# Detecting change points in neural population activity with contrastive metric learning

**Carolina Urzay (curzay@gatech.edu)**
Georgia Institute of Technology
Atlanta, GA 30308 USA

**Nauman Ahad (nauman.ahad@gatech.edu)**
Georgia Institute of Technology
Atlanta, GA 30308 USA

**Mehdi Azabou (mazabou@gatech.edu)**
Georgia Institute of Technology
Atlanta, GA 30308 USA

**Aidan Schneider (aidans@wustl.edu)**
Washington University in St. Louis
St. Louis, MO 63130 USA

**Geethika Atmakuri (gatmakuri3@gatech.edu)**
Georgia Institute of Technology
Atlanta, GA 30308 USA

**Keith B. Hengen (khengen@wustl.edu)**
Washington University in St. Louis
St. Louis, MO 63130 USA

**Eva L. Dyer (evadyer@gatech.edu)**
Georgia Institute of Technology
Atlanta, GA 30308 USA

This work is licensed under the Creative Commons Attribution 3.0 Unported License.
To view a copy of this license, visit http://creativecommons.org/licenses/by/3.0

## Abstract

**Finding points in time where the distribution of neural responses changes (change points) is an important step in many neural data analysis pipelines. However, in complex and free behaviors, where we see different types of shifts occurring at different rates, it can be difficult to use existing methods for change point (CP) detection because they can't necessarily handle different types of changes that may occur in the underlying neural distribution. In this work, we introduce a new approach for finding changes in neural population states across diverse activities and arousal states occurring in free behavior. Our model follows a contrastive learning approach: we learn a metric for CP detection based on maximizing the Sinkhorn divergences of neuron firing rates across two sides of a labeled CP. We apply this method to a 12-hour neural recording of a freely behaving mouse to detect changes in sleep stages and behavior. We show that when we learn a metric, we can better detect change points and also yield insights into which neurons and sub-groups are important for detecting certain types of switches that occur in the brain.**

**Keywords:** Change Point Detection; Neural Population Activity; Metric Learning; Naturalistic Behavior Analysis; Contrastive Learning

## Introduction

The brain is constantly in a state of flux, with variations occurring as individuals switch their attention to new tasks (Luo & Maunsell, 2019) or change their mood and overall state (Shannon et al., 2013; Li, Poo, & Dan, 2009). Thus, in the analysis of data that spans multiple states or behaviors, identifying change points, or points in time where the distribution of neural responses shifts, becomes a critical task (Brecht, Schneider, Sakmann, & Margrie, 2004; Houweling & Brecht, 2008).

Change points have been studied in neural activity for single neurons (Ratnam, Goense, & Nelson, 2003) and in relatively simple tasks where there are a small number of predefined switches (e.g., Ready-Set-Go has three phases and two change points) (Koepcke, Ashida, & Kretzberg, 2016; Alt, Messer, Roeper, Schneider, & Koeppl, 2018). However, there is a lot less work that examines the detection of changes in multi-neuron recordings during naturalistic behavior, where many different types of CPs can occur without any prompt or warning. In these cases, we need methods that can detect any number of diverse types of shifts that occur in the brain.

In this work, we establish a new metric learning approach for detecting change points in neural population activity. Our method relies on a contrastive learning approach to learn a metric where samples on different sides of a labeled change point are repelled from one another and nearby points on one side of a change point are brought closer.

We apply our method to a 12 hour block of neural activity from hippocampus and show that we can better detect sleeping states and behavior from a small amount of labeled data. Critically, our method is also interpretable and can be used to reveal which neurons or interactions between neurons are important for revealing certain types of shifts in the population states. Our results suggest that by coupling metric learning with a contrastive sampling mechanism, we can build interpretable measures of shifts in brain state over long time periods.

## Methods

**Dataset description.** To study change point detection over long timescales and in naturalistic settings, we curated a dataset containing a 12 hour recording of behavior and neural activity data from the hippocampus of a mouse during free behavior. We spike-sorted the neural data with MountainSort (Chung et al., 2017; Buccino et al., 2020), then binned 42 neurons using 4-second windows resulting in a total of 10,800 samples of the population firing rates.

During the experiment, the animal moves in and out of different sleeping stages and natural behaviors. We performed sleep-scoring to obtain arousal states (wake, sleep REM and sleep nREM) using the recorded local field potentials (Ma, Turrigiano, Wessel, & Hengen, 2019). We also annotated the discrete behaviors of the mouse through a manual tagging and analysis of the in-cage video (running, moving in place, standing still) that is recorded simultaneously.

**Approach.** To detect CPs, we use an *online* approach which selects a window before and after a sample of interest, and computes a measure of divergence between both windows before proceeding to the next sample in time. By streaming through the data and computing the divergence in this manner, we can generate a change point statistic which tells us the likelihood that there is a change at a given point in time.

Methods like the Sinkhorn divergence can give more robust estimates of CPs (Cuturi, 2013; Cheng, Aeron, Hughes, Hussey, & Miller, 2020). In such methods, a CP is detected at time $t$ when the Sinkhorn divergence $S_{L,\gamma}$ between the past window $X_p^t$ and the future window $X_f^t$ is greater than a set threshold $\tau$ (e.g. $(S_{L,\gamma}(X_p^t, X_f^t) \geq \tau)$.

Here, we propose to extend this general approach by using some supervised data to *learn a metric* to better detect change points of interest. This metric can be used and plugged into our Sinkhorn divergence and then used in an online manner.

We use change point labels to divide the sequence such that sub-sequences on the same side of a change point are considered to be a similar pair while sub-sequences on opposite sides of a change point are considered to be a dissimilar pair. These similar, dissimilar pairs are used to obtain triplets $(X_i, X_i^s X_i^d)$, which are used to learn a sparse metric $L$ by minimizing

$$l(L) = \sum_{i \in \text{Triplets}} \left[ c - (\mathcal{S}_{L,\gamma}(X_i, X_i^d) - \mathcal{S}_{L,\gamma}(X_i, X_i^s)) \right]^+ + \|L\|_1, \tag{1}$$

where $S_{L,\gamma}$ is Sinkhorn divergence equipped with this learned metric, and $c$ is triplet loss margin (Ahad, Dyer, Hengen, Xie, & Davenport, 2022). The L-1 regularization term helps learn a sparse metric which makes it easy to interpret what features or neurons are responsible for driving different types of changes.

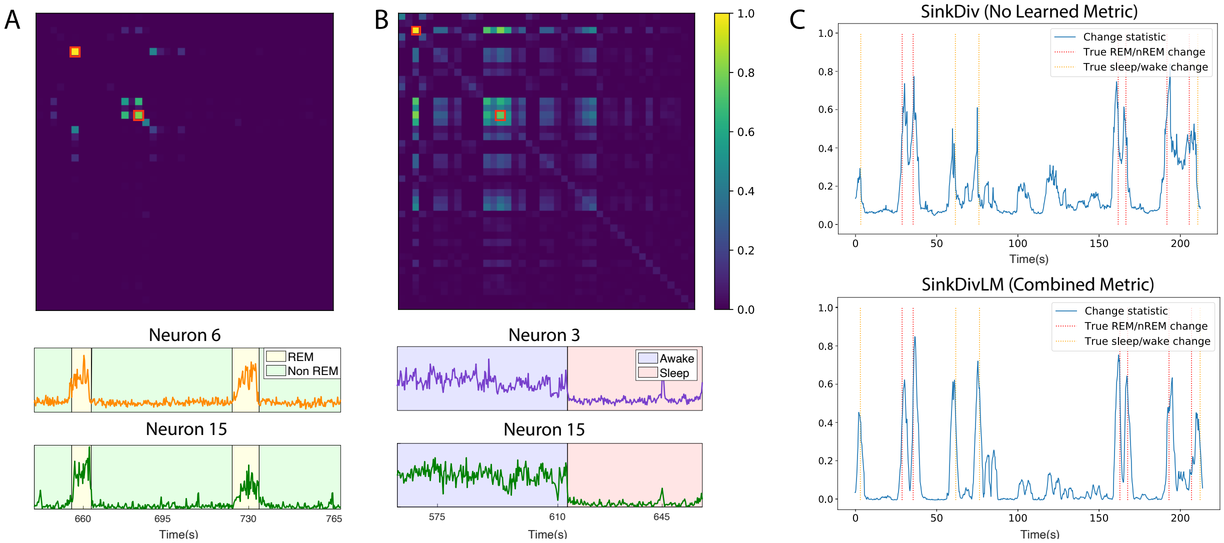

**Figure 1:** *Summary of results.* We show the learned metrics for the sleep substates (REM/nREM on top A, sleep/wake on top B) and the firing rate of the most relevant neurons in each (highlighted in red in learned metric above) at examples of change points (bottom A, bottom B). C shows the change statistics for SinkDiv (top) SinkDivLM when metrics from both sleep/wake and REM/nREM are combined (bottom).

## Results

As our dataset consists of distinct types of sleep state changes, where each of these changes persists for different time scales and is triggered by different neurons, it is difficult to learn a single sparse metric that captures many diverse kinds of CPs while also providing interpretability. To solve this, we learn two metrics, one for sleep/wake transitions and one for REM/nREM transitions (trained with 3 and 14 CPs, respectively), and combine the two sparse metrics to detect changes in all 3 sleep states with one model, by adding and normalizing them.

To show the improvements due to our metric learning approach, we compare our method with SinkDiv, which is Sinkhorn divergence without a learned metric. Since change point detection performance is dependent on detection threshold, we use area under the curve (AUC) as an evaluation metric as it captures detection performance at different thresholds, and is commonly used as a change point evaluation metric in literature (Liu, Yamada, Collier, & Sugiyama, 2013).

The results in Table 1 show the mean AUC for the SinkDiv baseline and our model, SinkDivLM. These demonstrate that learning a metric is an improvement over the baseline, since the score increases from 0.58 to 0.85 in the sleep/wake case and from 0.92 to 0.95 in the case of REM/nREM. As well, the combination of these metrics raise the score for identifying the changes among the 3 arousal substates, REM/nREM/wake.

Interestingly, SinkDivLM trained on sleep/wake labels does not outperform SinkDiv when tested with REM/nREM/wake labels but in the opposite case it does. A possible reason for this is that the neurons responsible for REM/nREM changes (6 and 15, as seen in Figure 1A) are shadowed by others in the sleep/wake metric, making the identification of these CPs a challenging task. However, the learned metric for REM/nREM attributes high relative importance to neurons 13 and 15, which are also relevant for the detection of sleep/wake changes.

**Table 1:** AUC of change point detection for SinkDiv baseline and SinkDivLM model reported on test sets with different truth CP labels.

| | SinkDiv | SinkDivLM |
|---|---|---|
| **Trained on sleep/wake** | | |
| Sleep/wake | 0.58 | 0.85 |
| REM/nREM/wake | 0.79 | 0.72 |
| **Trained on REM/nREM** | | |
| REM/nREM | 0.92 | 0.95 |
| REM/nREM/wake | 0.79 | 0.82 |
| **Combined sleep metrics** | | |
| REM/nREM/wake | 0.79 | 0.85 |
| **Trained on running/no running** | | |
| Running/no running | 0.51 | 0.65 |

To check whether the learnt Sinkhorn divergence metric could identify changes not only in arousal states but also in complex behavior, a high level behavior label (running) was isolated. The model was trained on the 3 points where the mouse started or stopped running (running/no running). In Table 1, we visualize how, as in the case with sleep CPs, SinkDivLM reaches a much higher AUC than SinkDiv.

## Conclusion

In this paper, we show that contrastive metric learning improves the performance of change point detection in two kinds of neural shifts: arousal states and natural behaviors. In addition to improving CP detection, a highlight of our approach is that it provides interpretability at the scale of neurons and local sub-circuits (interaction between neurons through off-diagonal components in learned metric). Such a tool has the potential to infer which behavioral changes are encoded in certain areas of the brain. These findings suggest a promising new direction towards building a tool for neuroscientists to analyze changes in neural activity during complex behavior.

## Acknowledgements

This work was supported, in part, by NSF award IIS-2039741, NIH award 1R01EB029852-01, and a generous gift from the McKnight Foundation.

## References

Ahad, N., Dyer, E. L., Hengen, K. B., Xie, Y., & Davenport, M. A. (2022). Learning sinkhorn divergences for supervised change point detection. *arXiv preprint arXiv:2202.04000*.

Alt, B., Messer, M., Roeper, J., Schneider, G., & Koeppl, H. (2018). Non-parametric bayesian inference for change point detection in neural spike trains. In *2018 ieee statistical signal processing workshop (ssp)* (p. 258-262). doi: 10.1109/SSP.2018.8450787

Brecht, M., Schneider, M., Sakmann, B., & Margrie, T. W. (2004). Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. *Nature*, *427*(6976), 704–710.

Buccino, A. P., Hurwitz, C. L., Garcia, S., Magland, J., Siegle, J. H., Hurwitz, R., & Hennig, M. H. (2020, nov). Spikeinterface, a unified framework for spike sorting. *eLife*, *9*, e61834. Retrieved from `https://doi.org/10.7554/eLife.61834` doi: 10.7554/eLife.61834

Cheng, K. C., Aeron, S., Hughes, M. C., Hussey, E., & Miller, E. L. (2020). Optimal transport based change point detection and time series segment clustering. In *Proc. ieee int. conf. on acoustics, speech, and signal processing (icassp).* Virtual Conference.

Chung, J. E., Magland, J. F., Barnett, A. H., Tolosa, V. M., Tooker, A. C., Lee, K. Y., … Greengard, L. F. (2017). A fully automated approach to spike sorting. *Neuron*, *95*(6), 1381-1394.e6. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0896627317307456` doi: https://doi.org/10.1016/j.neuron.2017.08.030

Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. In *Proc. int. conf. neural inf. process. syst. (neurips).* Lake Thaoe, Nevada.

Houweling, A. R., & Brecht, M. (2008). Behavioural report of single neuron stimulation in somatosensory cortex. *Nature*, *451*(7174), 65–68.

Koepcke, L., Ashida, G., & Kretzberg, J. (2016). Single and multiple change point detection in spike trains: Comparison of different cusum methods. *Frontiers in Systems Neuroscience*, *10*. Retrieved from `https://www.frontiersin.org/article/10.3389/fnsys.2016.00051` doi: 10.3389/fnsys.2016.00051

Li, C.-y. T., Poo, M.-m., & Dan, Y. (2009). Burst spiking of a single cortical neuron modifies global brain state. *Science*, *324*(5927), 643–646.

Liu, S., Yamada, M., Collier, N., & Sugiyama, M. (2013). Change-point detection in time-series data by relative density-ratio estimation. *Neural Netw.*, *43*, 72–83.

Luo, T. Z., & Maunsell, J. H. (2019). Attention can be subdivided into neurobiological components corresponding to distinct behavioral effects. *Proceedings of the National Academy of Sciences*, *116*(52), 26187–26194.

Ma, Z., Turrigiano, G. G., Wessel, R., & Hengen, K. B. (2019). Cortical circuit dynamics are homeostatically tuned to criticality in vivo. *Neuron*, *104*(4), 655-664.e4. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0896627319307378` doi: https://doi.org/10.1016/j.neuron.2019.08.031

Ratnam, R., Goense, J. B., & Nelson, M. E. (2003). Change-point detection in neuronal spike train activity. *Neurocomputing*, *52-54*, 849-855. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0925231202008159` (Computational Neuroscience: Trends in Research 2003) doi: https://doi.org/10.1016/S0925-2312(02)00815-9

Shannon, B. J., Dosenbach, R. A., Su, Y., Vlassenko, A. G., Larson-Prior, L. J., Nolan, T. S., … Raichle, M. E. (2013). Morning-evening variation in human brain metabolism and memory circuits. *Journal of neurophysiology*, *109*(5), 1444–1456.