Know Thyself by Knowing Others:

Learning Neuron Identity from Population Context (NuCLR)

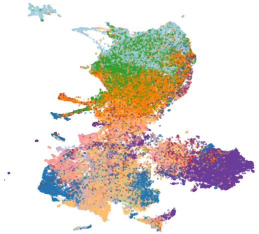

Authors:

Vinam Arora[1], Divyansha Lachi[1], Ian J. Knight[1],
Mehdi Azabou[2], Blake A. Richards[3], Cole L. Hurwitz[2],
Josh H. Siegle[4], Eva L. Dyer[1]

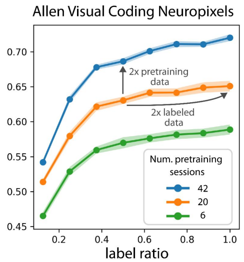

1 University of Pennsylvania
2 Columbia University
3 Mila
4 Allen Institute for Neural Dynamics

Neural Information Processing Systems
NeurIPS 2025

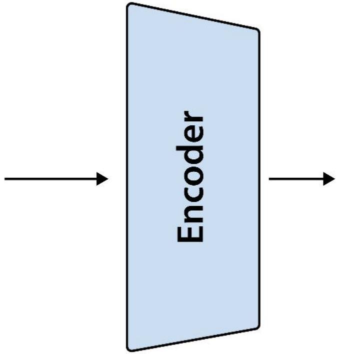

GOAL: Predict neuron-level properties directly from recorded activity.

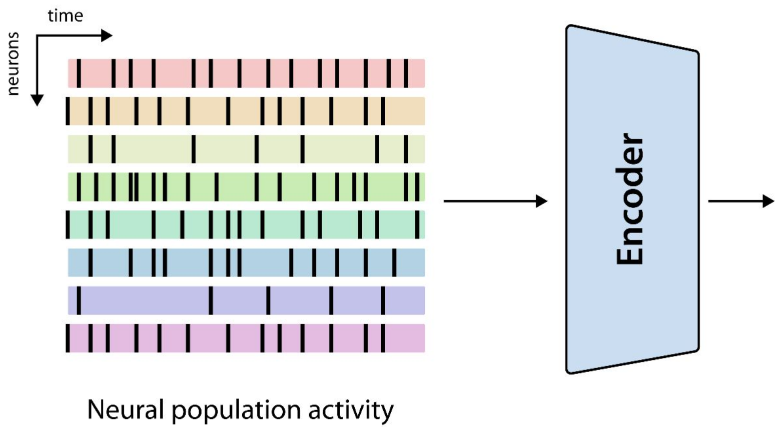

GOAL: Predict neuron-level properties directly from recorded activity.

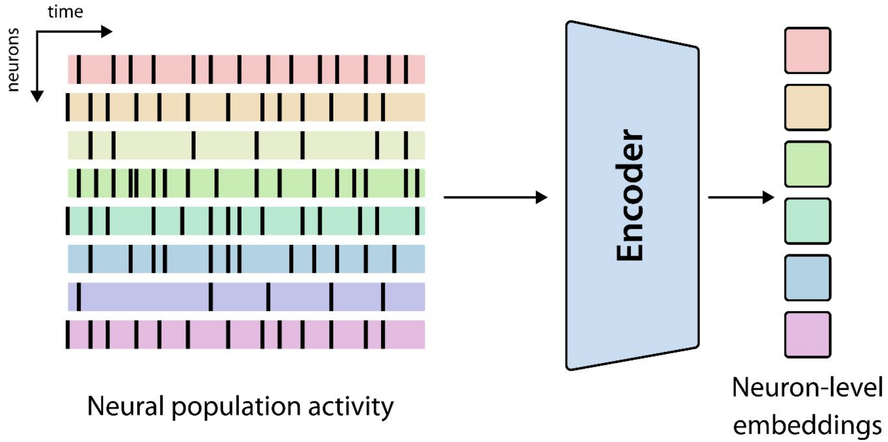

GOAL: Predict neuron-level properties directly from recorded activity.

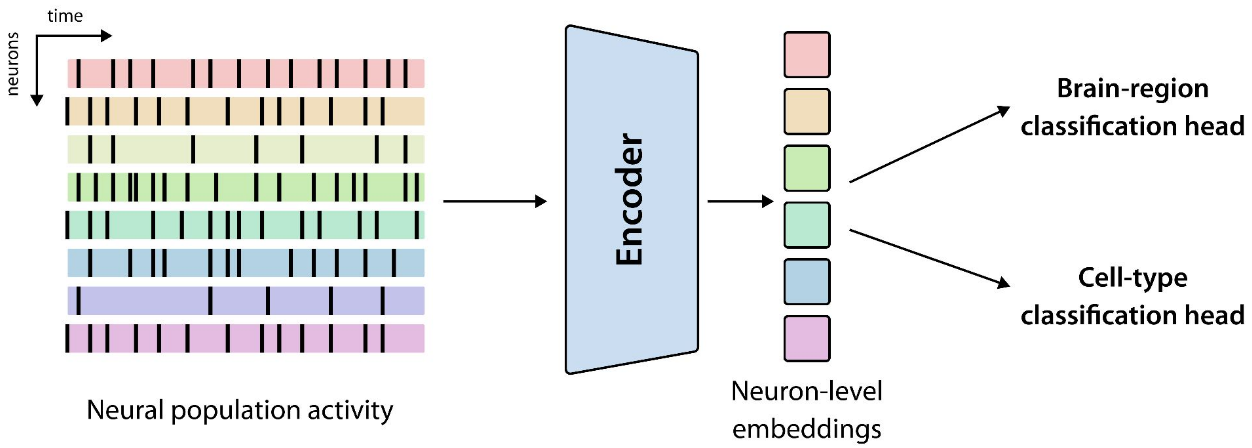

GOAL: Predict neuron-level properties directly from recorded activity.

Motivation

Neuroscientists can do more science
when they know which ***kind*** of neurons they recorded.

Motivation

Neuroscientists can do more science
when they know which ***kind*** of neurons they recorded.

Experimental methods are not scalable

Obtaining **Brain-Region** Information

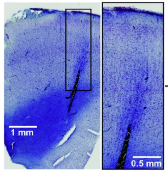

Coat probe with traceable ink

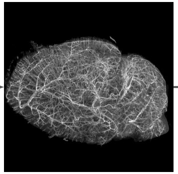

Brain clearing

Light sheet microscopy

Obtaining **Cell-type** Information

Genetically curated mice

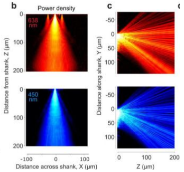

Special setups for optically stimulating target cell-types

Analysis

Motivation

Neuroscientists can do more science
when they know which ***kind*** of neurons they recorded.

Experimental methods are not scalable

Obtaining Brain-Region Information

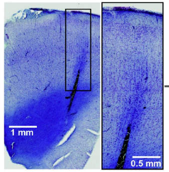

Coat probe with traceable ink

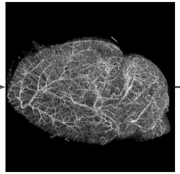

Brain clearing

Light sheet microscopy

Obtaining Cell-type Information

Genetically curated mice

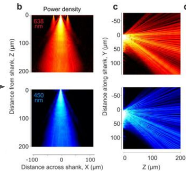

Special setups for optically stimulating target cell-types

Analysis

GOAL: Predict neuron-level properties directly from recorded activity.

Core Design Principles

Core Design Principles

1. Model interactions within the population

Cell-types express themselves in their neighborhood.

Neurons from the same brain-region have more correlated activity.

Core Design Principles

1. Model interactions within the population

Cell-types express themselves in their neighborhood.
Neurons from the same brain-region have more correlated activity.

2. Neuron identity is constant

A neuron's properties remain stable across time and stimuli.

Model Architecture

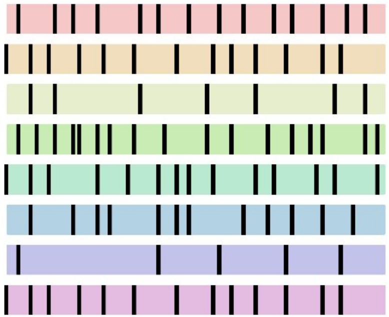

Model Architecture

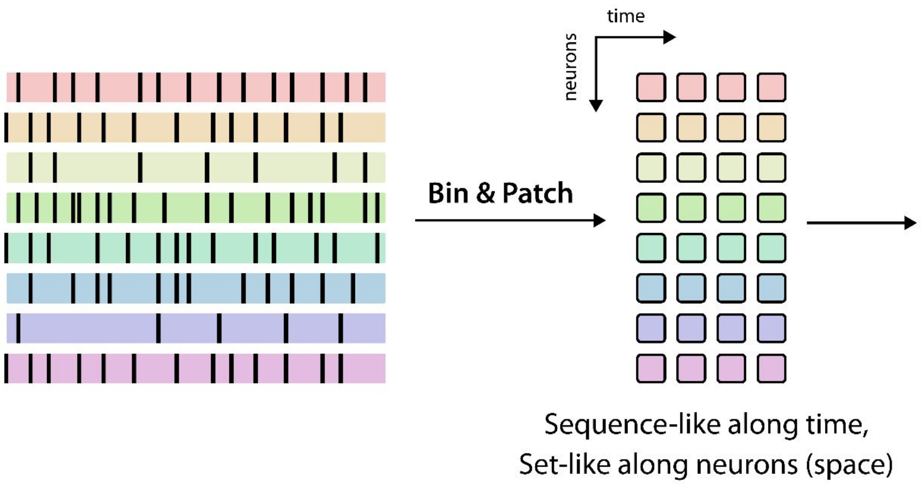

Model Architecture

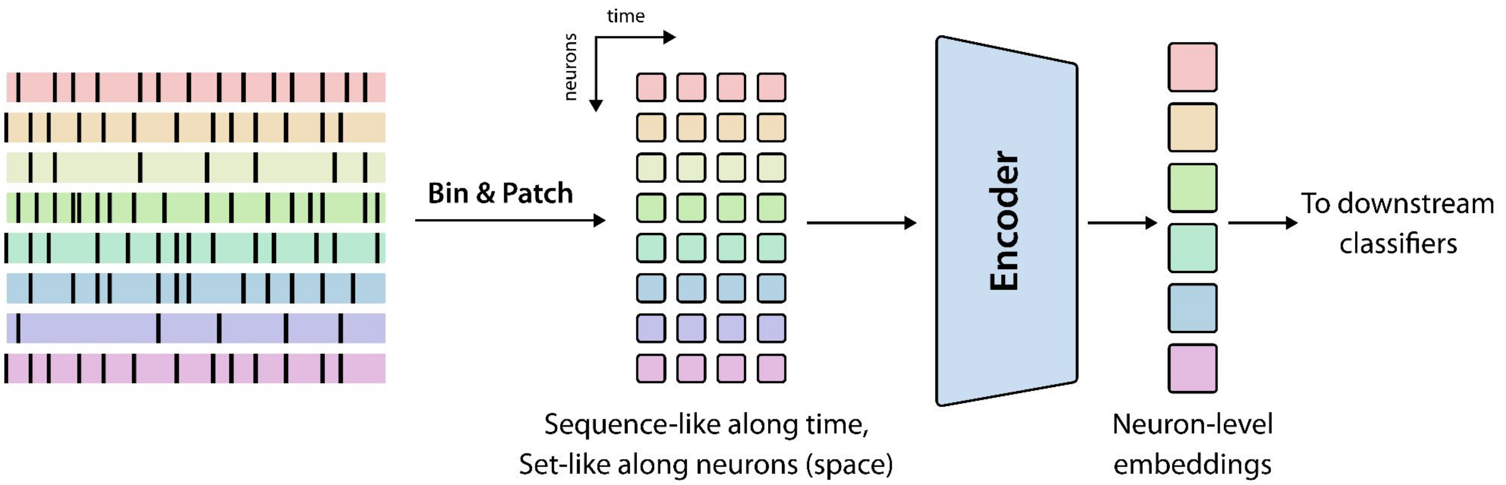

Model Architecture

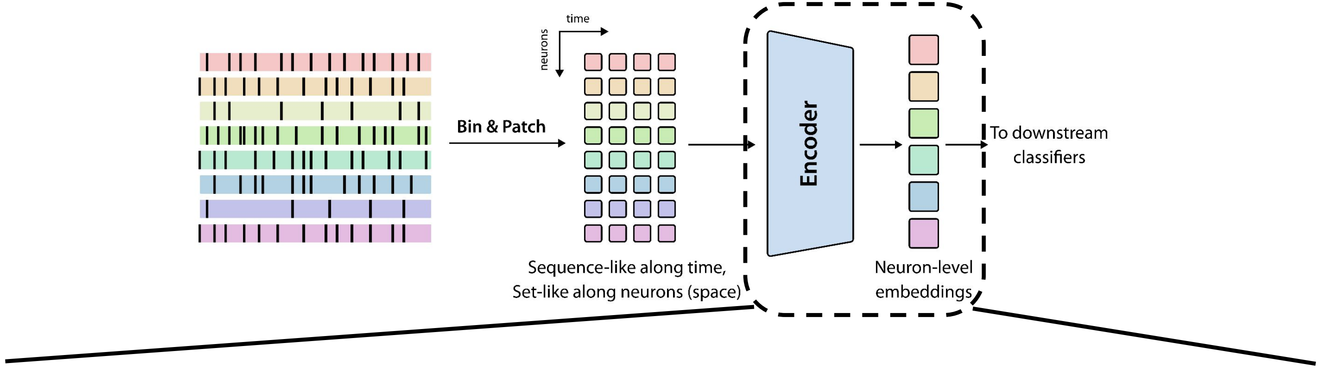

Model Architecture

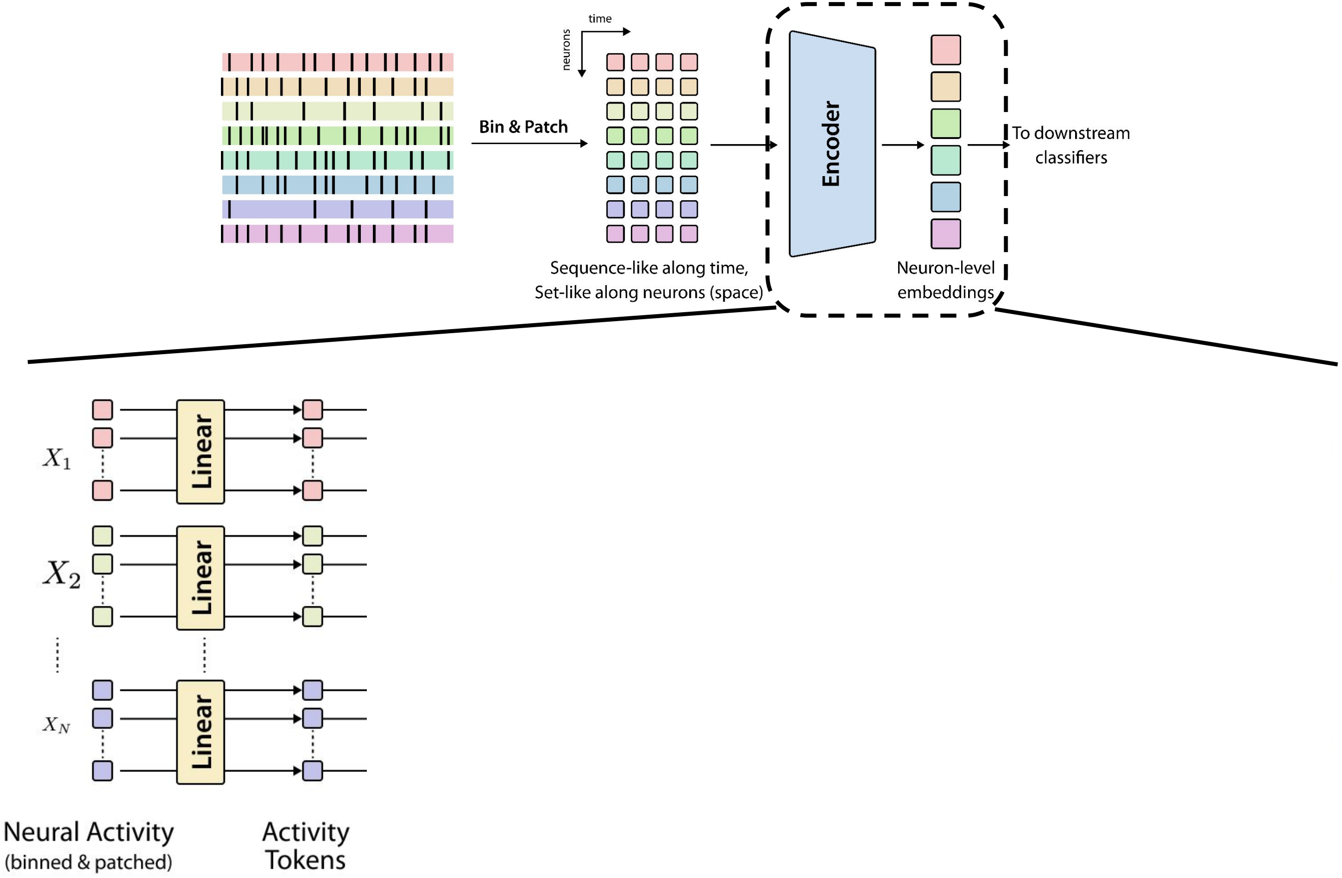

Model Architecture

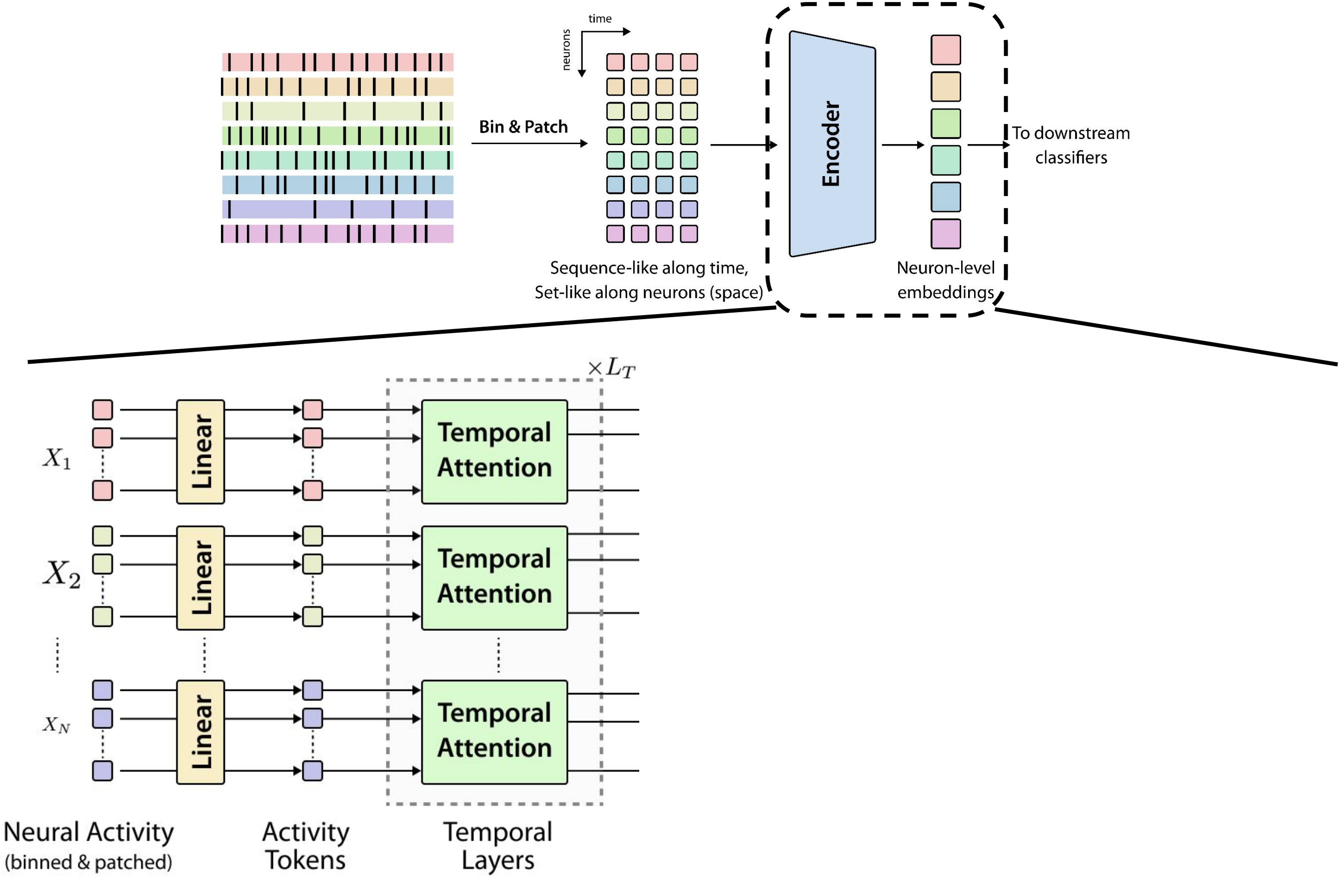

Model Architecture

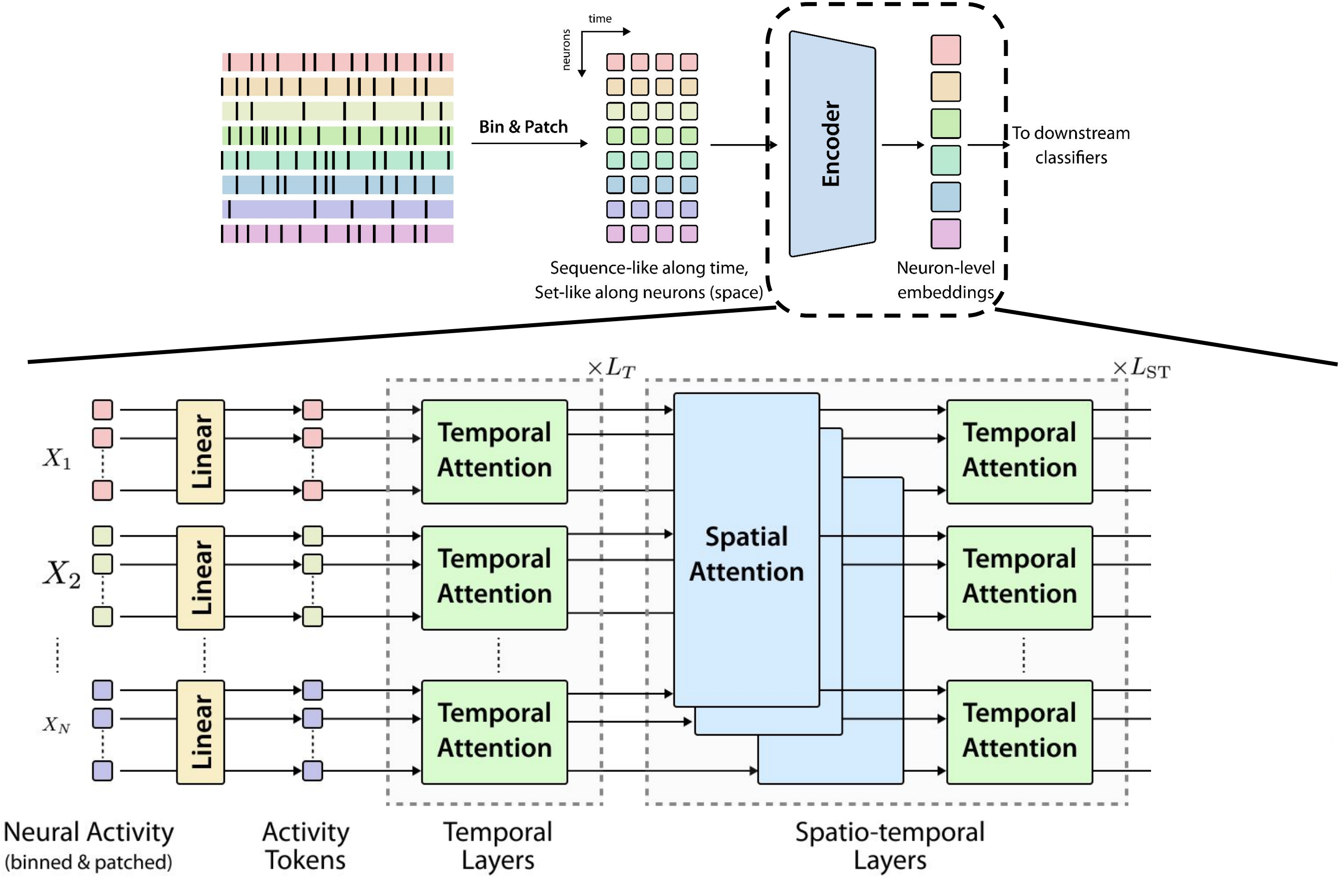

Model Architecture

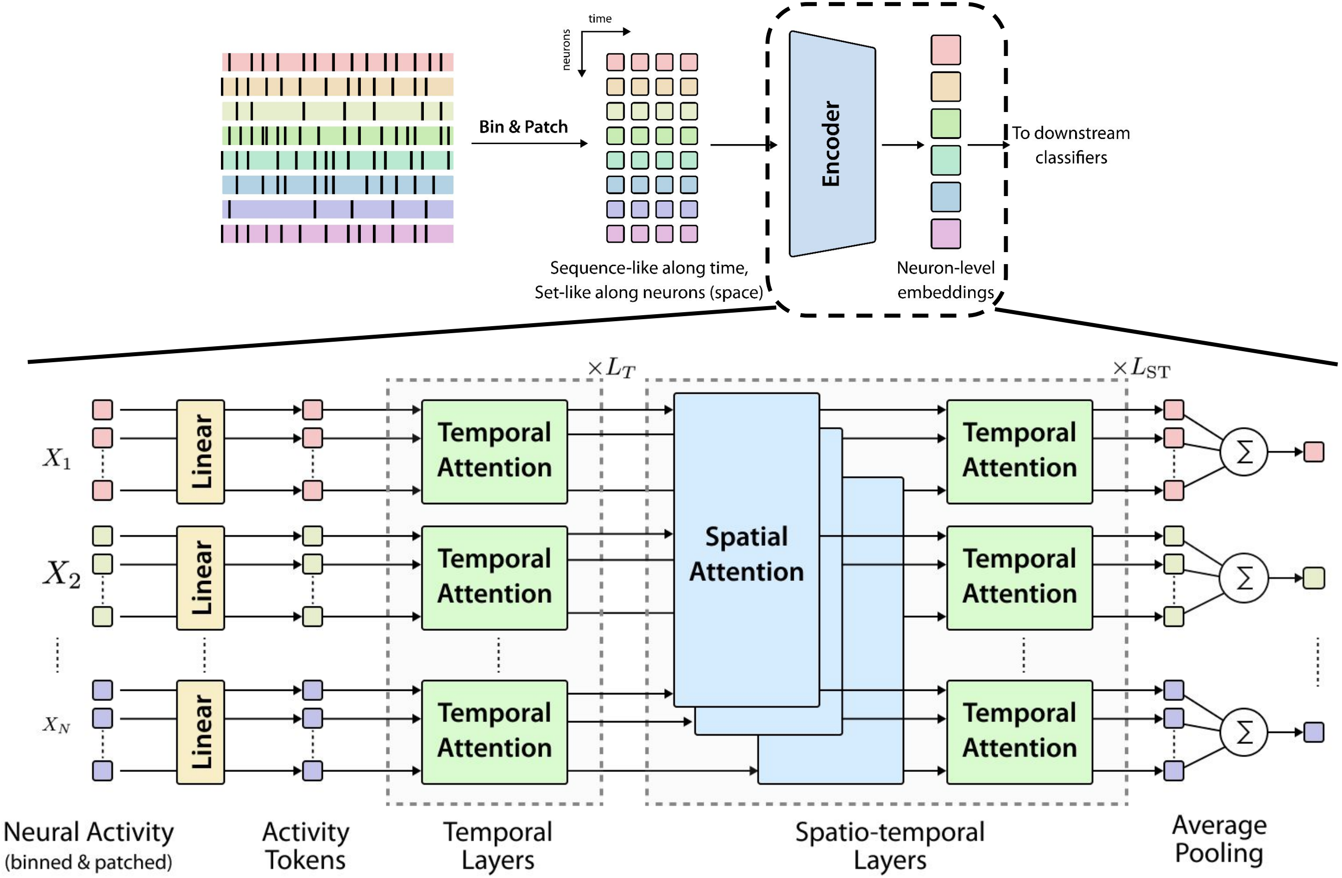

Contrastive Training Objective

Contrastive Training Objective

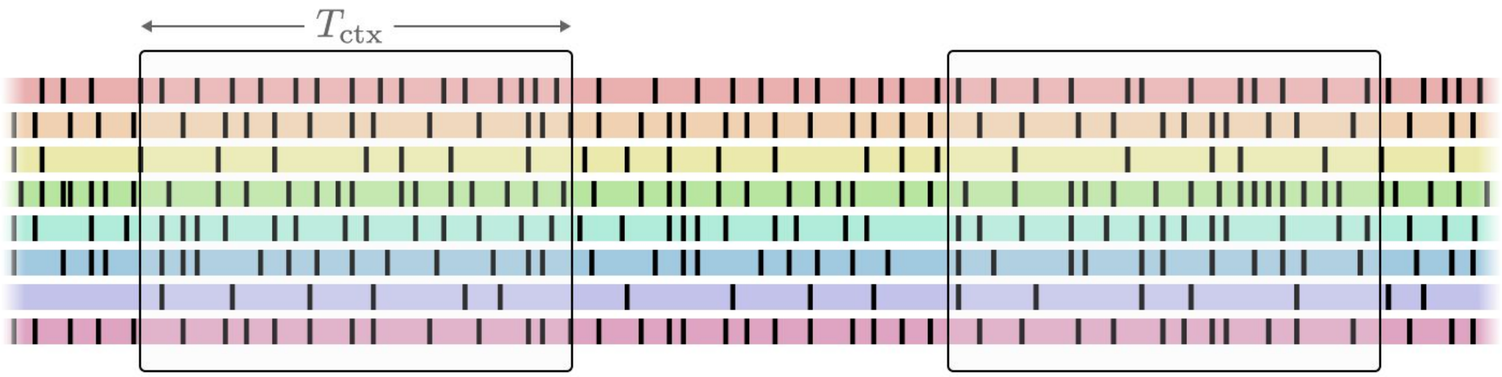

Contrastive Training Objective

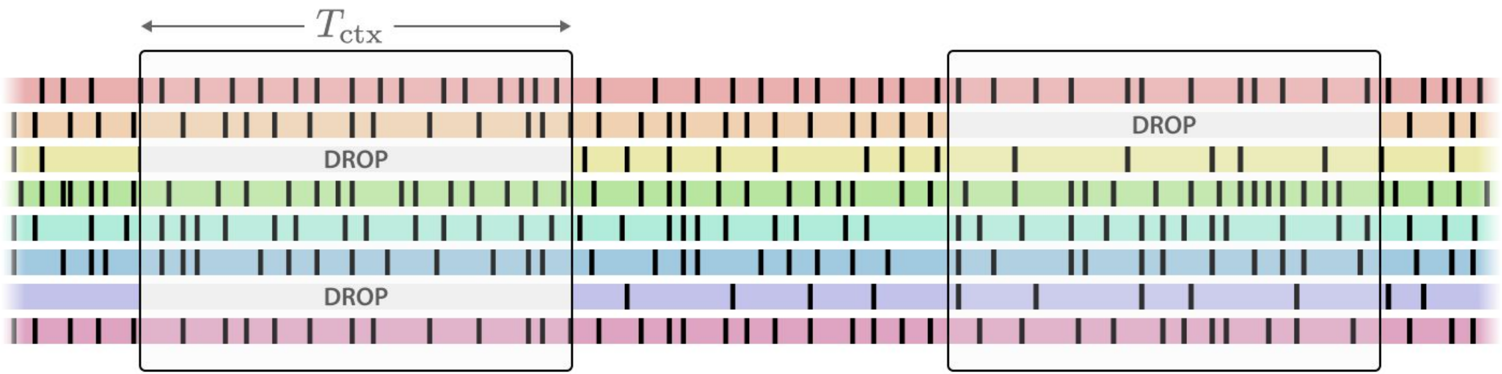

Contrastive Training Objective

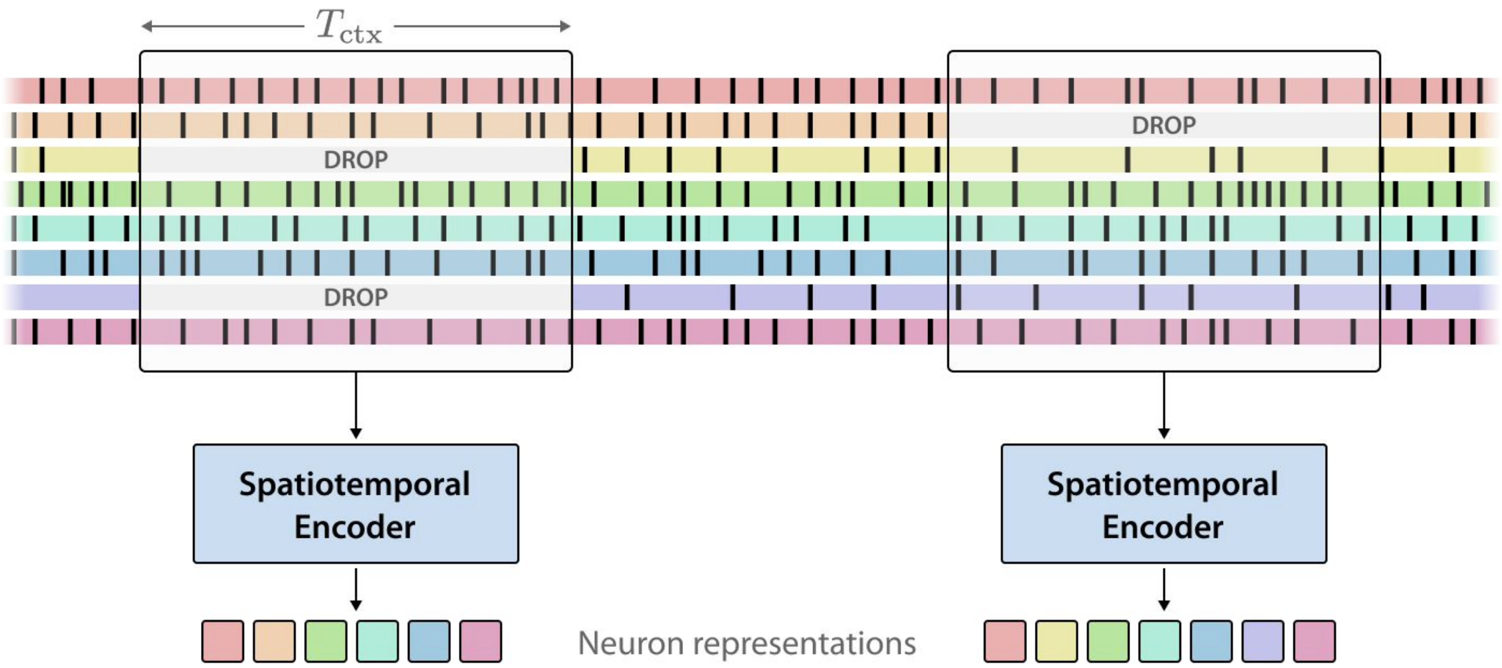

Contrastive Training Objective

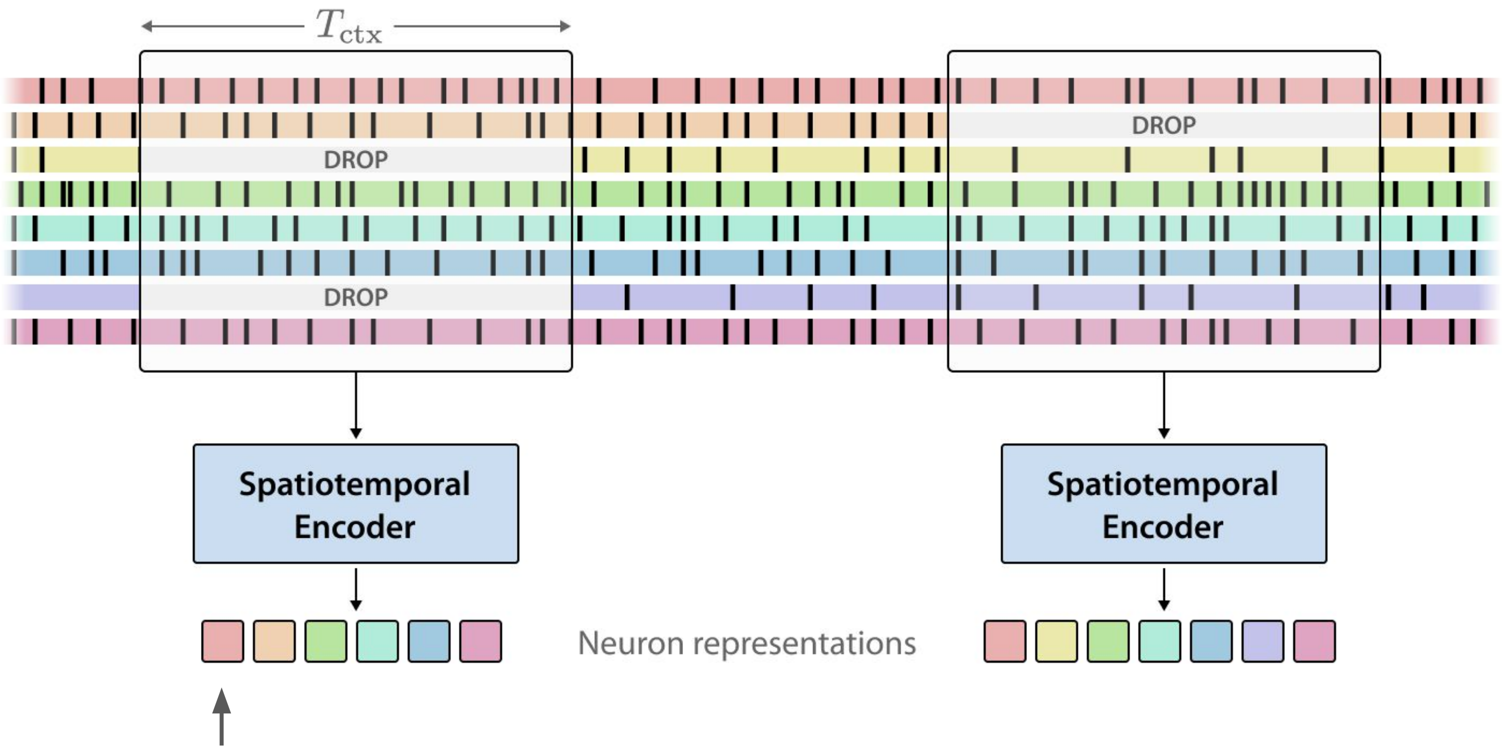

Contrastive Training Objective

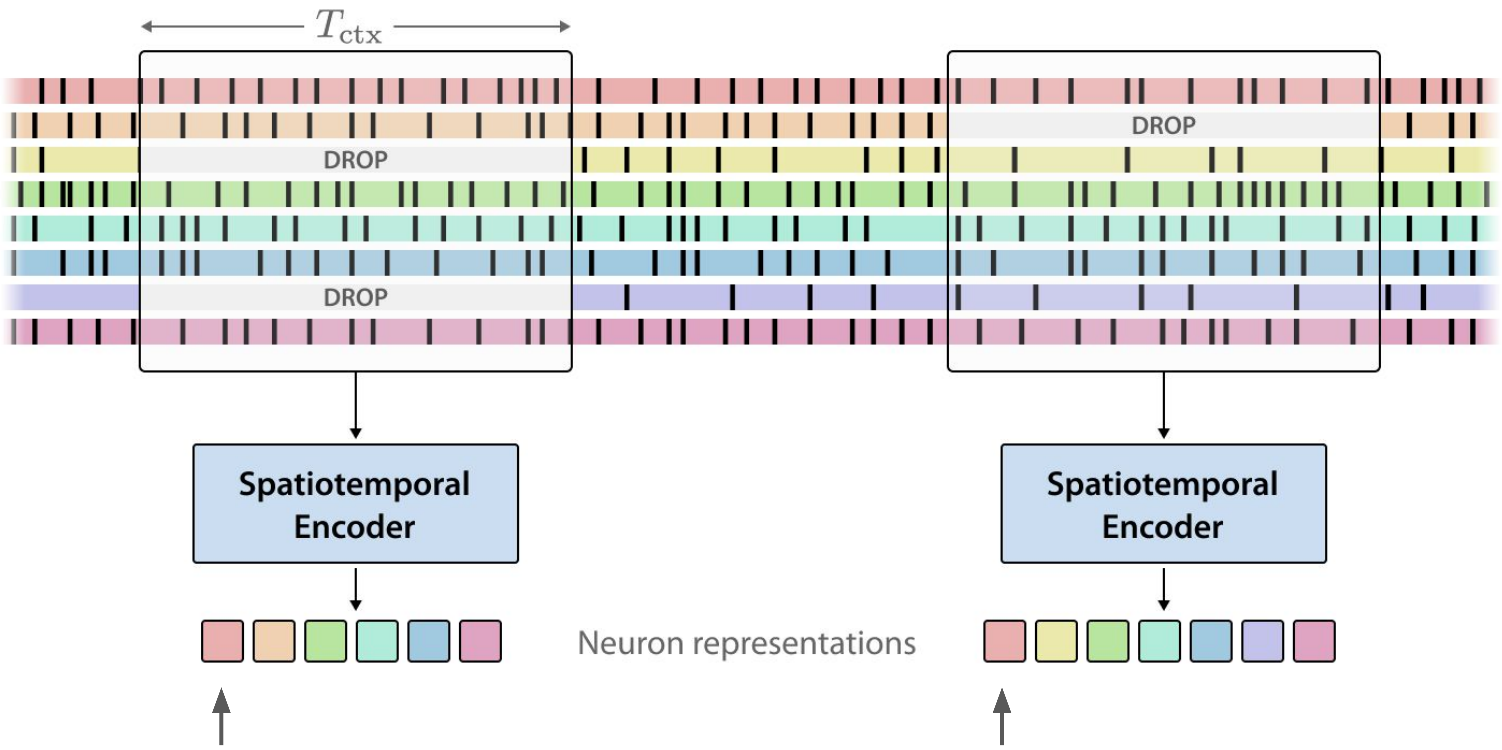

Contrastive Training Objective

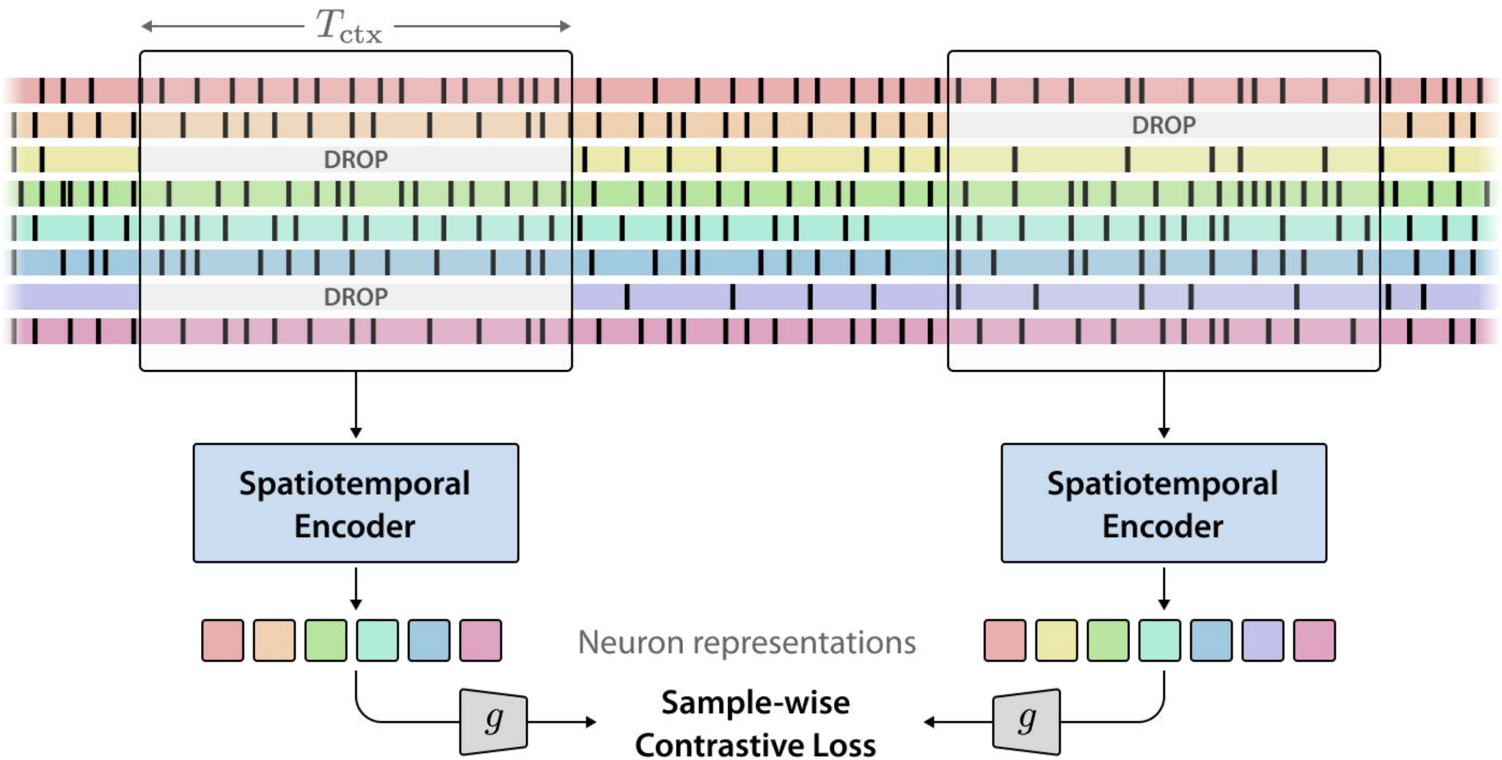

Downstream Tasks

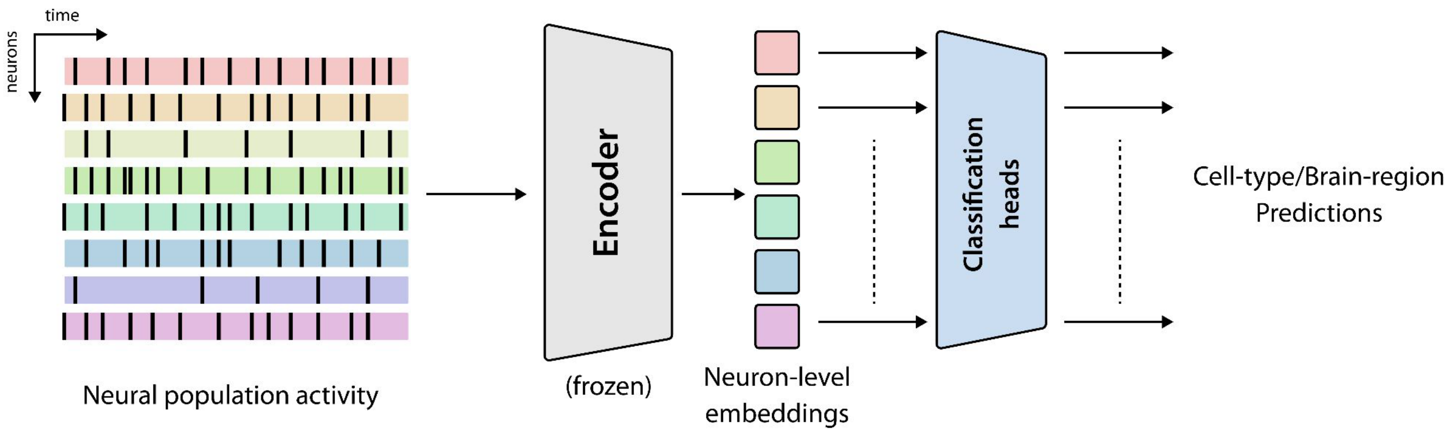

Encoder is frozen, and a classification head (usually Linear) is trained on top of the embeddings.

NuCLR achieves SoTA prediction

Cell-Type Prediction

	Dataset	# Classes	Setting	NuCLR	NeuPRINT	NEMO	LOLCAT	Δ wrt best baseline
Neuropixels	Allen VC	3	Transductive zero-shot	0.7218 ± 0.0113	0.3999 ± 0.0312	0.4256 ± 0.0114	N/A	+0.2962
			Inductive zero-shot	0.7200 ± 0.0267	N/A	0.4194 ± 0.0099	0.4121 ± 0.0800	+0.3006
Ca^{2+}	Bugeon et. al. (E vs. I)	2	Transductive	0.8080 ± 0.0007	0.6702 ± 0.0004	N/A	0.7205 ± 0.0127	+0.0875
			Transductive zero-shot	0.6979 ± 0.0368	0.6424 ± 0.0031	N/A	N/A	+0.0555
			Inductive zero-shot	0.7009 ± 0.0269	N/A	N/A	0.7463 ± 0.0095	-0.0454
	Bugeon et. al. (Subclass)	5	Transductive	0.6088 ± 0.0307	0.4949 ± 0.0260	N/A	0.2900 ± 0.0388	+0.1139
			Transductive zero-shot	0.3890 ± 0.0144	0.3385 ± 0.0221	N/A	N/A	+0.0505
			Inductive zero-shot	0.4440 ± 0.0495	N/A	N/A	0.2418 ± 0.0078	+0.2022

Brain-Region Prediction

	Dataset	# Classes	Setting	NuCLR	NeuPRINT	NEMO	LOLCAT	Δ wrt best baseline
Neuropixels	IBL	10	Transductive	0.6686 ± 0.0034	0.2684 ± 0.0285	0.4188 ± 0.0041	0.2851 ± 0.0008	+0.2498
			Transductive zero-shot	0.5343 ± 0.0115	0.2464 ± 0.0240	0.3804 ± 0.0011	N/A	+0.1539
			Inductive zero-shot	0.5295 ± 0.0040	N/A	0.3793 ± 0.0011	0.2532 ± 0.0016	+0.3006
	Steinmetz et. al.	5	Transductive	0.9594 ± 0.0027	0.4315 ± 0.0139	0.6989 ± 0.0044	0.3205 ± 0.0055	+0.2605
			Transductive zero-shot	0.7338 ± 0.0226	0.3832 ± 0.0378	0.6681 ± 0.0016	N/A	+0.0657
			Inductive zero-shot	0.5810 ± 0.0110	N/A	0.5595 ± 0.0048	0.3191 ± 0.0311	+0.0215

NuCLR achieves SoTA prediction

Cell-Type Prediction

	Dataset	# Classes	Setting	NuCLR	NeuPRINT	NEMO	LOLCAT	Δ wrt best baseline
Neuropixels	Allen VC	3	Transductive zero-shot	0.7218 ± 0.0113	0.3999 ± 0.0312	0.4256 ± 0.0114	N/A	+0.2962
			Inductive zero-shot	0.7200 ± 0.0267	N/A	0.4194 ± 0.0099	0.4121 ± 0.0800	+0.3006
Ca^{2+}	Bugeon et. al. (E vs. I)	2	Transductive	0.8080 ± 0.0007	0.6702 ± 0.0004	N/A	0.7205 ± 0.0127	+0.0875
			Transductive zero-shot	0.6979 ± 0.0368	0.6424 ± 0.0031	N/A	N/A	+0.0555
			Inductive zero-shot	0.7009 ± 0.0269	N/A	N/A	0.7463 ± 0.0095	-0.0454
	Bugeon et. al. (Subclass)	5	Transductive	0.6088 ± 0.0307	0.4949 ± 0.0260	N/A	0.2900 ± 0.0388	+0.1139
			Transductive zero-shot	0.3890 ± 0.0144	0.3385 ± 0.0221	N/A	N/A	+0.0505
			Inductive zero-shot	0.4440 ± 0.0495	N/A	N/A	0.2418 ± 0.0078	+0.2022

Brain-Region Prediction

	Dataset	# Classes	Setting	NuCLR	NeuPRINT	NEMO	LOLCAT	Δ wrt best baseline
Neuropixels	IBL	10	Transductive	0.6686 ± 0.0034	0.2684 ± 0.0285	0.4188 ± 0.0041	0.2851 ± 0.0008	+0.2498
			Transductive zero-shot	0.5343 ± 0.0115	0.2464 ± 0.0240	0.3804 ± 0.0011	N/A	+0.1539
			Inductive zero-shot	0.5295 ± 0.0040	N/A	0.3793 ± 0.0011	0.2532 ± 0.0016	+0.3006
	Steinmetz et. al.	5	Transductive	0.9594 ± 0.0027	0.4315 ± 0.0139	0.6989 ± 0.0044	0.3205 ± 0.0055	+0.2605
			Transductive zero-shot	0.7338 ± 0.0226	0.3832 ± 0.0378	0.6681 ± 0.0016	N/A	+0.0657
			Inductive zero-shot	0.5810 ± 0.0110	N/A	0.5595 ± 0.0048	0.3191 ± 0.0311	+0.0215

NuCLR scales with pretraining data size

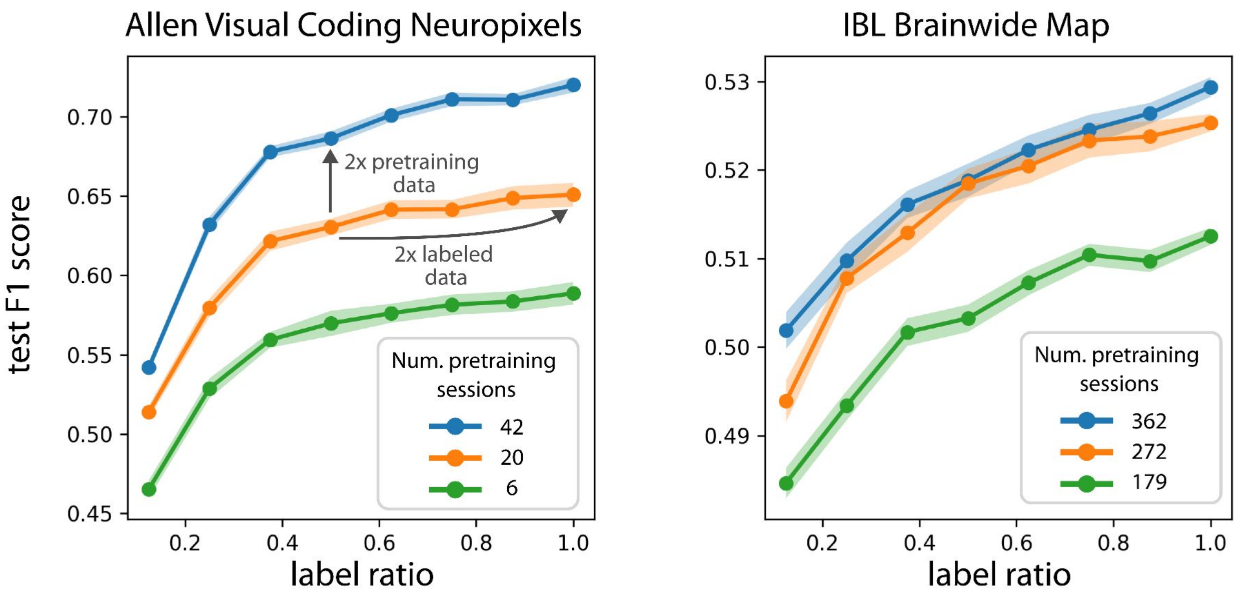

NuCLR scales with pretraining data size

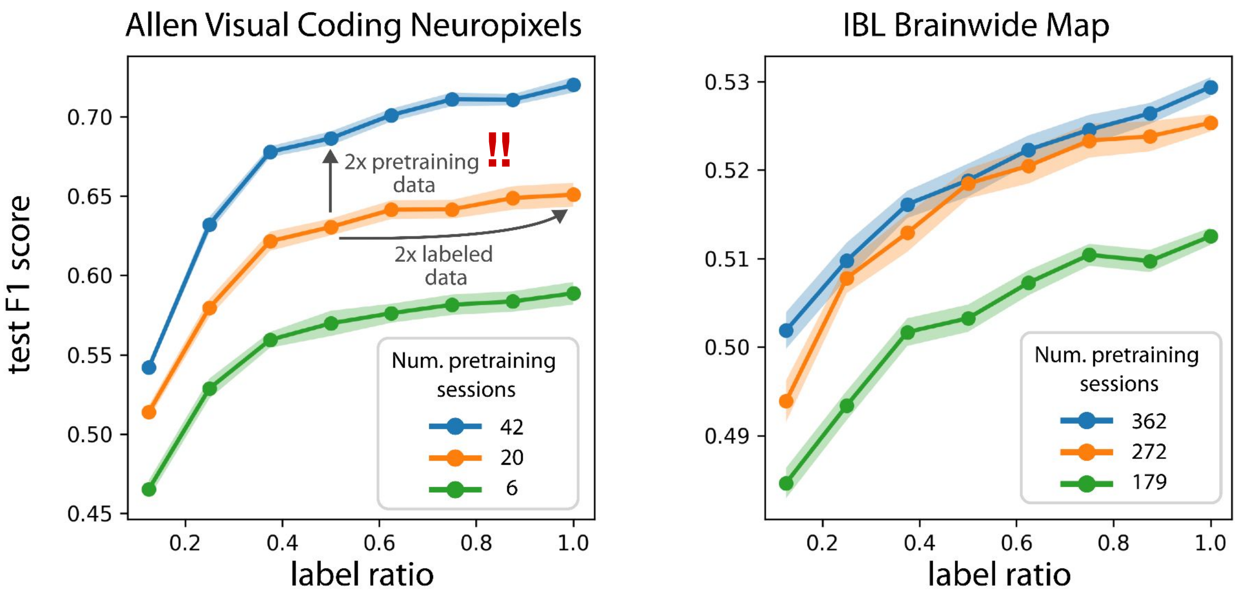

In some situations, it **can be more effective to collect more unlabelled (pre-training) data** than increasing the number of tagged neurons.

Conclusions

- Analyzing just neural activity can take you pretty far.
- Population context is quite important.
- NuCLR is a promising and scalable approach for learning neuron representations.

Thank you!

More about me: **`https://vinam.dev`**

Feel free to reach out at: **`vinam@seas.upenn.edu`**