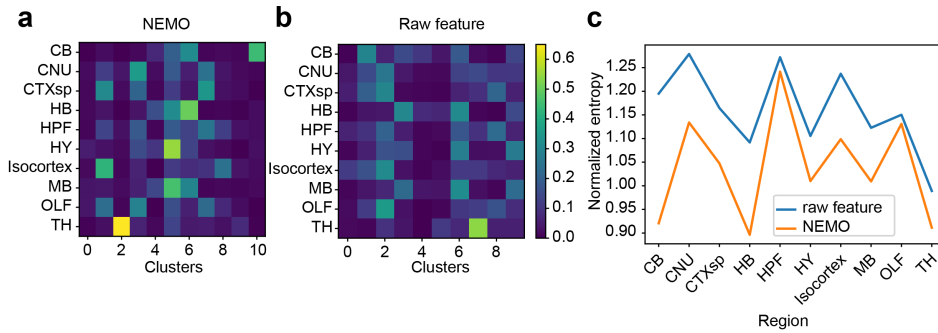
Supplementary Figure 10: Visualization of clustering results across repeated site in the IBL dataset (IBL et al., 2022). The Neuropixels probes target the same brain regions (including posterior parietal cortex, hippocampus, and thalamus) in these insertions. The color of the labels on top of each column indicates the lab ID of each insertion. Our results reveal that the clusters are highly distinguishable by region, with each region containing a distinct group of neurons. Moreover, the dominant cluster IDs for the same region remain consistent across different insertions.

less spatially organized compared to the clusters clustered using NEMO. As shown in Supplementary Figure 12, the clustering result based on NEMO overall shows lower entropy, which indicates higher region-selectivity.



Supplementary Figure 11: IBL neuron clustering using raw features. (a) A UMAP visualization of raw features on the training data colored by anatomical brain region. (b) The same UMAP as shown in (a) but instead colored by cluster labels using a graph-based approach (Louvain clustering). (c) We tuned the neighborhood size in UMAP and the resolution for the clustering. These parameters were selected by maximizing the modularity index which minimized the number of clusters. (d) 2D brain slices across three brain views with the location of individual neurons colored using the cluster IDs shown in (b). The black lines show the region boundaries of the Allen mouse atlas.

1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565

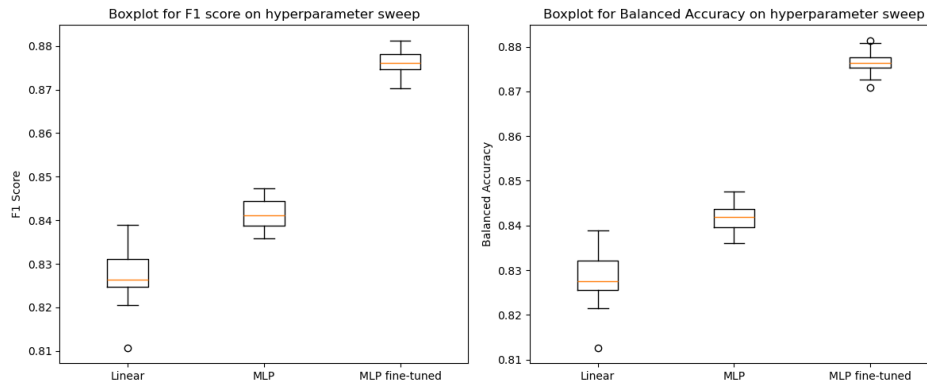


Supplementary Figure 12: Region distribution of clustering results using different NEMO features and raw features (normalized by region). For each region, we get a cluster distribution vector. We then computed the normalized entropy of that distribution. Cluster result based on NEMO overall shows lower entropy, which indicates higher region-selectivity.

L SENSITIVITY TO PARAMETERS

L.1 HYPERPARAMETER SWEEP

For assess the importance of hyperparameters in NEMO and the VAE baseline, we do a sweep over the learning rate (1e-4,5e-4,1e-3), the latent dimension size (256,512,1024), and the dropout (shifting the original dropout by -0.1,0,0.1) on the single channel NP-Ultra model with linear and MLP decoders. The dropout shift is added to the default waveform and acg dropouts. The boxplots of the resulting balanced accuracy and macro-averaged F1 scores for NEMO presented in Figure 13. Table 4 show the average performance across all hyperparameters for the VAE and NEMO. From these results, it can be seen that the hyperparameters do not have a large affect on the overall performance of either mention (for the grid search we tested).




Supplementary Figure 13: Boxplot of Macro-averaged F1 score and Balanced accuracy over 5 seeds of NEMO with linear, MLP and MLP fine-tuned decoders for the NEMO hyperparameter sweep.

L.2 AUGMENTATION ABLATION STUDY

To assess the importance of augmentations, we conducted an ablation experiment in two parts: (1) starting with the full set of augmentations, we removed one augmentation at a time, and (2) starting with no augmentations, we added one augmentation at a time. The resulting balanced accuracy and macro-averaged F1 scores are presented in Supplementary Figure 14. The results indicate that additive Gaussian noise in the ACG images is the most impactful augmentation, while the absence of augmentations significantly degrades performance. For other augmentations, the combined effect is large while the individual contribution is smaller.

Table 4: Average performance of the VAE and NEMO model for different hyperparameters. The accuracy and F1-scores are reported for three conditions: (i) a linear layer and (ii) MLP on top of the frozen pretrained representations (for VAE and NEMO), and (iii) after MLP finetuning.

Model	Linear		MLP		MLP fine-tuned	
	Acc	F1	Acc	F1	Acc	F1
VAE (latent)	0.76 ± 0.01	0.77 ± 0.02	0.75 ± 0.01	0.76 ± 0.01	0.789 ± 0.00	0.78 ± 0.00
VAE (rep)	0.78 ± 0.00	0.78 ± 0.00	0.79 ± 0.00	0.79 ± 0.00	0.80 ± 0.00	0.80 ± 0.00
NEMO (500d rep)	0.83 ± 0.01	0.83 ± 0.01	0.84 ± 0.00	0.84 ± 0.00	0.88 ± 0.00	0.88 ± 0.00

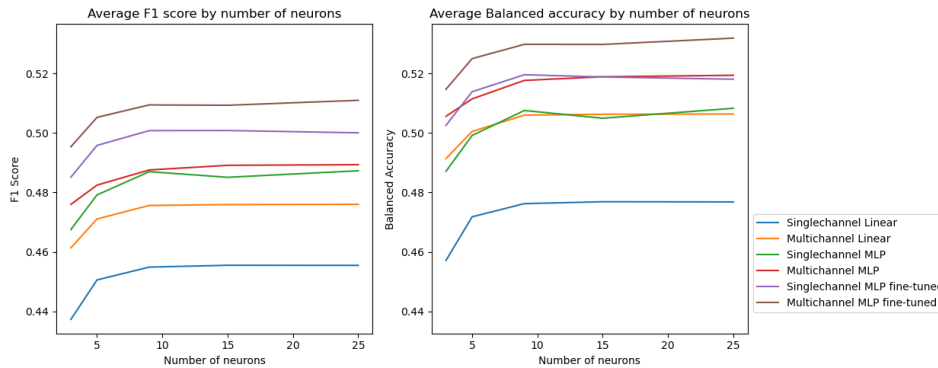


Supplementary Figure 14: To assess the impact of our augmentations, we performed two analyses on the UHD cell-type classification dataset: (1) removing one augmentation at a time, starting from the full set of augmentations, and (2) applying one augmentation at a time, starting from no augmentations. For each condition, we calculated the mean and standard deviation of (a) balanced accuracy and (b) macro-averaged F1 score with linear classifier.

L.3 MULTI-NEURON NEIGHBOR SWEEP

To evaluate the significance of the number of neurons we choose in our multi-neuron brain region classification model, we conduct a sweep over the maximum number of neighboring neurons (3, 5, 9, 15, 25) that can be selected within a 60-micron radius on the IBL dataset. As illustrated in Supplementary Figure 15, the overall classification performance improves as the maximum number of neighboring neurons increases, reaching a saturation point when this number reaches ~ 10 .

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631



1632
1633
1634
1635
1636
1637

Supplementary Figure 15: Macro-averaged F1 score and Balanced accuracy over 5 seeds of NEMO with multineuron classifier over different number of max neighboring neurons for single and multi-channel IBL with linear, frozen MLP, and MLP fine-tuning.

1638
1639

L.4 IBL CLUSTERING WITH DIFFERENT RANDOM SEEDS

1640
1641
1642
1643
1644
1645
1646

We trained five instances of NEMO with different random seeds and clustered the neurons using the same hyperparameters. The clustering results are visualized in Supplementary Figure 16, where UMAP was applied for dimensionality reduction, with colors distinguishing clusters. To ensure direct comparison, we maintain consistent UMAP visualizations across clustering results. The average adjusted Rand Index (ARI) for these outcomes is 0.48 ± 0.04 , while the adjusted mutual information (AMI) score is 0.58 ± 0.02 . These metrics indicate a moderate level of agreement among the clustering results.

1647
1648
1649
1650

The boundaries of the clusters are subtle, introducing stochasticity into the clustering results. In addition, clustering the full brain dataset results in a coarse segmentation, while significant variability is expected within individual regions. Therefore, these clustering results should mainly be interpreted as exploratory.

1651
1652

M STATISTICAL TESTS

1653
1654
1655
1656
1657

To evaluate performance differences between models, we conducted two-sample one-tailed t-tests (significance threshold $p < .05$) on the macro-averaged F1 scores and balanced accuracies across five random seeds. The results are summarized in a significance matrix comparing models: red indicates the model in the column is significantly better than the model in the row, blue indicates it is significantly worse, and white indicates no significant difference.

1659

1660
1661

M.1 CELL-TYPE CLASSIFICATION

1662
1663
1664
1665
1666
1667

For both linear and MLP classifiers, NEMO significantly outperforms all baselines except SimCLR (Supplementary Figure 17). Specifically, NEMO significantly outperforms SimCLR on balanced accuracy when using the fine-tuned MLP and the frozen MLP, but the difference is not significant with the linear classifier. On the F1 score, NEMO significantly outperforms SimCLR with the frozen MLP but not with the fine-tuned MLP or the linear classifier (though the p-value is 0.0505 for the fine-tuned MLP, narrowly missing the $p < .05$ threshold).

1668

1669
1670

M.2 BRAIN REGION CLASSIFICATION


1671
1672
1673

NEMO consistently outperforms all other baselines for both linear and MLP classifiers (Supplementary Figure 18). For single-modality brain region classification with the MLP classifier (Supplementary Figure 19), NEMO with the fine-tuned MLP either significantly outperforms or shows no significant difference compared to all other baselines.

1674 M.3 COMPARISON ACROSS LABEL RATIOS
1675

1676 We further compared classifier performance under varying label ratios. For both the linear classifier
1677 (Supplementary Figures 20 and 21) and MLP classifier (Supplementary Figures 22 and 23), NEMO
1678 significantly outperforms other baselines at each label ratio. Notably, with only 50% of the labels,
1679 the linear classifier with NEMO either significantly outperforms or shows no significant difference
1680 compared to all baselines using 100% of the labels. Similarly, the MLP classifier with NEMO
1681 achieves comparable or significantly better performance than all baselines using 100% of the labels
1682 with only 80% of the labels.

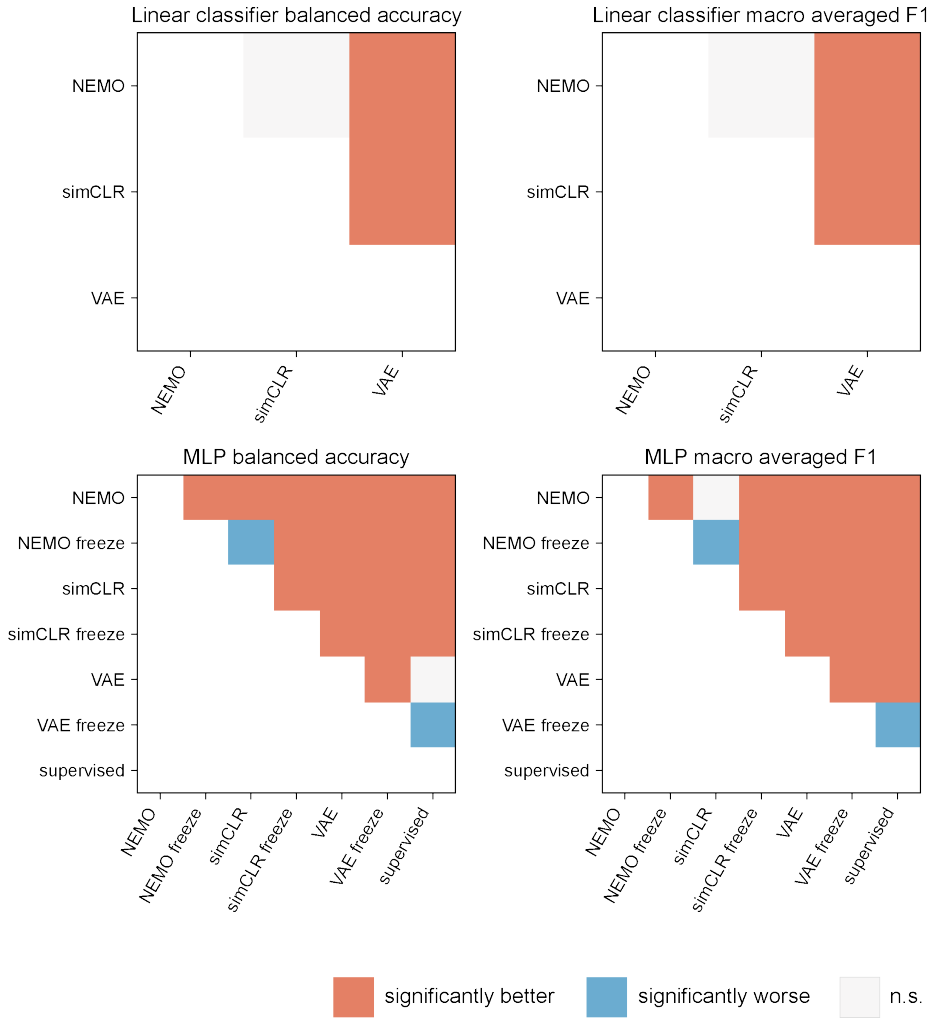
1683 Overall, these results demonstrate that NEMO consistently outperforms or matches the performance
1684 of other baselines across multiple evaluation criteria, including cell-type classification and brain re-
1685 gion decoding tasks. Its robustness is evident across classifiers, label ratios, and metrics, establishing
1686 NEMO as a highly effective model for these challenging neurobiological classification tasks.
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703



1726 Supplementary Figure 16: Clustering results from five NEMO instances trained with different random
1727 seeds, visualized using UMAP with consistent dimensionality reduction across runs. Colors
indicate distinct clusters.

1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

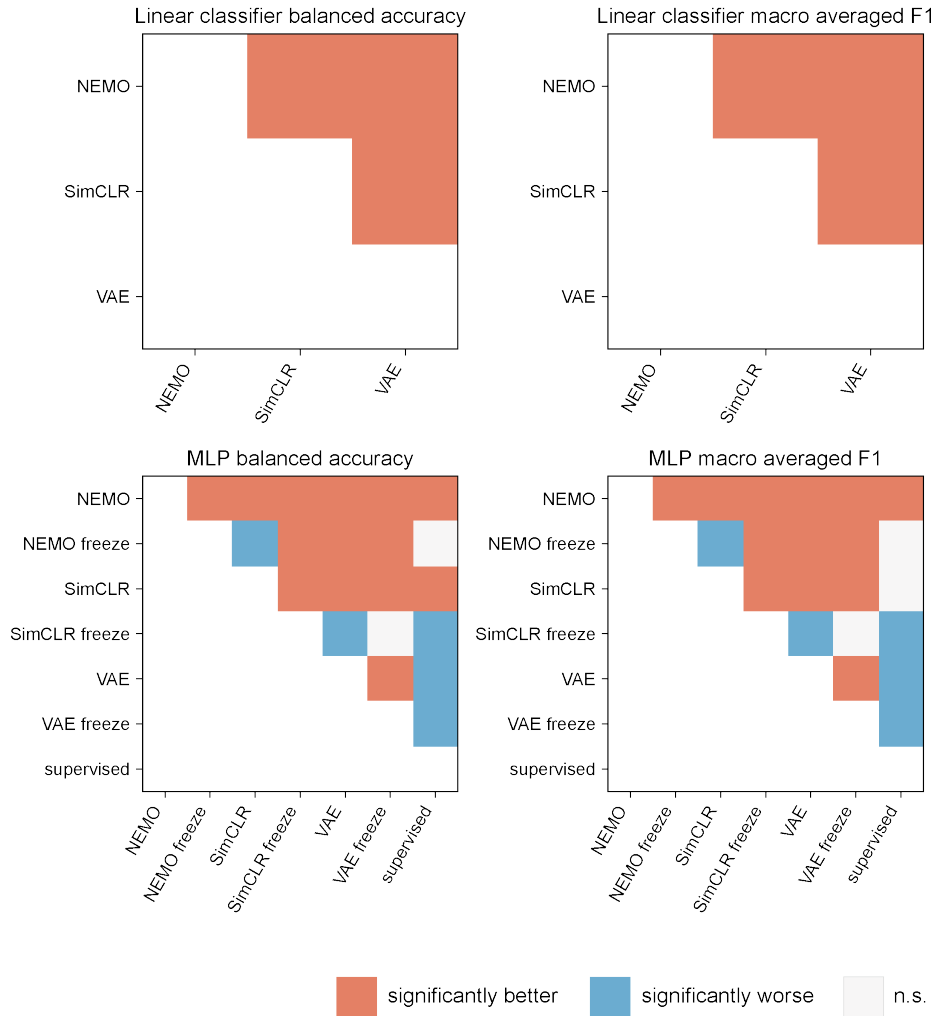
NP-Ultra cell-type classification



Supplementary Figure 17: One-tailed t-test results for NP-Ultra cell-type classification. For both linear and MLP classifiers, NEMO significantly outperforms or matches all baselines.

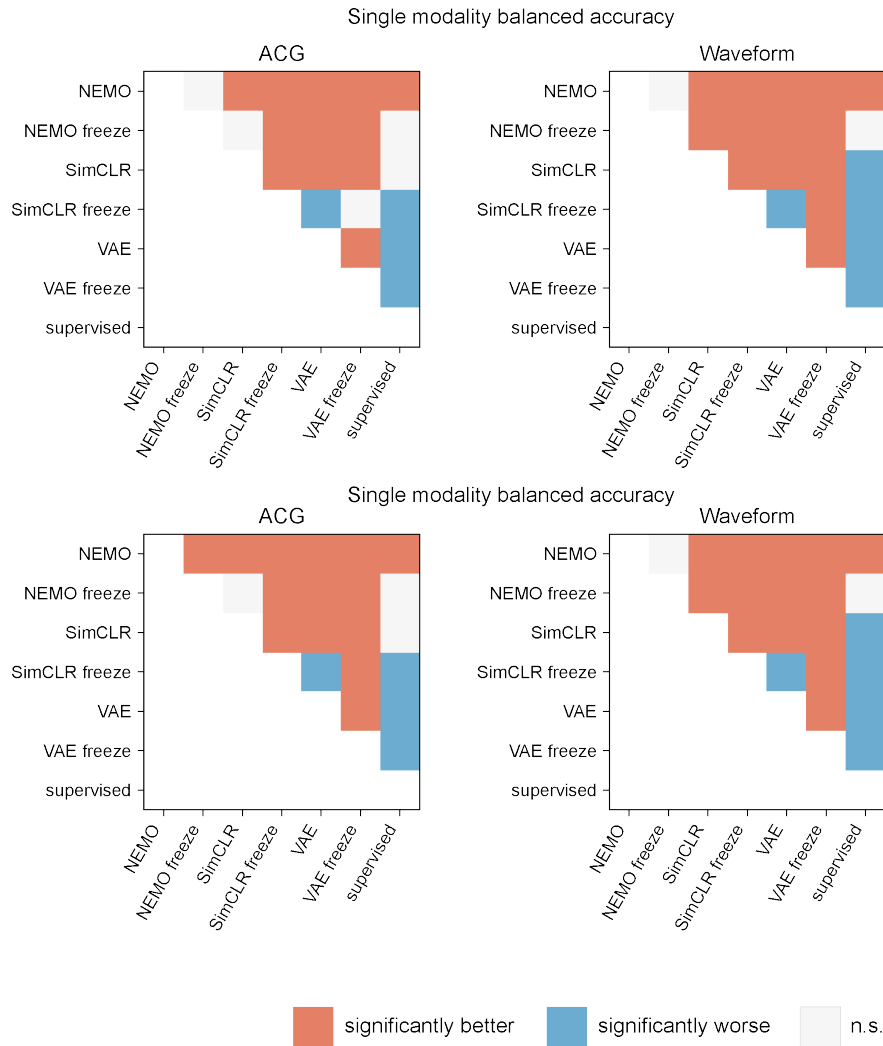
1782
 1783
 1784
 1785
 1786
 1787
 1788
 1789
 1790
 1791
 1792
 1793
 1794
 1795
 1796
 1797
 1798
 1799
 1800
 1801
 1802
 1803
 1804
 1805
 1806
 1807
 1808
 1809
 1810
 1811
 1812
 1813
 1814
 1815
 1816
 1817
 1818
 1819
 1820
 1821
 1822
 1823
 1824
 1825
 1826
 1827
 1828
 1829
 1830
 1831
 1832
 1833
 1834
 1835

IBL brain region classification




Supplementary Figure 18: One-tailed t-test results for IBL brain region classification. For both linear and MLP classifiers, NEMO significantly outperforms all baselines.

1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889




Supplementary Figure 19: One-tailed t-test results for IBL brain region classification with single modality and MLP classifier with end-to-end fine-tuning. For both modality, NEMO significantly outperforms or matches all baselines.

1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943




Supplementary Figure 20: One-tailed t-test results on balanced accuracy for IBL brain region classification with linear classifier across different label ratios. For all label ratios, NEMO significantly outperforms all other baselines. With only 50% of the labels, the linear classifier with NEMO either significantly outperforms or shows no significant difference compared to all baselines using 100% of the labels.

1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997




Supplementary Figure 21: One-tailed t-test results on macro averaged F1 score for IBL brain region classification with linear classifier across different label ratios. For all label ratios, NEMO significantly outperforms all other baselines. With only 50% of the labels, the linear classifier with NEMO significantly outperforms all baselines using 100% of the labels.

1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051



Supplementary Figure 22: One-tailed t-test results on balanced accuracy for IBL brain region classification with MLP classifier across different label ratios. For almost all label ratios (except 0.01), NEMO significantly outperforms all other baselines. With only 80% of the labels, the linear classifier with NEMO either significantly outperforms or shows no significant difference compared to all baselines using 100% of the labels.

2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105



Supplementary Figure 23: One-tailed t-test results on macro F1 score for IBL brain region classification with MLP classifier across different label ratios. For almost all label ratios (except 0.01), NEMO significantly outperforms all other baselines. With only 50% of the labels, the linear classifier with NEMO either significantly outperforms or shows no significant difference compared to all baselines using 100% of the labels.