



This is a repository copy of *Time warping reveals hidden features of neuronal population responses*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/186951/>

Version: Accepted Version

Article:

Okun, M. orcid.org/0000-0002-7795-5513 (2020) Time warping reveals hidden features of neuronal population responses. *Neuron*, 105 (2). pp. 203-204. ISSN 0896-6273

<https://doi.org/10.1016/j.neuron.2019.12.010>

Article available under the terms of the CC-BY-NC-ND licence
(<https://creativecommons.org/licenses/by-nc-nd/4.0/>).

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Time warping reveals hidden features of neuronal population responses

Michael Okun

Centre for Systems Neuroscience and Department of Neuroscience, Psychology and Behaviour,
University of Leicester, Leicester LE1 7RH, UK

m.okun@le.ac.uk

Neuronal population responses can vary across trials as a result of differences in the relative timing of internal brain processes. In this issue of *Neuron*, Williams et al. (2019) present an algorithm for inferring and inverting such trial-to-trial differences, thereby revealing an a priori hidden, precise temporal structure of population responses.

Have you ever noticed how the same sensory signal, e.g., calling a child's name, on different occasions elicits responses with a highly variable delay? To some extent such variability persists even in laboratory conditions, when subjects perform many trials of the same task. This variability is tightly linked to trial-to-trial differences in the ongoing activity of the neural networks responsible for the behavior. For neuroscientists, such variability represents a significant complication in their attempts to understand neuronal responses and their underlying mechanisms.

Consider the spiking of a single neuron across several trials. When its response timing is highly stereotyped, as in Fig. 1A, it is logical to conclude that this neuron's response is locked to the preceding sensory stimulus. In the case depicted in Fig. 1B, however, it is more plausible that the spikes are locked to some internal trigger that follows the stimulus with a lag that varies across trials. In this example the neuron's response is short and vigorous, thus it is relatively straightforward to realign the trials, thereby inferring the trial-to-trial differences in the latency of the hypothesized internal trigger. Unfortunately, in many brain areas neurons with such strong and unambiguous responses are rather uncommon. On the other hand, modern experimental methods allow us to observe the responses of many neurons simultaneously. While neurons' responses can be highly diverse, presuming the lag in each trial is shared across the entire neuronal population, could it be inferred from the joint responses of all the recorded neurons?

In this issue of *Neuron*, Williams et al. (2019) describe a new method for inferring differences in time alignment of population responses across trials. This method is not limited to lags, and is also able to handle cases where trials are differently stretched relative to each other, i.e., when the trial-to-trial differences are not only in the onset of the response but also in the pace of its unfolding. The method's core assumption is that responses are fully described by a fixed firing rate template that on every trial is distorted (warped) in time in a linear or piecewise-linear manner (Fig. 1C). Counterintuitively, because of the highly stochastic nature of spiking data, the restricted set of warping functions is able to fit neuronal population responses substantially better than the generic Dynamic Time Warping method used by computer scientists since the 1970s (Kruskal and Liberman, 1983). Devising an algorithm for fitting this warping model to data required the authors to find efficient strategies for parameter search, regularization and cross-validation. These algorithmic details will not be elaborated upon here. The warping algorithm was tested on synthetic data and then applied to several population recording datasets, as described next.

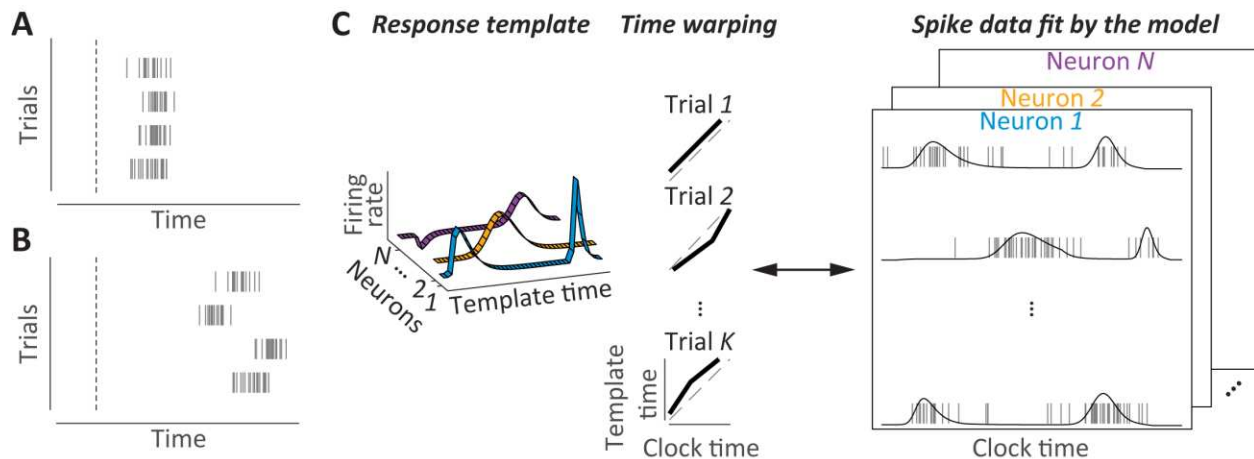


Figure 1. Time warping model of population responses. (A, B) Responses of a single neuron to an external stimulus (dashed line) in four example trials. While the time and number of individual spikes vary across trials, the timing of the response as a whole is rather precise in A, but highly variable in B. (C) The warping model presumes that the responses of a neuronal population are described by a fixed template unique to each neuron and a time warping function unique to each trial (left). The combination of the two provides the firing rate of each neuron on every trial (continuous lines, right). The response template and the time warping functions are optimized to produce predictions with the best fit to the observed spike trains.

First, the algorithm was applied to population recordings of odor responses in mouse olfactory bulb. This is an excellent test case, as it was formerly established that these responses are locked to the (internally controlled) sniff onset, rather than to the odor presentation time, which is controlled by the experimenter and is thus independent of the animal's breathing (Shusterman et al., 2011). The algorithm realigned the odor responses in the different trials, essentially inferring from the neuronal activity the trial-to-trial differences in the lag between the odor and sniff onsets. In the analyzed dataset the sniff times were actually measured, therefore the alignment produced by the algorithm could be compared to the alignment by sniff onset, and the two were found to be in good agreement.

A similar test case was provided by population recordings in motor and premotor cortices of primates that performed an arm reaching task with a variable delay between the appearance of the reach target and the go cue. Although some neurons in these cortical areas respond to target onset, the neuronal response is primarily locked to the movement onset (Churchland et al., 2012; Riehle and Requin, 1989). When the warping algorithm was provided with spike trains aligned by target onset, it was able to produce a corrected alignment with accurately inferred trial-to-trial differences in movement delay. Importantly, linear warping was better than pure time lags (in terms of explained variance on held-out data), as its two degrees of freedom allowed it to capture neuronal responses to both movement and target onsets.

The ability of the warping algorithm to reveal hidden structure of population dynamics was further demonstrated by detecting oscillations that are not locked to stimuli or behavior. Previously, these so called induced oscillations were primarily studied at the level of continuous field potential (FP) signals (Pfurtscheller and Lopes da Silva, 1999), where they are significantly easier to detect. In

the aforementioned reaching task, the warping algorithm revealed that during the delay period, premotor cortex population rate (the summed spiking activity of all the recorded neurons) oscillates in beta frequency. Because the oscillation phase was not locked to external events, its existence was not apparent from standard analyses, such as population rate peri-stimulus histogram. The authors then applied the warping of each trial, inferred from the spiking activity, to the FP from the same area, finding that oscillations in spiking are coherent with beta FP oscillations, a well-documented feature of neural activity in motor brain areas (Engel and Fries, 2010).

The final example of the warping algorithm involved a task where rats had to tap a lever twice, with a delay that had to fall within a narrow range of rewarded durations (780 ± 80 ms, Kawai et al., 2015). The warping algorithm revealed that spiking response of many motor cortex neurons around the time of the first tap was strongly oscillatory, with a ~ 6 -7 Hz period. This population response structure was not apparent from simple alignments of the spiking activity to behavior. Furthermore, unlike the previous case, the oscillations of different neurons were out of phase with each other, thus the population rate did not oscillate and neither did the FP.

Going forward, there are several important ways in which the model can be generalized. One generalization is to allow for several time warping processes on every trial. This would be of particular importance for analyzing simultaneous recordings from multiple brain areas. An additional generalization would allow for some degree of freedom in the population response template, i.e., rather than being fixed, the response template for each trial would belong to a discrete or continuous set of permissible templates. Such generalizations might require models using a different mathematical formalism, e.g., shared latent Gaussian processes (Duncker and Sahani, 2018). That being said, the simple and elegant warping model of the authors is already effective in revealing the timing structure of neuronal population responses in a wide range of cases, as compellingly demonstrated by the above examples. The publicly available code of the model will be a highly valuable resource for other researchers interested in subjecting their data to a similar analysis.

References

- Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. *Nature* 487, 51–56.
- Duncker, L., and Sahani, M. (2018). Temporal alignment and latent Gaussian process factor inference in population spike trains. In *Advances in Neural Information Processing Systems*, pp. 10445–10455.
- Engel, A.K., and Fries, P. (2010). Beta-band oscillations—signalling the status quo? *Current Opinion in Neurobiology* 20, 156–165.
- Kawai, R., Markman, T., Poddar, R., Ko, R., Fantana, A.L., Dhawale, A.K., Kampff, A.R., and Ölveczky, B.P. (2015). Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* 86, 800–812.

Kruskal, J., and Liberman, M. (1983). The symmetric time-warping problem: From continuous to discrete. *Time Warps, String Edits, and Macromolecules: The Theory and Practice of Sequence Comparison*.

Pfurtscheller, G., and Lopes da Silva, F.H. (1999). Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clinical Neurophysiology* *110*, 1842–1857.

Riehle, A., and Requin, J. (1989). Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. *J Neurophysiol* *61*, 534–549.

Shusterman, R., Smear, M.C., Koulakov, A.A., and Rinberg, D. (2011). Precise olfactory responses tile the sniff cycle. *Nat Neurosci* *14*, 1039–1044.

Williams, A.H., Poole, B., Maheswaranathan, N., Dhawale, A.K., Fisher, T., Wilson, C.D., Brann, D.H., Trautmann, E.M., Ryu, S., Shusterman, R., et al. (2019). Discovering Precise Temporal Patterns in Large-Scale Neural Recordings through Robust and Interpretable Time Warping. *Neuron*.