Neuron Intrinsic Representation From Dynamics Using Contrastive learning

Abstract

Unlocking the secrets of neuronal activity—the information conveyed by neural dynamics—remains one of the grand challenges of neuroscience. The Platonic Representation Hypothesis posits that behind different modalities of data (what we sense or detect), there exists a universal, modality-independent representation of reality. Inspired by this, we treat each neuron as a system, where we can detect the neuron's multi-segment activity data under different peripheral conditions. We believe that, similar to the Platonic idea, there exists a time-invariant representation behind the different segments of the same neuron, which reflects the intrinsic properties of the neuron's system. The optimization objective for obtaining the intrinsic representation of neurons should satisfy two criteria: (I) segments from the same neuron should have a higher similarity than segments from different neurons; (II) the representations should generalize well to out-of-domain data. To achieve this, we employ contrastive learning, treating different segments from the same neuron as positive pairs and segments from different neurons as negative pairs. During the implementation, we chose the VICReg, which uses only positive pairs for optimization but indirectly separates dissimilar samples via regularization terms. To validate the efficacy of our method, we first applied it to simulated neuron population dynamics data generated using the Izhikevich model. We successfully confirmed that our approach captures the intrinsic properties of each neuron as defined by preset hyperparameters. We then applied our method to two real-world neuron dynamics datasets, including spatial transcriptomics-derived neuron type annotations and the brain regions where each neuron is located. The learned representations from our model not only predict neuron type and location but also show robustness when tested on out-of-domain data (unseen animals). This demonstrates the potential of our approach in advancing the understanding of neuronal systems and offers valuable insights for future neuroscience research.

1 Introduction

Unraveling the intricacies of neuronal activity and the information encoded within neural dynamics stands as a monumental challenge in the field of neuroscience. Plato's cave allegory and Platonic Representation Hypothesis suggests the existence of a universal, modality-independent representation of world that transcends the modalities through which we perceive it. Drawing inspiration from this philosophical concept, we propose a novel perspective on neuronal activity, treating each neuron as a distinct 'world' capable of generating multi-segment activity data under various peripheral conditions, the multi-segment activity data of individual neuron as what we perceive.

Our approach is predicated on the belief that, akin to the Platonic ideal, there exists a time-invariant representation that underlies the diverse segments of activity within the same neuron. This representation is posited to encapsulate the intrinsic properties of the neuron's system, providing a stable and consistent framework for understanding neuronal function.

To extract this intrinsic representation, we have formulated an optimization objective that adheres to two criteria. Firstly, segments of activity from the same neuron should exhibit a higher degree of similarity compared to segments from different neurons. This criterion ensures that our method can effectively distinguish between the unique signatures of individual neurons. Secondly, the represen-

tations derived from our method should possess a high degree of generalizability, extending their applicability to out-of-domain data, which includes neuronal activity data from different species or experimental conditions that were not part of the training set.

To achieve these objectives, we have employed contrastive learning, a powerful technique in machine learning that involves treating segments from the same neuron as positive pairs and segments from different neurons as negative pairs. This method leverages the contrast between similar and dissimilar samples to learn an effective representation. In our implementation, we have chosen VICReg, a variant of contrastive learning that optimizes solely on positive pairs but incorporates regularization terms to indirectly separate dissimilar samples. This approach ensures that the learned representations are not only discriminative but also robust to variations in the data. However, achieving the above requires the following two innovative designs for neuron data: Firstly, we use CEBRASchneider et al. (2023b) to integrate the single neuronal peripheral information (such as activity of neighboring neuronal populations and behavioral data for each segment of a single neuron). This process encodes the peripheral information associated with each segment of an individual neuron. CEBRA, which has been widely demonstrated to provide high-performance learnable latent embeddings for joint representation of surrounding neuronal data and external information, but before, it was applied to groups of neurons instead of considering it from the perspective of a single neuron. Secondly, aims to output a consistent representation of a neuron's identity across varying experimental conditions, utilizing these conditions as positive instances in a contrastive learning paradigm. While multi-session data can readily supply such instances, the reality is that many neurons are recorded in a single session only. To address this issue, we design a data augmentation method specific to neuronal data. This method involves extracting segments of varying lengths from the same session to serve as positive sample pairs. In this way, the learned representation can capture intrinsic properties across different time scales.

To rigorously assess the effectiveness of the representations learned by , we model neuron population dynamic data using the Izhikevich model Izhikevich (2003), where different neurons are assigned distinct hyperparameters representing different firing modes as intrinsic properties. These neurons are randomly connected to form a network, and after stimulating the network, we obtain neuron population data. The representations learned by through self-supervised learning on this neuron population data have been confirmed to contain hyperparameter information. In recognition of the noise and complexity in real-world neuronal datasets, we turn to a publicly available dataset of mouse brain neuron populations. These neurons, after recording dynamic activities, are separated and labeled with cell type annotations using transcriptomics. The representations learned by on this dataset can be efficiently utilized for downstream cell type prediction tasks. In addition, using another real public dataset containing ten mice and 39 datasets, we demonstrate that the intrinsic information contained in the representation obtained by is consistent across animals, and that this representation has strong cross-domain capability when performing downstream tasks.

2 RELATED WORK

Single Neuron Models: Single neuron models are essential in understanding the fundamental properties of neuronal dynamics and behavior. The Leaky Integrate-and-Fire model Liu & Wang (2001), developed in the 1950s, is a minimalistic model that simulates the integration of synaptic inputs and the leakage of membrane potential over time. The Hodgkin-Huxley model Nelson & Rinzel (1995), introduced in 1952, provides a detailed description of action potential generation using complex equations to represent ionic currents across the neuronal membrane, offering deep insights into neuronal excitability. The FitzHugh-Nagumo model Izhikevich & FitzHugh (2006), proposed in 1961, simplifies the Hodgkin-Huxley model to focus on excitability and action potential dynamics while reducing computational complexity. Finally, the Izhikevich model, developed in 2003, combines simplicity with versatility, effectively capturing a wide range of neuronal firing patterns and balancing computational efficiency with biological realism. The hyperparameters set for these models can be viewed as prior assumptions about the inherent properties of the neurons. In the subsequent sections, we select the Izhikevich model to generate a synthetic dataset for testing , using the set hyperparameters as the true labels.

Neural Latent Representation Learning: Neural Latent Representation Learning has been pivotal in transforming high-dimensional neuronal data into lower-dimensional embeddings that en-

capsulate instantaneous information Bengio et al. (2013). A remarkable impact has been made in neuroscience —from the linear dimensionality reduction techniques such as Principal Component Analysis (PCA) Maćkiewicz & Ratajczak (1993)—to the non-linear visualization methods like Uniform Manifold Approximation and Projection (UMAP) McInnes et al. (2018) and t-Distributed Stochastic Neighbor Embedding (t-SNE) Kobak & Berens (2019)—and most recently, to the advanced, data-driven deep learning strategies. The focus was first on reducing the dimensionality of neuronal data alone. It later expanded to include joint dimensionality reduction with behavioral information and external stimuli (e.g., pi-VAE Prakash (2024)). CEBRA represents a culmination of these advancements, integrating various techniques into a unified framework. CEBRA is used in this study to integrate the stimuli information experienced by each neuron as input for . Neuprint is a method used to extract invariant information from neurons Mi et al. (2023); however, it still adheres to traditional neuron modeling approaches, which can only implicitly represent neurons. This results in challenging training processes and suboptimal representation performance.

Contrastive Learning for Voice Representations: The basic approach of the solution presented in this paper draws inspiration from similar tasks, such as extracting inherent representations of speakers from voice data Torres et al. (2024). This type of work has been implemented on song-singer datasets, where contrastive learning is used to bring together the voices of the same singer while pushing apart the voices of different singers. However, while we borrowed the underlying idea, the specific methods had to be uniquely designed to accommodate the characteristics of neural data.

3 METHOD

3.1 GOAL

Our goal is to develop a method for learning intrinsic neuron representations on neuron population data. These representations should fulfill three key criteria: (I) Neurons with similar functional roles should exhibit higher similarity in their representations compared to those with different roles; (II) The learned representations should be robust to variations in neuronal activity patterns due to different stimuli or environmental conditions; and (III) The representations should be adaptable and generalize effectively to new and unseen neuronal activity patterns.

3.2 ARCHITECHTURE

The ideal embedding space for neuron representations should cluster activity segments of the same neuron while also ensuring semantic consistency by placing similar neurons close to each other within the space. In line with the criteria outlined in Section 3.1, We conducted advanced contrastive learning loss functions, VICReg Bardes et al. (2021). We also carefully designed the data sampling methods to generate multi-segments activity data for each neuron for training purposes.

Data sampling: Without loss of generality, we take bi-segment data as an example. Consider two scenarios: 1) If a neuron has recordings from multiple sessions, we randomly extract two segments (X_{seg}, X'_{seg}) from different sessions. 2) If a neuron has recordings from only a single session, we randomly extract two non-overlapping segments (X_{seg}, X'_{seg}) from that session. We repeat this process B times for a batch size of B, obtaining a positive pair batch $(X_{seg}^{(1)}, X_{seg}^{(2)})$.

Model: We should integrate surrounding information into each segment. In neurobiology, the stimulation received by the neuron corresponding to a segment (X_{st}), the animal's behavior at the corresponding time (X_{be}), and the session information (X_{se}) are all the surrounding information (SI) of this segment. CEBRA is a self-supervised learning algorithm designed to obtain interpretable and consistent embeddings of high-dimensional recordings using auxiliary variables. In simple terms, it serves as an encoder for the dynamic activity information of neuronal populations. Its advantage lies in its ability to encode both neuronal activity data (N*T) and corresponding auxiliary variables, such as behavior and external stimuli (M*T), into a lower-dimensional representation (D*T), where D represents the latent space dimensions. In CEBRA, D does not correspond to individual neurons, but rather provides a dimensionality reduction method that combines high-dimensional temporal activity data of neuronal populations with corresponding auxiliary variable information, aiming to uncover the relationships between neuronal activity and these variables. Our work focuses on ob-

Figure 1: Optimization objective for obtaining intrinsic representations of neurons as follows: clips(segments) from the same neuron should have a higher average similarity than clips from different neurons. Figure 1 illustrates how this objective can be achieved using a contrastive learning approach, where different clips from the same neuron are treated as positive pairs and clips from different neurons are treated as negative pairs. Each neuron is considered separately, and the activity information from the remaining neurons in the population is treated as peripheral information, which acts as auxiliary variables. The peripheral information for each neuron is processed and encoded using CEBRA.

taining time-invariant representations for individual neurons. We cleverly utilize CEBRA from a different perspective as a preprocessing step for our input data. From the viewpoint of individual neurons, we use CEBRA to integrate the single neuronal peripheral information (such as activity of neighboring neuronal populations and behavioral data) for each segment of a single neuron.

$$X_{si} = \text{CEBRA}(X_{st}, X_{be}, X_{se}) \tag{1}$$

We then denote the pair (X_{si}, X_{seg}) for the same time as X in the following sections. The encoder $G(\cdot)$ maps the extracted input X into a latent representation H', which is then aggregated into time-invariant feature embeddings H using adaptive average pooling. These embeddings H are further mapped into a lower-dimensional space Z by a projection layer $P(\cdot)$. The full model $F(\cdot)$ combines feature extraction, encoding, and projection. During training, $F(X)$ produces the projections Z. After training, the projection layer is removed, retaining only the embeddings H. The similarity between embeddings is measured using cosine similarity.

Contrastive Learning - VICReg: VICReg aims to enhance the quality of learned embeddings by incorporating three types of losses: variance loss, invariance loss, and covariance loss.

The invariance loss, ensures that embeddings of segments from the same neuron are close:

$$\mathcal{L}_{\text{invar}}(Z^{(1)}, Z^{(2)}) = \frac{1}{N} \sum_i \left\| Z_i^{(1)} - Z_i^{(2)} \right\|^2, \tag{2}$$

The variance loss regularizes the standard deviation of the embeddings to be near a target value μ, which helps prevent embedding dimensions from becoming non-informative. Given $\mathbf{d_j}(\mathbf{z}) \in \mathbb{R}^B$, a vector of batch values at dimension j, the variance loss is defined as:

$$\mathcal{L}_{\text{var}}(\mathbf{z}) = \frac{1}{D} \sum_{j=1}^{D} \max\left(0, \mu - S\left(\mathbf{d_j}(\mathbf{z}), \epsilon\right)\right), \tag{3}$$

where D represents the number of dimensions in $\mathbf{z}$, and S denotes the regularized standard deviation, $S(x, \epsilon) = \sqrt{\text{Var}(x) + \epsilon}$.

The covariance loss promotes orthogonality among embedding dimensions by decorrelating them:

$$\mathcal{L}_{\text{cov}}(\mathbf{z}) = \frac{1}{D_z} \sum_{i \neq j} (C(\mathbf{z}))_{i,j}^2, \tag{4}$$

where $C(\mathbf{z}) = \frac{1}{N-1} \sum_{i=1}^{N} (\mathbf{z}_i - \bar{\mathbf{z}}) (\mathbf{z}_i - \bar{\mathbf{z}})^T$ is the covariance matrix of $\mathbf{z}$, and $\bar{\mathbf{z}} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{z}_i$.

4 EXPERIMENTS

4.1 DATA

Simulated Data: Since the neuron intrinsic property is hardly available in vivo neuronal recordings, we applied to synthetic data where we can access the ground-truth intrinsic property. To make the synthetic data exhibit dynamics similar to that of real neurons, we simulated the data following the Izhikevich model. The Izhikevich model is a spiking neuron model that combines biological realism with computational efficiency. It is designed to capture the rich dynamics of real neurons while remaining computationally simple. The model is defined by the following differential equations:

$$\frac{dV}{dt} = 0.04V^2 + 5V + 140 - u + I, \tag{5}$$

$$\frac{du}{dt} = a(bV - u), \tag{6}$$

where V is the membrane potential of the neuron, u is a recovery variable, I is the input current, and a, b, and c are hyperparameters. The model also includes a spike-reset mechanism:

$$\text{if } V \geq 30 \text{ mV, then } \begin{cases} V \leftarrow c \\ u \leftarrow u + d \end{cases} \tag{7}$$

a , b , c and d can be regarded as intrinsic properties of each neuron.

Real Data - Bugeon: We utilized a rare, real-world multimodal dataset Bugeon et al. (2022) to test , which comprises two main components: 1) in vivo long-term recordings of multiple neurons in the mouse primary visual cortex using two-photon imaging technology, and 2) spatial transcriptomics of ex vivo slices from the recorded brain tissue to measure the expression levels of 72 selected genes in these neurons. Based on gene expression levels, neurons were assigned two main types of labels: (i) excitatory and inhibitory classification, and (ii) subclass labels (Lamp5, Pvalb, Vip, Sncg, and Sst) for some inhibitory neurons. This dataset contains recordings from four mice (Mouse A, B, C, and D). Mouse A has data from 6 sessions, Mouse B from 3 sessions, Mouse C from 5 sessions, and Mouse D from 3 sessions. During each session, the mice were exposed to three types of visual stimuli: Blank, Drifting Gratings, and Natural Scenes. A total of 9,278 unique neurons were recorded in the dataset.

Real Data - Steinmetz: The Steinmetz dataset comprises 39 Neuropixels recordings, capturing data from 400 to 700 neurons across various regions of the mouse brain of 10 mice during a visual behavior task Steinmetz et al. (2019). This dataset is excellent for exploratory analyses and is well-supported by extensive tutorial resources [1], alongside a wealth of experimental and behavioral variables included within.

---

[1] https://www.youtube.com/watch?v=WXn4-FpVaOo

4.2 EVALUATION

Evaluation on Simulated Data: The process is divided into two steps: (i) Using , perform self-supervised contrastive learning on the dynamic data of all neurons to obtain a representation for each neuron. (ii) Employ a 5-fold cross-validation approach to use these neuron representations as input to a classifier for predicting the pre-defined neuron type labels from the simulation.

Evaluation on Real Data - Bugeon: The process is divided into three steps: (i) perform self-supervised contrastive learning on the dynamic data of all neurons from four mouse in the real dataset to obtain a representation for each neuron; (ii) Based on neurobiological prior knowledge and spatial transcriptomic gene expression information, assign cell type labels to each neuron. The labels fall into two categories: (a) excitatory and inhibitory, and (b) Lamp5, Pvalb, Vip, Sncg, and Sst; (iii) Implement a 4-fold (with the folds based on the identity of the mice) cross-validation approach where the neuron representations are used as input to a classifier to predict the neuron's type labels for both categories (a) and (b).

Out-of-Domain Evaluation - Steinmetz dataset: The process is divided into three steps: (i) Perform self-supervised contrastive learning on the dynamic data of all neurons from all mice in the real dataset to obtain a representation for each neuron. (ii): As before, assign brain area location labels to the neurons. (iii): Implement a 4-fold (with the folds based on the identity of the mice) cross-validation approach where the neuron representations are used as input to a classifier to predict the neuron's location.

4.3 COMPARISON OF METHODS

LOLCAT: This methodSchneider et al. (2023a) follows a supervised learning paradigm. It directly uses activity data from a subset of neurons to train a classifier to predict neuron labels, and then validates on the remaining neurons. Consequently, the representations learned in the intermediate layers of the model contain only label information and do not fully capture the intrinsic properties of the neurons.

PCA: This method employs a self-supervised approach. Principal Component Analysis (PCA) reduces the dimensionality of each neuron's activity data by projecting it onto a lower-dimensional space, thereby providing a representation based on the most significant components of the activity data.

UMAP: Similar to PCA, Uniform Manifold Approximation and Projection (UMAP) is a self-supervised method that reduces the dimensionality of each neuron's activity data. UMAP preserves local and global structures in the data to create a meaningful lower-dimensional representation.

NeurPrint: This self-supervised method involves implicit representation learning through back-propagation. Due to its implicit nature, the model can be difficult to converge and may require substantial training data and time to achieve effective results, and the representations may not align well with the Platonic Representation Hypothesis.

5 RESULTS

5.1 RESULTS ON SIMULATED DATA

In order to prove that representations learned only from neuronal activity contain invariant properties information, we designed to verify with the help of simulation data, we set the time-invariant hyperparameters of neurons as their intrinsic properties based on the Izhikevich model , and then gave stimulation to obtain activity data (Figure 2). Our model uses only the activity data to perform time-invariant representations of neurons, and validates them by demonstrating that representations can be used to distinguish predefined intrinsic hyperparameters. We evaluate the effectiveness of our proposed method, , for intrinsic hyperparameters (five categories, representing five firing modes of neurons) classification using the obtained intrinsic representation, comparing it with four different algorithms: PCA, UMAP, NeuPRINT, and LOLCAT. The classification performance is assessed across multiple neuron types using three common metrics: Precision, Recall, and F1-score. Higher values of these metrics indicate better classification performance.

Figure 2: The top two subplots show the hyperparameters of five firing modes of neurons when simulating data; The lower left subfigure shows the neuron firing: the index of RS is 0-400, the index of IB is 400-800, the index of CH is 800-1200, the index of FS is 1200-1600, the index of LTS is 1600-2000; The lower right subgraph shows the results of the visualization of neuronal activity data using UMAP

From the results presented in the Table 5.1, we observe that: i) consistently achieves the highest performance across most neuron firing modes, especially in the Regular Spiking (RS) and fast spiking (FS) categories, where it significantly outperforms all baselines with F1-scores of 0.884 and 0.881, respectively; ii) Although some baselines, such as LOLCAT and NeuPRINT, demonstrate competitive results in certain neuron firing modes like low-threshold spiking (LTS), their performance falls short when compared to . For example, in the case of Chattering (CH) neurons, achieves an F1-score of 0.671, outperforming LOLCAT's 0.652 and NeuPRINT's 0.611. These results highlight the robustness and accuracy of across varying neuron firing modes, making it the most effective method in this comparison. These neuron Firing Modes are determined by preset hyperparameters, which alse proves that the neuron intrinsic representation from captures hyperparameter information.

5.2   REAL DATA - NEURON PLATONIC INTRINSIC REPRESENTATION CONTAINS MOLECULAR INFORMATION

In this section, we utilize a public multimodal dataset, as described Bugeon, to train and evaluate our model. We present the performance metrics of neuron type classification for three methodologies: NeuPRINT, LOLCAT, and the proposed (PCA and UMAP cannot handle the processing of behavioral information so they are excluded). The metrics assessed include precision (Prec.), recall (Rec.), and F1 score, which are critical for evaluating the effectiveness of classification across various neuron types categorized into subclasses and classes.

Subclass Performance: The results for the subclasses LAMP5, Vip, Pvalb, and Sst are summarized in Table 5.2. Subclass Analysis: For the LAMP5 subclass, F1 score, which balances precision and recall, was highest for NeurPIR (0.569 ± 0.014), indicating a more balanced approach to classification. In the Vip subclass, NeurPIR again led in F1 score (0.662 ± 0.035), suggesting its superior ability to classify Vip neurons accurately. For Pvalb neurons, The F1 score was highest for NeurPIR (0.604 ± 0.043), indicating its effectiveness in classifying Pvalb neurons. In the Sst subclass, NeuPRINT showed the highest precision (0.704 ± 0.168), but this was accompanied by a large standard deviation, indicating variability in performance. The F1 score was highest for NeurPIR (0.492 ± 0.080), despite the challenges in classifying Sst neurons. Overall, the results highlight the effectiveness of NeuPIR in accurately classifying distinct neuronal subtypes, offering improved precision and F1 scores in comparison to existing methodologies. These findings underscore that the neuron representation obtained by our method clearly contains neuron type information. The neuron

Firing Modes	Metric	PCA	UMAP	NeuPRINT	LOLCAT	NeuPIR
regular spiking (RS)	Prec.	0.689	0.473	0.836	0.783	**0.872**
	Rec.	**0.918**	0.590	0.908	0.910	0.898
	F1.	0.786	0.525	0.870	0.841	**0.884**
intrinsically bursting (IB)	Prec.	0.534	0.332	0.646	**0.681**	0.678
	Rec.	0.375	0.310	0.645	0.573	**0.693**
	F1.	0.440	0.320	0.644	0.620	**0.684**
chattering (CH)	Prec.	0.506	0.335	0.648	0.616	**0.722**
	Rec.	0.603	0.248	0.580	**0.698**	0.630
	F1.	0.548	0.285	0.611	0.652	**0.671**
fast spiking (FS)	Prec.	0.778	0.602	0.826	0.853	**0.853**
	Rec.	0.620	0.638	0.820	0.733	**0.913**
	F1.	0.689	0.618	0.823	0.787	**0.881**
low-threshold spiking (LTS)	Prec.	0.970	0.944	0.968	**0.993**	0.990
	Rec.	0.935	0.965	**0.993**	0.983	0.990
	F1.	0.952	0.954	0.980	0.988	**0.990**

Table 1: This table presents the performance metrics—precision (Prec.), recall (Rec.), and F1 score—across five methods (PCA, UMAP, NeurPrint, LOLCAT, and) for different neuron types: regular spiking (RS), intrinsically bursting (IB), chattering (CH), fast spiking (FS), and low-threshold spiking (LTS). The results indicate that and NeurPrint consistently achieve higher precision and F1 scores for most neuron types, while UMAP shows relatively lower performance, particularly for IB and CH neurons.

type here is determined by molecular information, so it can be said that the intrinsic representation contains molecular information.

Neuron Type	Metric	NeuPRINT	LOLCAT	NeurPIR
Subclass				
LAMP5	Prec.	0.481 ± 0.035	0.344 ± 0.023	0.487 ± 0.007
	Rec.	0.667 ± 0.037	0.694 ± 0.033	0.684 ± 0.028
	F1.	0.559 ± 0.036	0.460 ± 0.024	0.569 ± 0.014
Vip	Prec.	0.614 ± 0.044	0.592 ± 0.078	0.657 ± 0.048
	Rec.	0.652 ± 0.034	0.406 ± 0.028	0.668 ± 0.027
	F1.	0.632 ± 0.034	0.480 ± 0.042	0.662 ± 0.035
Pvalb	Prec.	0.559 ± 0.046	0.428 ± 0.026	0.602 ± 0.042
	Rec.	0.604 ± 0.017	0.403 ± 0.017	0.607 ± 0.055
	F1.	0.580 ± 0.032	0.415 ± 0.019	0.604 ± 0.043
Sst	Prec.	0.704 ± 0.168	0.477 ± 0.130	0.681 ± 0.072
	Rec.	0.323 ± 0.084	0.170 ± 0.077	0.388 ± 0.080
	F1.	0.441 ± 0.105	0.248 ± 0.100	0.492 ± 0.080
Class				
Ex	Prec.	0.685 ± 0.050	0.555 ± 0.010	0.720 ± 0.025
	Rec.	0.774 ± 0.020	0.854 ± 0.009	0.817 ± 0.031
	F1.	0.726 ± 0.036	0.673 ± 0.008	0.765 ± 0.009
In	Prec.	0.737 ± 0.037	0.682 ± 0.020	0.790 ± 0.020
	Rec.	0.640 ± 0.073	0.314 ± 0.027	0.680 ± 0.048
	F1.	0.685 ± 0.058	0.430 ± 0.028	0.729 ± 0.023

Table 2: Performance metrics for neuron type classification, 4-fold cross-validation was used, with the folds based on the identity of the mice This table delineates the precision (Prec.), recall (Rec.), and F1 score for various neuron types categorized into subclasses (LAMP5, Vip, Pvalb, Sst) and classes (Ex, In). The metrics demonstrate the comparative performance of these methods in identifying and classifying distinct neuronal subtypes.

Figure 3: a: Steinmetz dataset contains 39 subdataset from 10 mice. Each subdataset records 400-700 neurons each from across the mouse brain during a visual behavior task. Each neuron with its brain area label. b: show of part of subdataset3. c: show of part of subdataset26.

Region	Method	Precision(In/Out)	Recall(In/Out)	F1(In/Out)
Vis	NeuPrint	0.923 ± 0.029 / 0.751 ± 0.047	0.614 ± 0.052 / 0.467 ± 0.105	0.736 ± 0.042 / 0.572 ± 0.086
	LOLCAT	0.836 ± 0.049 / 0.775 ± 0.014	0.450 ± 0.059 / 0.425 ± 0.053	0.584 ± 0.060 / 0.547 ± 0.043
	NeuPIR	0.900 ± 0.016 / 0.825 ± 0.017	0.652 ± 0.024 / 0.514 ± 0.012	0.756 ± 0.013 / 0.633 ± 0.013
Thal	NeuPrint	0.758 ± 0.025 / 0.669 ± 0.025	0.762 ± 0.005 / 0.724 ± 0.019	0.760 ± 0.013 / 0.695 ± 0.016
	LOLCAT	0.717 ± 0.047 / 0.698 ± 0.073	0.738 ± 0.013 / 0.681 ± 0.010	0.726 ± 0.025 / 0.688 ± 0.040
	NeuPIR	0.765 ± 0.025 / 0.725 ± 0.049	0.772 ± 0.018 / 0.761 ± 0.027	0.768 ± 0.015 / 0.742 ± 0.034
Hipp	NeuPrint	0.708 ± 0.033 / 0.626 ± 0.026	0.785 ± 0.013 / 0.671 ± 0.022	0.744 ± 0.023 / 0.647 ± 0.016
	LOLCAT	0.668 ± 0.034 / 0.595 ± 0.028	0.742 ± 0.018 / 0.703 ± 0.026	0.703 ± 0.027 / 0.644 ± 0.020
	NeuPIR	0.723 ± 0.018 / 0.644 ± 0.032	0.782 ± 0.019 / 0.743 ± 0.019	0.751 ± 0.013 / 0.689 ± 0.017
Mid	NeuPrint	0.669 ± 0.026 / 0.611 ± 0.012	0.814 ± 0.005 / 0.748 ± 0.031	0.734 ± 0.017 / 0.672 ± 0.013
	LOLCAT	0.584 ± 0.018 / 0.569 ± 0.024	0.768 ± 0.010 / 0.726 ± 0.029	0.663 ± 0.013 / 0.637 ± 0.009
	NeuPIR	0.669 ± 0.028 / 0.634 ± 0.027	0.789 ± 0.018 / 0.740 ± 0.013	0.724 ± 0.021 / 0.683 ± 0.021

Table 3: Performance Comparison of Methods Across Brain Regions (In-Distribution vs Out-of-Distribution)

5.3 REAL DATA - SHOWS ROBUSTNESS ON OUT-OF-DOMAIN (UNSEEN ANIMAL) DATA

In this section, our main focus is to validate the consistency of the intrinsic representations obtained by the model across different animals. The Steinmetz dataset comprises neural activity data from ten rats, with each neuron labeled by the brain region it is in. It can be intuitively seen from Figure

3 that the response patterns of neurons in the same brain area are significantly different in different mice, but we hope that the representation we obtain through the model still contains consistent information and contains consistent brain area information. This corresponds to the evaluation of the generalizability of the representations obtained by the model on cross-modal (here, cross-animal) data in deep learning. We use a task of classifying brain regions to validate the intrinsic representations. We use Ne u to perform self-supervised training on all the neurons from the all mice to obtain the intrinsic representations. For the downstream task of brain region classification, we use the intrinsic representations from the neurons of five mice chosen at random as the training set and the validation set, while the intrinsic representations from the neurons of the remaining mice serve as the cross-animal test set.

As shown in Figure 5.2, the model's generalizability is evaluated through brain region classification across different neural representations. In-domain validation results demonstrate that consistently outperforms other methods like NeuPRINT and LOLCAT across all brain regions, including visual cortex (vis ctx), thalamus (thal), hippocampus (hipp), and midbrain. Specifically, achieves validation accuracies close to 0.8 in most regions, with LOLCAT slightly trailing behind.

When examining out-of-domain performance (cross-animal), we notice a general drop in accuracy for all methods. , however, still retains a higher degree of accuracy across all brain regions compared to NeuPRINT and LOLCAT, showing the model's ability to capture more robust intrinsic representations across different animals. This consistent outperformance across both in-domain and out-of-domain tests suggests that the representations obtained by better generalize across animals while still preserving critical brain region information.

6 CONCLUSION AND DISCUSSION

In this paper, we presented Neuron Platonic Inherent Representation (), a framework designed to extract intrinsic, time-invariant characteristics from neurons based on their dynamic activity patterns. We addressed the challenge of decoupling these inherent properties from the fluctuating input-output behaviors of neurons, providing a more nuanced understanding of their intrinsic information. leverages multi-segment contrastive learning tailored to neural population data, enabling it to generate robust, consistent representations of neurons across varying experimental conditions and animal subjects. Through simulations using Izhikevich-modeled neurons, we validated that successfully encodes key hyperparameters representing intrinsic properties. Furthermore, using real-world datasets of neuronal populations from the mouse brain, demonstrated its ability to produce representations that facilitate accurate cell type prediction, illustrating its utility in biological tasks. Importantly, exhibited strong generalization capabilities, maintaining consistent performance across different animals and exhibiting cross-domain robustness when evaluated on out-of-domain test sets.

In conclusion, provides a novel, scalable approach for extracting and leveraging intrinsic neuronal properties, offering significant potential for re-evaluating existing neuroscience data and advancing our understanding of neural computation. Future work may focus on further refining 's ability to handle even more diverse datasets, as well as applying it to other domains where extracting intrinsic properties from complex systems is crucial.

7 REPRODUCIBILITY STATEMENT

To enhance the reproducibility of this study, we provide an Appendix section comprising 4 subsections that offer detailed supplementary information. Appendix A.1 presents the pseudo-code of Synthetic Data. Appendix A.2 presents python code for downloading and organizing the steinmetz dataset. Appendix A.3 presents the pseudo-code of sownstream task. Appendix A.4 presents the Confusion matrices for different models.

REFERENCES

Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

Stephane Bugeon, Joshua Duffield, Mario Dipoppa, Anne Ritoux, Isabelle Prankerd, Dimitris Nicoloutsopoulos, David Orme, Maxwell Shinn, Han Peng, Hamish Forrest, et al. A transcriptomic axis predicts state modulation of cortical interneurons. *Nature*, 607(7918):330–338, 2022.

Eugene M Izhikevich. Simple model of spiking neurons. *IEEE Transactions on neural networks*, 14(6):1569–1572, 2003.

Eugene M Izhikevich and Richard FitzHugh. Fitzhugh-nagumo model. *Scholarpedia*, 1(9):1349, 2006.

Dmitry Kobak and Philipp Berens. The art of using t-sne for single-cell transcriptomics. *Nature communications*, 10(1):5416, 2019.

Ying-Hui Liu and Xiao-Jing Wang. Spike-frequency adaptation of a generalized leaky integrate-and-fire model neuron. *Journal of computational neuroscience*, 10:25–45, 2001.

Andrzej Maćkiewicz and Waldemar Ratajczak. Principal components analysis (pca). *Computers & Geosciences*, 19(3):303–342, 1993.

Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

Lu Mi, Trung Le, Tianxing He, Eli Shlizerman, and Uygar Sümbül. Learning time-invariant representations for individual neurons from population dynamics. *Advances in Neural Information Processing Systems*, 36:46007–46026, 2023.

Mark Nelson and John Rinzel. The hodgkin-huxley model. *The book of genesis*, 2, 1995.

Rohan Prakash. Interpreting latent manifolds of high-dimensional neural activity using pi-vae. Master's thesis, New York University Tandon School of Engineering, 2024.

Aidan Schneider, Mehdi Azabou, Louis McDougall-Vigier, David F Parks, Sahara Ensley, Kiran Bhaskaran-Nair, Tomasz Nowakowski, Eva L Dyer, and Keith B Hengen. Transcriptomic cell type structures in vivo neuronal activity across multiple timescales. *Cell reports*, 42(4), 2023a.

Steffen Schneider, Jin Hwa Lee, and Mackenzie Weygandt Mathis. Learnable latent embeddings for joint behavioural and neural analysis. *Nature*, 617(7960):360–368, 2023b.

Nicholas A Steinmetz, Peter Zatka-Haas, Matteo Carandini, and Kenneth D Harris. Distributed coding of choice, action and engagement across the mouse brain. *Nature*, 576(7786):266–273, 2019.

Bernardo Torres, Stefan Lattner, and Gael Richard. Singer identity representation learning using self-supervised techniques. *arXiv preprint arXiv:2401.05064*, 2024.

A APPENDIX

A.1 SIMULATION OF NEURONAL POPULATION USING THE IZHIKEVICH MODEL

1. Define the Izhikevich model parameters for each neuron type:
   - Regular Spiking (RS)
   - Intrinsically Bursting (IB)
   - Chattering (CH)
   - Fast Spiking (FS)
   - Low-Threshold Spiking (LTS)

2. Initialize the total number of excitatory (N_e) and inhibitory (N_i) neurons.

3. Assign neuron types to indices:

- 25% Regular Spiking (RS)
- 25% Intrinsically Bursting (IB)
- 25% Chattering (CH)
- 25% Fast Spiking (FS)

4. Initialize synaptic connection matrix S.

5. Set initial values for membrane potential v and recovery variable u.

6. For each time step t from 0 to T:

   - Calculate input current I.
   - If $t > 0$, add synaptic contributions from previously fired neurons.
   - Update membrane potential v and recovery variable u using Euler's method.
   - Check for spikes and record firing events.
   - Reset membrane potential and increment recovery variable for fired neurons.
   - Store current and potentials for analysis.

7. Update neuron activity data.

8. Plot firing events.

9. Print shapes of activity arrays for each neuron type.

A.2 PYTHON CODE FOR DOWNLOADING AND ORGANIZING THE STEINMETZ DATASET

```python
# @title Data retrieval
import os, requests

fname = []
for j in range(3):
    fname.append('steinmetz_part%d.npz' % j)
url = ["https://osf.io/agvxh/download"]
url.append("https://osf.io/uv3mw/download")
url.append("https://osf.io/ehmw2/download")

for j in range(len(url)):
    if not os.path.isfile(fname[j]):
        try:
            r = requests.get(url[j])
        except requests.ConnectionError:
            print("!!! Failed to download data !!!")
        else:
            if r.status_code != requests.codes.ok:
                print("!!! Failed to download data !!!")
            else:
                with open(fname[j], "wb") as fid:
                    fid.write(r.content)
# @title Data loading
alldat = np.array([])
for j in range(len(fname)):
    alldat = np.hstack((alldat,
                        np.load('steinmetz_part%d.npz' % j,
                                allow_pickle=True)['dat']))
```

A.3 PSEUDO-CODE OF DOWNSTREAM TASK

Algorithm 1 Neural Network Classification with K-Fold Cross-Validation

---

1: **Import necessary libraries:** NumPy, scikit-learn, Seaborn, Matplotlib
2: **Ignore warnings**
3:
4: **procedure** MAIN
5:     **Initialize labels** as NumPy array
6:     **Encode labels** using LabelEncoder
7:     **Standardize neuron features** using StandardScaler
8:
9:     **Initialize StratifiedKFold** with 5 splits
10:    **Initialize statistics dictionaries** for precision, recall, F1-score
11:    **Initialize empty list for confusion matrices**
12:
13:    **for** each fold in K-Fold **do**
14:        **Split data into training and test sets**
15:        **Create MLPClassifier model**
16:        **Fit model to training data**
17:        **Predict labels for test data**
18:        **Generate classification report and confusion matrix**
19:        **Append confusion matrix to list**
20:
21:        **for** each cell type in classes **do**
22:            **Record precision, recall, and F1-score**
23:        **end for**
24:    **end for**
25:
26:    **for** each cell type in classes **do**
27:        **Calculate and print average metrics**
28:    **end for**
29:
30:    **Compute cumulative confusion matrix**
31:    **Print cumulative confusion matrix**
32:
33:    **Optional: Plot cumulative confusion matrix using Seaborn**
34: **end procedure**

A.4 Confusion matrices for different models.

Figure 4: Confusion matrices for different neuron classification models. Each heatmap illustrates the performance of the respective model in classifying neuron subclasses and excitatory/inhibitory neurons.

A.5 Explanation of Izhikevich model hyperparameters

These hyperparameters characterize the neuron's unique dynamics and response characteristics:

- a: Defines the time scale of the recovery variable u, indicating the speed at which the neuron recovers from spiking events.
- b: Influences the sensitivity of u to the membrane potential V, affecting the neuron's firing patterns.
- c: Determines the reset value for V after a spike, specifying the neuron's post-spike membrane potential.
- d: Indicates the amount by which u is incremented after a spike, shaping the neuron's recovery process.