
Overcoming the Domain Gap in Contrastive Learning of Neural Action Representations

Semih Günel^{1,2}, Florian Aymanns², Sina Honari¹, Pavan Ramdya^{*2}, Pascal Fua^{*1}

¹CVLab, EPFL, {name.surname}@epfl.ch

²Neuroengineering Lab, EPFL, {name.surname}@epfl.ch

Abstract

A fundamental goal in neuroscience is to understand the relationship between neural activity and behavior. For example, the ability to extract behavioral intentions from neural data, or neural decoding, is critical for developing effective brain machine interfaces. Although simple linear models have been applied to this challenge, they cannot identify important non-linear relationships. Thus, a self-supervised means of identifying non-linear relationships between neural dynamics and behavior, in order to compute neural representations, remains an important open problem. To address this challenge, we generated a new multimodal dataset consisting of the spontaneous behaviors generated by fruit flies, *Drosophila melanogaster*—a popular model organism in neuroscience research. The dataset includes 3D markerless motion capture data from six camera views of the animal generating spontaneous actions, as well as synchronously acquired two-photon microscope images capturing the activity of descending neuron populations that are thought to drive actions. Standard contrastive learning and unsupervised domain adaptation techniques struggle to learn *neural action representations* (embeddings computed from the neural data describing action labels) due to large inter-animal differences in both neural and behavioral modalities. To overcome this deficiency, we developed simple yet effective augmentations that close the inter-animal domain gap, allowing us to extract behaviorally relevant, yet domain agnostic, information from neural data. This multimodal dataset and our new set of augmentations promise to accelerate the application of self-supervised learning methods in neuroscience.

1 Introduction

Recent technological advances have enabled large-scale simultaneous recordings of neural activity and behavior in animals including rodents, macaques, humans and the vinegar fly, *Drosophila melanogaster* [1, 2, 3, 4, 5, 6]. In parallel, recent efforts have been made it possible to perform markerless predictions of 2D and 3D animal poses [7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17]. Video and pose data have been used to segment and cluster temporally related behavioral information [18, 19, 20, 21, 22]. To capture a similarly low dimensional representation of neural activity, most efforts have focused on the application of recurrent state space models [23, 24, 25], or variational autoencoders [26, 4]. By contrast, there has been relatively limited work aimed at extracting behavioral information from neural data [27, 28, 29] and most efforts have focused on identifying linear relationships between these two modalities using simple correlation analysis, or generalized linear models [30, 31, 32]. However, *neural action representations*—the mapping of behavioral information within neural data—which are particularly crucial for brain-machine interfaces and closed-loop experimentation [33, 34] are highly nonlinear. Therefore, devising a systematic approach for uncovering complex non-linear relationships between behavioral and neural modalities remains an important challenge.

*Equal contribution

Contrastive learning is one promising approach to address this gap. It has been used to extract information from multimodal datasets in a self-supervised way, for modalities including audio, speech, and optical flow [35, 36, 37, 38, 39, 40]. Contrastive learning also has been applied to unimodal datasets, including the study of human motion sequences [41, 42, 43], medical imaging [44, 45], video understanding [46, 47], and pose estimation [48, 49]. Thus, contrastive learning holds great promise for application in neuroscience.

One of the largest barriers to applying contrastive learning to behavioral-neural multimodal datasets is the fact that their statistics (e.g., neuron locations and sizes, body part lengths and ranges of motion) often differ dramatically across animals. This makes it difficult to train models that can generalize across subjects. We confront this domain gap when comparing neural imaging datasets from two different flies (**Supplementary Fig. S3; Supplementary Videos 1-2**). Although multimodal domain adaptation methods for downstream tasks such as action recognition exist [35], they assume supervision in the form of labeled source data. However, labeling behavioral-neural datasets requires expensive and arduous manual labor by trained scientists, and thus often leaving the vast majority of data unlabeled. Similarly, it is non-trivial to generalize few-shot domain adaptation methods to multimodal tasks [50, 51]. Thus, the field of neuroscience needs new computational approaches that can extract information from ever-increasing amounts of unlabeled multimodal datasets that also suffer from extensive domain gaps across subjects.

Here, we address this challenge by extracting domain agnostic action representations from neural data. We measure representation quality using an action recognition task, in which we apply a linear classification head and transfer our pretrained weights to classify action labels. Therefore, we call our representations *neural action representations*. To best reflect real world conditions, during the unsupervised pre-training phase, we assume access to the paired behavioral-neural data for all domains but without any action labels. Then, we show that a strong domain gap exists across data taken from different animals, rendering standard contrastive methods ineffective. To address this challenge, we propose a set of simple augmentations that can perform domain adaptation and extract useful representations. We find that the resulting model outperforms baseline approaches, including linear models, previous neural representation learning approaches and common domain adaptation techniques. Finally, to accelerate the uptake and development of these and other self-supervised methods in neuroscience, we will release our new multimodal *Drosophila* behavioral-neural dataset along with associated dense action labels for spontaneously-generated behaviors

2 Methods

2.1 Problem Definition

We assume a paired set of data $\mathcal{D}_s = \{(\mathbf{b}_i^s, \mathbf{n}_i^s)\}_{i=1}^{n_s}$, where \mathbf{b}_i^s and \mathbf{n}_i^a represent the behavioral and neural information respectively, with n_s being the number of samples for animal $s \in \mathcal{S}$. We quantify behavioral information \mathbf{b}_i^s as a set of 3D poses corresponding to a set of frames i from animal s , and neural information \mathbf{n}_i^s as a set of two-photon microscope images capturing the activity of neurons. We assume that the two modalities are always synchronized (paired), and therefore describe the same set of events. Our goal is to learn a parameterized image encoder function f_n , which maps a set of neural images \mathbf{n}_i^s to a low-dimensional representation. We aim for our learned representation to be representative of the underlying behavioral label, while being modality-agnostic and not representative of the underlying animal identity information s , and therefore effectively removing the domain gap across animals and modalities. We assume that we are not given behavioral labels during unsupervised training.

2.2 Contrastive Representation Learning

For each input pair $(\mathbf{b}_i^s, \mathbf{n}_i^s)$, we first draw a random view $(\tilde{\mathbf{b}}_i^s, \tilde{\mathbf{n}}_i^s)$ with a sampled transformation function $t_n \sim \mathcal{T}_n$ and $t_b \sim \mathcal{T}_b$, where \mathcal{T}_n and \mathcal{T}_b represent a family of stochastic image transformation functions for behavioral and neural data, respectively. Next, the encoder functions f_b and f_n transform input data into low-dimensional vectors \mathbf{h}_b and \mathbf{h}_n , followed by non-linear projection functions g_b and g_n , which further transform data into the vectors \mathbf{z}_b and \mathbf{z}_n . During training, we sample a minibatch of N input pairs $(\mathbf{b}_i^s, \mathbf{n}_i^s)$, and train with the symmetric loss function

$$\mathcal{L}_{NCE} = - \sum_{i=1}^N \log \frac{\exp(\langle \mathbf{z}_b^i, \mathbf{z}_n^i \rangle / \tau)}{\sum_{k=1}^N \exp(\langle \mathbf{z}_b^i, \mathbf{z}_n^k \rangle / \tau)} + \log \frac{\exp(\langle \mathbf{z}_n^i, \mathbf{z}_b^i \rangle / \tau)}{\sum_{k=1}^N \exp(\langle \mathbf{z}_n^i, \mathbf{z}_b^k \rangle / \tau)} \quad (1)$$

where $\langle \mathbf{z}_b^i, \mathbf{z}_n^i \rangle$ is the cosine similarity between behavioral and neural modalities and $\tau \in \mathbb{R}^+$ is the temperature parameter. The loss function maximizes the mutual information between two modalities [52]. The symmetric version of the contrastive loss function was previously used in multimodal self-supervised learning [45, 53]. An overview of our method for learning f_n is shown in **Supplementary Fig S2**. Although standard contrastive learning bridges the gap between different modalities, it does not bridge the gap between different animals, a fundamental challenge that we address in this work through augmentations described in the following section.

Swapping Augmentation: Given a set of consecutive 3D poses \mathbf{b}_i^s , for each $k \in \mathbf{i}$, we stochastically replace \mathbf{b}_k^s with one of its nearest pose neighbors in the set of domains $\mathcal{D}_{\mathcal{S}/s}$, where \mathcal{S} is the set of all animals. To do so, we first randomly select a domain $\hat{s} \in \mathcal{S}/s$ and define a probability distribution $\mathbf{P}_{\mathbf{b}_k^s}^{\hat{s}}$ over the domain $\mathcal{D}_{\hat{s}}$ with respect to \mathbf{b}_k^s ,

$$\mathbf{P}_{\mathbf{b}_k^s}^{\hat{s}}(\mathbf{b}_l^{\hat{s}}) = \frac{\exp(-\|\mathbf{b}_l^{\hat{s}} - \mathbf{b}_k^s\|_2)}{\sum_{\mathbf{b}_m^{\hat{s}} \in \mathcal{D}_{\hat{s}}} \exp(-\|\mathbf{b}_m^{\hat{s}} - \mathbf{b}_k^s\|_2)}. \quad (2)$$

We then replace each 3D pose \mathbf{b}_k^s by first uniformly sampling a new domain \hat{s} , and then sampling from the above distribution $\mathbf{P}_{\mathbf{b}_k^s}^{\hat{s}}$, therefore resulting in $\tilde{\mathbf{b}}_k^s \sim \mathbf{P}_{\mathbf{b}_k^s}^{\hat{s}}$. In practice, we calculate the distribution \mathbf{P} only over the first \mathbf{N} nearest neighbors of \mathbf{b}_k^s , in order to sample from a distribution of the most similar poses. We empirically set \mathbf{N} to 128. Swapping augmentation reduces the identity information in the behavioral data without perturbing it to the extent that semantic action information is lost. Each transformed behavioral sample $\tilde{\mathbf{b}}_i^s$ is composed of multiple domains. This forces the behavioral encoding function f_b to leave identity information out, therefore merging multiple domains in the latent space. Swapping augmentation is similar to synonym replacement augmentation used in natural language processing [54], where randomly selected words in a sentence are replaced by their synonyms. To the best of our knowledge, we are the first to use swapping augmentation in the context of time-series analysis or for domain adaptation.

Neural Calcium Imaging Data Augmentation: Our neural data was obtained using two-photon microscopy and calcium imaging. The resulting images are only a function of the underlying neural activity, and have temporal properties that differ from the true neural activity. For example, calcium signals from a neuron change much more slowly than the neuron’s actual firing rate. Consequently, a single neural image \mathbf{n}_t includes decaying information concerning neural activity from the recent past, and thus carries information about previous behaviors. This makes it harder to decode the current behavioral state. We aimed to prevent this overlap of ongoing and previous actions. Specifically, we wanted to teach our network to be invariant with respect to past behavioral information by augmenting the set of possible past actions. To do this, we generated new data $\tilde{\mathbf{n}}_i^s$, that included previous neural activity \mathbf{n}_k^s . To mimic calcium indicator decay dynamics, given a neural data sample \mathbf{n}_i^s of multiple frames, we sample a new neural frame \mathbf{n}_k^s from the same domain, where $k \notin \mathbf{i}$. We then convolve \mathbf{n}_k^s with the temporally decaying calcium convolutional kernel \mathcal{K} , therefore creating a set of images from a single frame \mathbf{n}_k^s , which we then add back to the original data sample \mathbf{n}_i^s . This results in $\tilde{\mathbf{n}}_i^s = \mathbf{n}_i^s + \mathcal{K} * \mathbf{n}_k^s$ where $*$ denotes the convolutional operation. In the Appendix, we explain calcium dynamics and our calculation of the kernel \mathcal{K} in more detail.

3 Experiments

In this section we introduce a new dataset consisting of *Drosophila melanogaster* neural and behavioral recordings as well as the set of downstream evaluation metrics.

3.1 Dataset

Motion Capture and Two-photon Dataset (MC2P): We acquired data from tethered adult female flies, (*Drosophila melanogaster*). This dataset consists of neural activity recorded using a two-photon microscope [55] from the axons of descending neurons passing through the animal’s cervical connective. It also includes behavioral data recorded using multi-view infrared cameras (**Supplementary Fig. S1; Supplementary Videos 1-2**). Specifically, behavioral video data of size 480×960 pixels were acquired at 100 frames-per-second (fps) using a six circular camera network with the animal at its center. The neural data was recorded using a two-photon microscope, yielding images of 480×736 pixels at 16 fps. Eight animals and 133 trials were recorded, resulting in 8.2 hours of recordings with 2,975,000 behavioral and 476,000 neural frames. The dataset includes manual and dense action labels of eight behaviors: *forward walking*, *pushing*, *hindleg grooming*, *abdominal grooming*, *rest*,

foreleg grooming, antennal grooming, and eye grooming. We report the statistics of our dataset in **Supplementary Fig. S5**. See the Appendix for more details.

3.2 Evaluation

To evaluate our unsupervised pretrained neural encoder f_n , we froze its parameters and trained a randomly initialized linear classification layer with SGD. To compare data efficiency, for each setting we evaluated image encoders with 50% and 100% of the data. We report aggregated results over 4-fold cross-validation evaluations and report the average in each task. We evaluated models on the following tasks:

Single-Animal Action Recognition: We performed action recognition on a single domain by training and testing on the same animal. We repeated the same experiment on each of four animals, and report the mean accuracy.

Multi-Animal Action Recognition: We evaluated models on their ability to reduce the domain gap. We trained the linear classifier on N-1 animals and tested on the left-out one, leaving each animal out one at a time.

Identity Recognition: We classified animal identity from among the eight animals. We sampled 1000 random data-points uniformly across animals and applied 4-fold cross validation. In the case that the learned representations are domain (subject) invariant, we expect that the linear classifier will not be able to detect the domain of the representations, resulting in a lower identity recognition accuracy.

4 Results

We present action recognition results from neural imaging data in **Table 1** and identity recognition task results in **Table S2**. For the supervised baseline, we trained an MLP with manually annotated action labels using cross-entropy loss, with the raw neural data as input, and show the results in the "Raw" section of **Table 1**. For the "Self-Supervised" section, before using the proposed augmentations, the contrastive method SimCLR performed worse than convolutional and recurrent regression-based methods including the current state-of-art BehaveNet [27]. Although domain adaptation methods MMD (Maximum Mean Discrepancy) and GRL (Gradient Reversal Layer) close the domain gap and lower identity recognition accuracy, they do not position semantically similar points near one another (**Supplementary Fig. S4**). As a result, domain adaptation-based methods do not result in significant improvements in the action recognition task. Although regression-based methods suffer less from the domain gap problem, they do not produce as discriminative representations as contrastive learning based methods. The same trend is observed in **Table S2**. Our proposed set of augmentations close the domain gap, while significantly improving the action recognition baseline for self-supervised methods, for both single-animal and multi-animal tasks. We include detailed information about the baselines in the Appendix.

5 Conclusion

We introduced an unsupervised *neural action representation* framework. We extended previous methods by establishing set of augmentations that we show overcomes the multimodal domain gap in our *Drosophila* behavioral-neural dataset. Finally, we will share in order to dataset to accelerate the application of self-supervised learning methods in neuroscience. In future work, we aim to extend our work for domain generalization.

	Tasks → Percentage of Data →	Single-Animal ↑		Multi-Animal ↑	
		0.5	1.0	0.5	1.0
Random Guess		16.6	16.6	16.6	16.6
Neural (Linear)	Raw	29.3	32.5	18.4	18.4
Neural (MLP)		–	–	18.4	18.4
SimCLR [56]	Self-Supervised	54.3	57.6	46.9	50.6
Regression (Recurr.)		53.6	59.7	49.4	51.2
Regression (Conv.)		52.6	59.6	50.6	55.8
BehaveNet [27]		54.6	60.2	50.5	56.8
Ours		57.9	63.3	54.8	61.9
SimCLR [56] + MMD	Domain Ada.	53.6	57.8	50.1	53.1
SimCLR [56] + GRL		53.5	56.3	49.9	52.3
Regression (Conv.) + MMD		54.5	60.7	52.6	55.4
Regression (Conv.) + GRL		55.5	60.2	51.8	55.7

Table 1: **Action Recognition Accuracy.** Single- and multi-animal action recognition results on the MC2P dataset. Behavioral and Neural MLP results for the single-animal task are removed because single animals often do not have enough labels for every action.

References

- [1] Daniel A. Dombeck, Anton N. Khabbaz, Forrest Collman, Thomas L. Adelman, and David W. Tank. Imaging large-scale neural activity with cellular resolution in awake, mobile mice. *Neuron*, 56(1):43 – 57, 2007.
- [2] Johannes D Seelig, M Eugenia Chiappe, Gus K Lott, Anirban Dutta, Jason E Osborne, Michael B Reiser, and Vivek Jayaraman. Two-photon calcium imaging from head-fixed *Drosophila* during optomotor walking behavior. *Nature Methods*, 7(7):535–540, 2010.
- [3] Chen C-L, Hermans L, Meera C Viswanathan, Denis Fortun, Florian Aymanns, Michael Unser, Anthony Cammarato, Michael H Dickinson, and Pavan Ramdya. Imaging neural activity in the ventral nerve cord of behaving adult *drosophila*. *Nature communications*, 9(1):4390, 2018.
- [4] Chethan Pandarinath, Daniel J. O’Shea, Jasmine Collins, Rafal Jozefowicz, Sergey D. Stavisky, Jonathan C. Kao, Eric M. Trautmann, Matthew T. Kaufman, Stephen I. Ryu, Leigh R. Hochberg, Jaimie M. Henderson, Krishna V. Shenoy, L. F. Abbott, and David Sussillo. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10):805–815, 2018.
- [5] A. S. Ecker, P. Berens, G. A. Keliris, M. Bethge, N. K. Logothetis, and A. S. Tolias. Decorrelated neuronal firing in cortical microcircuits. *Science*, 327(5965):584–587, 2010.
- [6] Uros Topalovic, Zahra M. Aghajan, Diane Villaroman, Sonja Hiller, Leonardo Christov-Moore, Tyler J. Wishard, Matthias Stangl, Nicholas R. Hasulak, Cory S. Inman, Tony A. Fields, Vikram R. Rao, Dawn Eliashiv, Itzhak Fried, and Nanthia Suthana. Wireless programmable recording and stimulation of deep brain activity in freely moving humans. *Neuron*, 108(2):322–334.e9, 2020.
- [7] Talmo D Pereira, Diego E Aldarondo, Lindsay Willmore, Mikhail Kislin, Samuel S H Wang, Mala Murthy, and Joshua W Shaevitz. Fast animal pose estimation using deep neural networks. *Nature Methods*, 16(1):117–125, 2019.
- [8] Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience*, 21(9):1281–1289, 2018.
- [9] Semih Günel, Helge Rhodin, Daniel Morales, João Campagnolo, Pavan Ramdya, and Pascal Fua. DeepFly3D, a deep learning-based approach for 3D limb and appendage tracking in tethered, adult *Drosophila*. *eLife*, 8:3686, 2019.
- [10] Praneet C. Bala, Benjamin R. Eisenreich, Seng Bum Michael Yoo, Benjamin Y. Hayden, Hyun Soo Park, and Jan Zimmermann. Automated markerless pose estimation in freely moving macaques with openmonkeystudio. *Nature Communications*, 11(1):4560, 2020.
- [11] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.
- [12] Jacob M Graving, Daniel Chae, Hemal Naik, Liang Li, Benjamin Koger, Blair R Costelloe, and Iain D Couzin. Deepposekit, a software toolkit for fast and robust animal pose estimation using deep learning. *eLife*, 8:e47994, 2019.
- [13] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: Regional multi-person pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [14] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [15] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

- [16] Adam Gosztolai, Semih Günel, Victor Lobato-Ríos, Marco Pietro Abrate, Daniel Morales, Helge Rhodin, Pascal Fua, and Pavan Ramdya. Liftpose3d, a deep learning-based approach for transforming two-dimensional to three-dimensional poses in laboratory animals. *Nature Methods*, 18(8):975–981, 2021.
- [17] Siyuan Li, Semih Günel, Mirela Ostrek, Pavan Ramdya, Pascal Fua, and Helge Rhodin. Deformation-aware unpaired image translation for pose estimation on laboratory animals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [18] Jennifer J Sun, Ann Kennedy, Eric Zhan, David J Anderson, Yisong Yue, and Pietro Perona. Task programming: Learning data efficient behavior representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [19] Cristina Segalin, Jalani Williams, Tomomi Karigo, May Hui, Moriel Zelikowsky, Jennifer J. Sun, Pietro Perona, David J. Anderson, and Ann Kennedy. The mouse action recognition system (mars): a software pipeline for automated analysis of social behaviors in mice. *bioRxiv*, 2020.
- [20] Katherine Overman, Daniel Choi, Kawai Leung, Joshua Shaevitz, and Gordon Berman. Measuring the repertoire of age-related behavioral changes in drosophila melanogaster. *bioRxiv*, 2021.
- [21] Talmo D. Pereira, Joshua W. Shaevitz, and Mala Murthy. Quantifying behavior to understand the brain. *Nature Neuroscience*, 23(12):1537–1549, 2020.
- [22] Robert Evan Johnson, Scott Linderman, Thomas Panier, Caroline Lei Wee, Erin Song, Kristian Joseph Herrera, Andrew Miller, and Florian Engert. Probabilistic models of larval zebrafish behavior reveal structure on many scales. *Current Biology*, 30(1):70–82.e4, 2020.
- [23] Josue Nassar, Scott W. Linderman, M. Bugallo, and Il-Su Park. Tree-structured recurrent switching linear dynamical systems for multi-scale modeling. *arXiv*, 2019.
- [24] Scott Linderman, Annika Nichols, David Blei, Manuel Zimmer, and Liam Paninski. Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in c. elegans. *bioRxiv*, 2019.
- [25] Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski. Bayesian Learning and Inference in Recurrent Switching Linear Dynamical Systems. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.
- [26] Yuanjun Gao, Evan Archer, L. Paninski, and J. Cunningham. Linear dynamical neural population models through nonlinear embeddings. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [27] Eleanor Batty, Matthew Whiteway, Shreya Saxena, Dan Biderman, Taiga Abe, Simon Musall, Winthrop Gillis, Jeffrey Markowitz, Anne Churchland, John P Cunningham, et al. Behavenet: nonlinear embedding and bayesian neural decoding of behavioral videos. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [28] Omid G. Sani, Hamidreza Abbaspourazad, Yan T. Wong, Bijan Pesaran, and Maryam M. Shanechi. Modeling behaviorally relevant neural dynamics enabled by preferential subspace identification. *Nature Neuroscience*, 24(1):140–149, 2021.
- [29] Joshua I. Glaser, Ari S. Benjamin, Raeed H. Chowdhury, Matthew G. Perich, Lee E. Miller, and Konrad P. Kording. Machine learning for neural decoding. *eNeuro*, 7(4), 2020.
- [30] Alice A. Robie, Jonathan Hirokawa, Austin W. Edwards, Lowell A. Umayam, Allen Lee, Mary L. Phillips, Gwyneth M. Card, Wyatt Korff, Gerald M. Rubin, Julie H. Simpson, Michael B. Reiser, and Kristin Branson. Mapping the neural substrates of behavior. *Cell*, 170(2):393–406.e28, 2017.

- [31] Simon Musall, Matthew T. Kaufman, Ashley L. Juavinett, Steven Gluf, and Anne K. Churchland. Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience*, 22(10):1677–1686, 2019.
- [32] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):255–255, 2019.
- [33] Shixian Wen, Allen Yin, Po-He Tseng, Laurent Itti, Mikhail A. Lebedev, and Miguel Nicolelis. Capturing spike train temporal pattern with wavelet average coefficient for brain machine interface. *Scientific Reports*, 11(1):19020, 2021.
- [34] Celia K S Lau, Meghan Jelen, and Michael D Gordon. A closed-loop optogenetic screen for neurons controlling feeding in drosophila. *G3 (Bethesda)*, 11(5), 05 2021.
- [35] Jonathan Munro and Dima Damen. Multi-modal Domain Adaptation for Fine-grained Action Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [36] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [37] Humam Alwassel, Dhruv Mahajan, Bruno Korbar, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [38] Yuki M. Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations (ICLR)*, 2020.
- [39] Relja Arandjelović and Andrew Zisserman. Objects that sound. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [40] Yuki M. Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [41] Yuejiang Liu, Qi Yan, and Alexandre Alahi. Social nce: Contrastive learning of socially-aware motion representations. *arXiv*, 2020.
- [42] Kun Su, Xiulong Liu, and Eli Shlizerman. Predict & cluster: Unsupervised skeleton based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [43] Lilang Lin, Sijie Song, Wenhan Yang, and Jiaying Liu. MS2L: Multi-task self-supervised learning for skeleton based action recognition. In *Proceedings of the ACM International Conference on Multimedia*, 2020.
- [44] Krishna Chaitanya, Ertunc Erdil, Neerav Karani, and Ender Konukoglu. Contrastive learning of global and local features for medical image segmentation with limited annotations. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [45] Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D. Manning, and Curtis P. Langlotz. Contrastive learning of medical visual representations from paired images and text. *arXiv*, 2020.
- [46] Tian Pan, Yibing Song, Tianyu Yang, Wenhao Jiang, and Wei Liu. Videomoco: Contrastive video representation learning with temporally adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [47] I. Dave, Rohit Gupta, M. N. Rizve, and M. Shah. TCLR: Temporal contrastive learning for video representation. *arXiv*, 2021.
- [48] Sina Honari, Victor Constantin, Helge Rhodin, Mathieu Salzmann, and Pascal Fua. Unsupervised learning on monocular videos for 3d human pose estimation. *arXiv*, 2020.

- [49] Rahul Mitra, Nitesh B Gundavarapu, Abhishek Sharma, and Arjun Jain. Multiview-consistent semi-supervised learning for 3d human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2020.
- [50] Guoliang Kang, Lu Jiang, Yunchao Wei, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for single-and multi-source domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020.
- [51] Rui Wang, Zuxuan Wu, Zejia Weng, Jingjing Chen, Guo-Jun Qi, and Yu-Gang Jiang. Cross-domain contrastive learning for unsupervised domain adaptation. *arXiv*, 2021.
- [52] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv*, 2019.
- [53] Xin Yuan, Zhe Lin, Jason Kuen, Jianming Zhang, Yilin Wang, Michael Maire, Ajinkya Kale, and Baldo Faieta. Multimodal contrastive training for visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [54] Jason Wei and Kai Zou. EDA: Easy data augmentation techniques for boosting performance on text classification tasks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.
- [55] Chin-Lin Chen, Laura Hermans, Meera C. Viswanathan, Denis Fortun, Florian Aymanns, Michael Unser, Anthony Cammarato, Michael H. Dickinson, and Pavan Ramdya. Imaging neural activity in the ventral nerve cord of behaving adult drosophila. *Nature Communications*, 9(1):4390, 2018.
- [56] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.
- [57] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations, (ICLR)*, 2015.
- [58] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- [59] Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- [60] Jessica Cande, Shigehiro Namiki, Jirui Qiu, Wyatt Korff, Gwyneth M Card, Joshua W Shaevitz, David L Stern, and Gordon J Berman. Optogenetic dissection of descending behavioral control in *Drosophila*. *eLife*, 7:970, 2018.
- [61] Florian Aymanns. utils2p. <https://doi.org/10.5281/zenodo.5501119>, Sep 2021.
- [62] Florian Aymanns. ofco: optical flow motion correction. <https://doi.org/10.5281/zenodo.5518800>, Sep 2021.
- [63] Jérôme Lecoq, Michael Oliver, Joshua H. Siegle, Natalia Orlova, and Christof Koch. Removing independent noise in systems neuroscience data using deepinterpolation. *bioRxiv*, 2020.
- [64] Victor Lobato-Rios, Pembe Gizem Özdil, Shraavan Tata Ramalingasetty, Jonathan Arreguit, Auke Jan Ijspeert, and Pavan Ramdya. Neuromechfly, a neuromechanical model of adult *drosophila melanogaster*. *bioRxiv*, 2021.
- [65] Eftychios A. Pnevmatikakis, Josh Merel, Ari Pakman, and Liam Paninski. Bayesian spike inference from calcium imaging data. *arXiv*, 2013.

- [66] Peter Rupprecht, Stefano Carta, Adrian Hoffmann, Mayumi Echizen, Antonin Blot, Alex C. Kwan, Yang Dan, Sonja B. Hofer, Kazuo Kitamura, Fritjof Helmchen, and Rainer W. Friedrich. A database and deep learning toolbox for noise-optimized, generalized spike inference from calcium imaging. *Nature Neuroscience*, 24(9):1324–1337, 2021.
- [67] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander J. Smola. A kernel method for the two-sample-problem. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2006.
- [68] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) Please see the Conclusion section.
 - (c) Did you discuss any potential negative societal impacts of your work? [\[Yes\]](#) Please see the Broader Impact Statement Section.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[N/A\]](#)
 - (b) Did you include complete proofs of all theoretical results? [\[N/A\]](#)
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#) We include instructions to download and use our dataset in the supplementary materials.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) Please see the appendix, particularly the implementation details section.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#) We use cross-validation and report the mean accuracy. Please see the appendix, the implementation details section.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#) Please see the appendix, the implementation details section.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[N/A\]](#)
 - (b) Did you mention the license of the assets? [\[Yes\]](#) We include the license of our dataset in the supplementary material.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[Yes\]](#) We include instructions to download and using our dataset in the supplementary materials.
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#)
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#)
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[N/A\]](#)