firing patterns, e.g. by producing stochastic drift in grid attractor networks. Recently, Hardcastle et al. (Neuron, 2015) proposed that border cells may provide one mechanism for correcting such drift. We construct a model in which experience-dependent Hebbian plasticity during exploration allows border cells to self-organize their responses, while also learning connectivity to grid cells which maintain their activity through an attractor network. We show that border cells in this learned network effectively correct for grid drift despite stochasticity of border cell firing. This error-correction is robust with respect to environmental shape including squares and circles. Furthermore, it survives insertion of barriers within an enclosure (consistent with Solstad et al., 2008) even though a given border cell can fire ambiguously at multiple boundaries with the same allocentric orientation (Solstad et al, 2008; Lever et al, 2009). In our mechanism, the learned border-grid connectivity pattern compensates for such ambiguities. Upon deformation of an environment, e.g., by shrinking or changing shape, the error correction initially fails and grid drift resumes. However, in our model the border-grid connectivity adapts to boundaries of the deformed environment, restoring error correction after a characteristic timescale. Our results demonstrate a class of self-organized mechanisms that achieve robust path integration. These mechanisms predict that: (a) disrupting synaptic plasticity between grid cells and border cells will cause grid patterns to drift in a random walk, and (b) deforming an environment will initially lead to grid drift, which is subsequently stabilized.

III-14. Time-warped PCA: simultaneous alignment and dimensionality reduction of neural data

Ben Poole[1] POOLE@CS.STANFORD.EDU
Alexander Williams[1] AHWILLIA@STANFORD.EDU
Niru Maheswaranathan[1] NIRUM@STANFORD.EDU
Byron Yu BYRONYU@CMU.EDU
Gopal Santhanam[2] GOPALS@STANFORD.EDU
Stephen Ryu SEOULMAN@STANFORD.EDU
Stephen A. Baccus[1] BACCUS@STANFORD.EDU
Krishna Shenoy[1] SHENOY@STANFORD.EDU
Surya Ganguli[1] SGANGULI@STANFORD.EDU

[1]Stanford University
[2]Google X

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 introduces 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 predefined 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.