Time-warped PCA: simultaneous alignment and dimensionality reduction of neural data

B. Poole, A. Williams, N. Maheswaranathan, B. Yu, G. Santhanam, S. I. Ryu, S. A. Baccus, K. Shenoy, S. Ganguli

Summary. Analysis of multi-trial neural data often relies on rigid alignment of neural activity to stimulus triggers or behavioral events. However, activity on a single trial may be shifted and skewed in time due to differences in attentional state, biophysical kinetics, and other unobserved latent variables. This temporal variability can inflate the apparent dimensionality of data and obscure our ability to recover inherently simple, low-dimensional structure. For example, small temporal shifts on each trial introduce illusory dimensions as revealed by principal component analysis (PCA). We demonstrate the prevalence of these issues in spike-triggered analysis of retinal ganglion cells and in primate motor cortical neurons during a reaching task. To address these challenges, we develop a novel method, *time-warped PCA* (twPCA), that simultaneously identifies time warps of individual trials and low-dimensional structure across neurons and time. Our method contains a single hyperparameter that trades off complexity of the temporal warps against the dimensionality of the aligned data. Furthermore, we identify the temporal warping in a data-driven, unsupervised manner, removing the need for explicit alignment with external variables. We apply twPCA to motor cortical data recorded from a monkey performing a center-out delayed reaching task. The learned warpings can explain 70% of the variability in reaction time. Time-warped PCA is broadly applicable to a variety of neural systems as a method for disentangling temporal variability across trials as well as discovering underlying neural dynamics and structure of interest.

Significance. Temporal variability in the form of shifts, stretches, and skewing of neural data across trials appears in a multitude of neural systems. For example, in the birdsong system (Long & Fee 2008), varying the temperature of neural tissue leads to stretching and squashing of neural activity and behavioral output. In the olfactory system, not accounting for a latent variable (the sniff cycle) results in temporal shifts across recorded trials (Shusterman et. al. 2011) that obscures the precision of olfactory responses. PCA applied to these datasets will conflate variability across trials due to temporal warping with underlying task-related firing variability. We have developed a general method that disentangles these two factors, allowing recovery of both the underlying neural components along with the temporal warping on each trial.

Additional Detail. Neural signals are often influenced by unobserved latent variables that cause recorded activity to be shifted in time. If the shifts in time are small, they can be well explained by the average signal and its temporal derivative (Fig 1). We provide evidence for shifts in two very different datasets: (a) in motor cortical neurons during center-out reaching in rhesus macaques (Fig 2A), and in spike-triggered covariance analysis of salamander retinal ganglion cells (Fig 2B). In both cases, one of the top two principal components resembles the first temporal derivative of the trial-averaged signal. The presence of derivatives in the principal components of neural data can also be found in existing literature (Fairhall et al 2006, Kobak et al 2016). However, the source of this phenomenon is rarely discussed, and, more importantly, its consequences for neural data analysis have not been systematically studied.

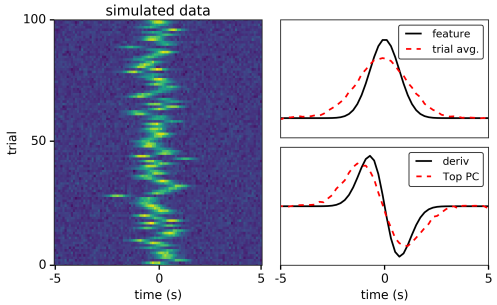

Figure 1: Temporal shifts in synthetic neural activity (left) distort the trial-average estimate of the signal and cause the top PC to ressemble the temporal derivative.

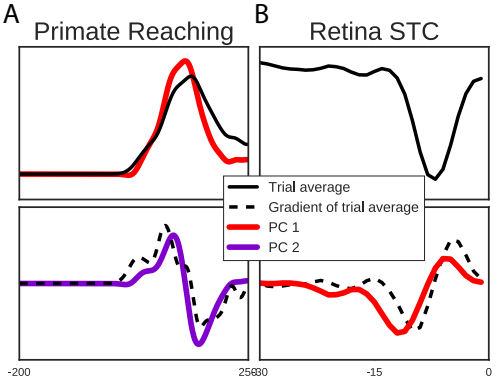

Figure 2: Neurons have PCs that resemble the temporal derivative of the average response across trials.

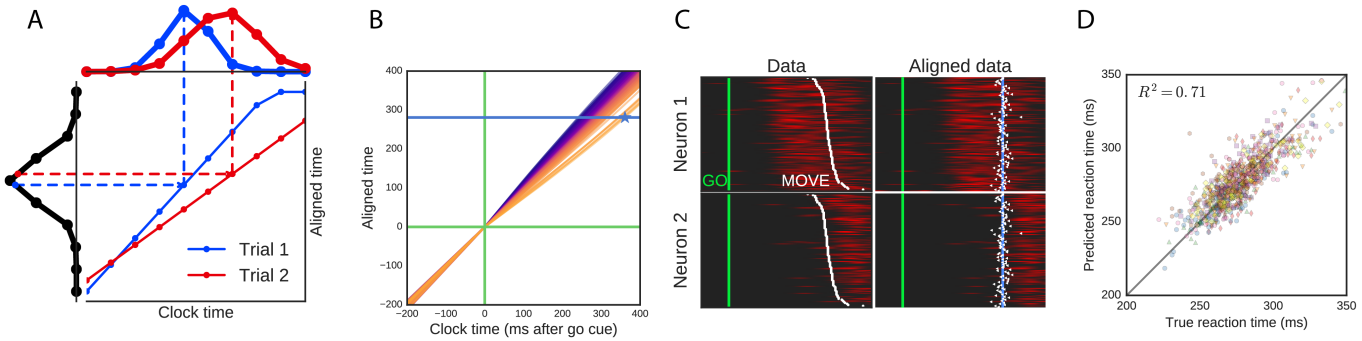

Figure 3: Time-warped PCA predicts reaction time in a center-out reach task. **A.** Schematic depicting the time-warping procedure for two trials. The aligned neural activity (black) is warped into the individual trial activity (blue, red). **B.** Time warps learned for 110 trials on one condition of the reaching data. Time warps were constrained to be 0 at the go cue (green lines). After the go cue, trials with different reaction times (yellow, slow to blue, fast) diverge. Given the true reaction time for one trial (blue star), we can predict the reaction time for all trials by computing the reaction time in the warped space (blue line) and intersecting that with each trial's warping function. **C.** Warping the data into the aligned space removes temporal variability. Each image depicts the activity of neuron with each row containing the trial activity. When sorting trials by reaction time (left), there is a clear shift in neural activity correlated with movement onset (white). Warping the data with the inverse of the learned warping functions from twPCA projects unaligned data into the aligned space. In this space, the movement onset variability is vastly reduced. **D.** Reaction times predicted by twPCA explain 71% of the variance in reaction times across trials from 8 conditions (colors). Only *one* trial of supervision per condition was used to identify the time in the aligned space corresponding to movement onset. Our approach outperforms earlier work from Petreska et al. (NIPS 2011) using latent-variable dynamical systems (52%) and population vector regression (48%).

Method. We apply twPCA to datasets of N neurons over K trials with each trial having length T. For each trial, k, $\mathbf{X}^{(k)} \in \mathbb{R}^{T\times N}$ is a matrix whose columns contain the activity of an individual neuron over time. Our goal is to identify a rank-R representation of the data, $\mathbf{L} = \mathbf{U}\mathbf{V}^T \in \mathbb{R}^{T\times N}$, where $\mathbf{U} \in \mathbb{R}^{T\times R}$ and $\mathbf{V} \in \mathbb{R}^{N\times R}$ are the neuron loadings and temporal dynamics for each of the R components. To compensate for trial-to-trial temporal variability, we model each trial as a temporally warped version of the low-rank representation $\mathbf{L}$, where $\tau^{(k)}(\mathbf{L})$ is a transformation from the data space to the aligned space for each trial k. The twPCA objective is to minimize the reconstruction error of the time-warped low-rank representation with the oserved data: $\underset{\mathbf{U},\mathbf{V},\tau^{(1)},\cdots,\tau^{(K)}}{\text{minimize}} \sum_{k=1}^{K} \left\| \mathbf{X}^{(k)} - \tau^{(k)}\left(\mathbf{U}\mathbf{V}^T\right)\right\|_F^2$. Each trial's transformation $\tau^{(k)}$ is parameterized by a vector in $\mathbb{R}^T$ that represents a *differentiable* transformation from the aligned space to the data space (Fig 3A). The same transformation is applied across all neurons in a trial, but varies between trials, allowing twPCA to explicitly capture trial-to-trial temporal variability.

Unlike dynamic time warping (Berndt & Clifford 1994), twPCA has a differentiable objective and can be optimized to jointly align many sequences, not just a pair. We can also regularize the warping function using known constraints. For example, in the monkey reach task, we force all warping functions to be aligned at the go cue, allowing us to incorporate experimental cues with flexible warping at other points in time (Fig 3B). While PCA has a known globally optimal solution, twPCA is a non-convex problem without guarantees. In practice, we found two heuristics were critical to the performance of twPCA: the transformations $\tau^{(k)}$ were initialized to identity (corresponding to no warping), and the components $\mathbf{U}, \mathbf{V}$ were initialized using the K-truncated SVD of the data matrix. We jointly optimized the objective over the low-rank representation of the aligned data $\mathbf{U}, \mathbf{V}$ and the per-trial transformations $\tau^{(1)}, \cdots, \tau^{(K)}$.

We applied twPCA to neural recordings of a monkey performing center-out reaches. Firing rates were extracted by binning spikes (5ms bins) and smoothing using a Gaussian filter (σ=20ms) from 1s before to 500ms after the go cue. We applied twPCA to each of the 8 conditions separately, and used a single trial from each condition to estimate the time of movement onset in the aligned space (see Figure 3 caption).