Sequential Components Analysis

From molecules to whole organisms, the dynamics of natural living systems depart from "thermodynamic equilibrium" [Gnesotto, 2018]. Statistically, the state trajectories produced by such systems are non-reversible, i.e. they do not have equal likelihood of occurring in the reverse temporal direction even in stationary regimes. For example, neural systems exhibit anisotropic waves of activity (e.g. during development), produce precisely ordered spike sequences (e.g. episodic memory), or generate rotational patterns of activity (e.g. motor control). Despite the prevalence and importance of sequential neural activity, there is a relative paucity of methods for exploring the spatio-temporal structure of irreversibility in multivariate time series. Here, we introduce Sequential Components Analysis (SCA), a simple yet effective and scalable method for doing this. SCA performs a systematic analysis of spatio-temporal covariances (all time pairs and unit pairs), and extracts the spatio-temporal modes of activity that contribute most to sequentiality. We highlight the main distinguishing features of the method using a toy example, and apply it to monkey M1 motor activity as well as rat hippocampal data where we show that sequential features separate navigational memories better than mere principal components.

Sequentiality and space-time decomposition We assume data samples, denoted as X, come as $T \times N$ matrices, where T is the number of time bins and N is the number of units (Fig. 1A, left). We say that a spatio-temporal process $X \in \mathbb{R}^{T\times N}$ is reversible (non-sequential), if, and only if $p(X_{ti}, X_{t'i'}) = p(X_{t'i}, X_{ti'})$ for any two distinct time points (t, t') and any two distinct units (i, i'). This definition generalizes the notion of "reversibility" (or "detailed balance", which only applies to stationary processes) to the case of non-stationary processes. By negation, X is said to be "sequential" if it is not reversible.

To study the sequentiality of a spatio-temporal process, we first compute its space-time $TN \times TN$ covariance matrix. For ease of notation, we call the centered data X, too. The space-time covariance matrix $C \in \mathbb{R}^{TN\times TN}$ is defined as:

$$C = \left\langle \mathrm{vec}\,(X)\,\mathrm{vec}\,(X)^T \right\rangle. \tag{1}$$

where $\mathrm{vec}\,(X)$ is the standard vectorization operator (vertical stacking of the columns of X) and $\langle\cdot\rangle$ denotes averaging over different spatio-temporal samples. While in general this matrix is prohibitively large to store or compute, we have developed a variant of the method that entirely bypasses the need to compute C directly (see below).

We then decompose C as the sum of two components: one that contains the symmetric parts of all $T \times T$ cross-covariance matrices between all pairs of units which we call $C^{(+)}$, and one that contains the anti-symmetric counterparts, which we call $C^{(-)}$ (Fig. 1A - center). The latter summarizes the spatio-temporal organization of sequentiality in the data, at least up to second-order moments (i.e. at least for a Gaussian process). Specifically, we define:

$$C^{(+)} = C + \sigma\,(C) \tag{2}$$

$$C^{(-)} = C - \sigma\,(C) \tag{3}$$

where, for any matrix $M \in \mathbb{R}^{NT\times NT}$, we use the notation $\sigma\,(M)$ to denote the matrix obtained by separately transposing each $T \times T$ block of M (there are N^2 of them). We note that $C^{(-)}$ is symmetric but not positive definite (e.g. its trace is zero).

We use these two components to define the degree of sequentiality of a spatio-temporal process as

$$s = \frac{\|C^{(-)}\|_F^2}{\|C^{(+)}\|_F^2} \tag{4}$$

where $\|\|_F^2$ is the squared Frobenius norm. This metric is upper-bounded by one, and can be used to compare the sequentiality of various datasets.

We show that the spatio-temporal structure of sequentiality can be discovered based on an analysis of $C^{(-)}$'s principal subspace. The methodology is straightforward: we seek the top singular vectors of $C^{(-)}$, and reshape them as $T \times N$ prototypical time series (Fig. 1A, right); these form a subspace of spatio-temporal activity patterns that capture the non-sequential aspects of the organization of space-time cross-covariances in the data.

Scaling the method and cross-validation For large N and/or large T, computing (or even storing) $C^{(-)}$ is intractable. To bypass having to compute $C^{(-)}$ directly, we derive an efficient computation of matrix-vector products of the form $C^{(-)}v$ for any TN-vector $v = \mathrm{vec}\,(V)$:

$$C^{(-)}v = \mathrm{vec}\left(\left\langle \mathrm{Tr}[XV^T]X - XV^TX \right\rangle\right) \tag{5}$$

which can be estimated in $\mathcal{O}(\min(KNT^2, KTN^2))$ based on a number K of samples. Using this identity, it is possible to very efficiently compute the principal subspace of $C^{(-)}$ using either randomized SVD algorithms [Halko et al., 2011] or more standard iterative optimization methods for matrix factorization. Here, we directly minimize

$$\mathcal{C}(Z) = \|C^{(-)}C^{(-)T} - ZZ^T\|_F^2 \tag{6}$$

over matrix $Z \in \mathbb{R}^{NT\times r}$, where r is the desired rank; the gradient of $\mathcal{C}$ involves Equation 5.

We have also developed a bootstrapping method (inspired by Machens, 2010) to estimate the noise floor in the singular values of $C^{(-)}$ – this allows us to determine how many of the top modes contribute significantly.

Extensions Our iterative method based on Equation 6 allows constraints to be added such as sparsity. SCA with an ℓ_1 penalty on Z can demix sequences that occur in distinct populations (not shown here). We also show that the rotational dynamics model underlying jPCA [Churchland(2012)] assumes a particularly simple SCA decomposition: pure rotations in a plane spanned by two orthogonal vectors (x_1, x_2) (starting with a random phase in each trial) give rise to a $C^{(-)}$ of the form:

$$C^{(-)} = (x_1 x_2^T - x_2 x_1^T) \otimes (sc^T - cs^T) \quad (7)$$

where s and $c \in \mathbb{R}^T$ are a pair of sine and cosine temporal waveforms. In this case, the top two modes of $C^{(-)}$ correspond to rotational motion of the activity vector in the (x_1, x_2) plane. The equation above generalizes to multiple rotations in orthogonal spatial planes, in which case $C^{(-)}$ is the sum of multiple Kronecker products similar to Equation 7.

Illustrative toy example To illustrate the method, we construct a toy dataset comprised of one rank-2 sequential mode (pure rotation), overwhelmed by two non-sequential modes of much higher variance; then scaled by a transient envelope, plus Gaussian observation noise (Fig. 1B, left). Methods such as PCA can discard important sequential modes if these, for instance, ride on top of larger global fluctuations shared across all neurons. In contrast, SCA is impervious to non-sequential modes irrespective of their variance, and here successfully recovers the hidden sequence (Fig. 1B, right).

Sequences in M1 activity during reaching Movements require largely sequential activation of muscles, which presumably should be reflected in neural activity. We apply SCA to M1 activity (Churchland lab, $N = 218$ neurons) recorded in monkeys while they executed $K = 108$ different reaches (Fig. 1C, left). The data was first averaged across trials for each condition; thus, we analyzed across-conditions covariances. SCA here confirms that the data is highly sequential ($s = 0.48$), with several significant modes which, following reordering of the neurons by peak time in each mode, reveal clear sequences. While $C^{(-)}$'s singular vectors (seq. modes) are orthogonal, this does not in principle guarantee that they correspond to different sequences. In this case though, the sequences appear genuinely different: there is no apparent correlation in the neurons' rank order between pairs of modes (Fig. 1C, right).

Encoding of spatial memories in sequential subspace: hippocampus, navigational task We apply SCA to unpublished rat hippocampal data recorded during a spatial memory task in a three-arm maze (Fig. 1D, left). The data consists of $N = 315$ neurons recorded over the 8-second delay period of the task ($K = 100$ trials), during which the rats have to run in a wheel before being allowed back into the maze for the next trial. It must then choose one of the three maze arms to collect a reward, before returning to the wheel. Critically, the rat must visit the arms in a specific order, and must therefore remember where he came from in the previous trial in order to decide where to go next. The data (courtesy of Brian Lustig, HHMI Janelia) was collected using silicone probes, binned in 20ms bins and filtered with a Gaussian kernel. We compared projections of the data into the top two SCA and PCA components (Fig. 1D, right). The sequential subspace better separates memories of recent navigation episodes, providing evidence that the brain may use sequential dynamics to encode such episodes in the hippocampus.

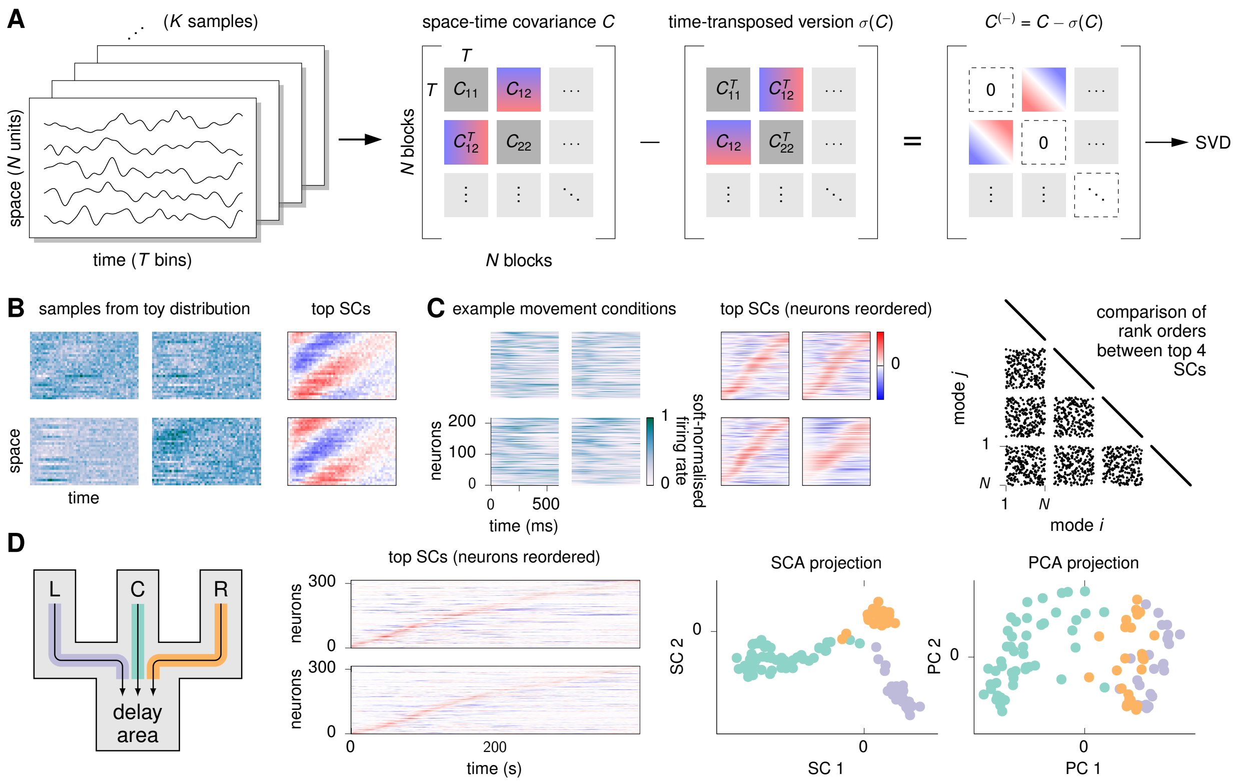

Figure 1. Please see text for full details. The SCA workflow (A) is applied to a toy dataset (B), to monkey M1 data (C), and to a delayed three-arm spatial memory task demonstrating that memories of recent navigation episodes are better separated in the top sequential subspace (D).