

A geometry of spike sequences: Fast, unsupervised discovery of high-dimensional neural spiking patterns based on optimal transport theory.

Boris Sotomayor-Gómez¹, Francesco P. Battaglia², Martin Vinck^{1*}

1 Ernst Strüngmann Institute (ESI) for Neuroscience in Cooperation with Max Planck Society, Frankfurt, Germany

2 Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Nijmegen, Netherlands

* martin.vinck@esi-frankfurt.de

Abstract

Neuronal coding and memory formation depend on temporal activation patterns spanning high-dimensional ensembles of neurons. To characterize these high-dimensional spike sequences, it is critical to measure their dissimilarity across different epochs (e.g. stimuli, brain states) in terms of all the relative spike-timing relationships. Such a dissimilarity measure can then be used for dimensionality reduction, clustering of sequences or decoding. Here, we present a new measure of dissimilarity between multi-neuron spike sequences based on optimal transport theory, called *SpikeShip*. SpikeShip computes the optimal transport cost (Earth Mover’s Distance) to make all the relative spike-timing relationships (across neurons) identical between two spiking patterns. It achieves this by computing the optimal transport of spikes across time, per neuron separately, and then decomposing this transport cost in a temporal rigid translation term and a vector of neuron-specific transport flows. This yields a suitable geometry for spike sequences. SpikeShip can be effectively computed for high-dimensional neuronal ensembles and has linear $\mathcal{O}(N)$ cost. Furthermore, it is explicitly based on the higher-order structure in the spiking patterns. SpikeShip opens new avenues for studying neural coding and memory consolidation by finding patterns in high-dimensional neural ensembles.

Introduction

Information in the brain is encoded by very high-dimensional “ensembles” of neurons. These neurons encode information by firing nerve impulses (“spikes”) that are discrete all-or-nothing events. Populations of neurons can produce specific spike patterns depending on sensory inputs or internal variables [1–7]. With the development of new recording techniques like Neuropixels [8], it is now possible to simultaneously record from thousands of single neurons with a high temporal resolution. These techniques offer new opportunities to uncover the relationship between high-dimensional multi-neuron

spiking patterns and sensory or internal variables. They also pose unique mathematical challenges and opportunities for the unsupervised discovery of the “dictionary” of neuronal “code-words”. For instance, how do we measure the similarity between two multi-neuron spiking patterns in terms of their temporal structure (spike sequence)?

In general, any notion of information encoding relies on the construction of a distance or dissimilarity measure in an N -dimensional space. For example, the distance between binary strings can be measured using the Hamming distance, which is an essential mathematical construct in error-correcting coding. In the brain, the distance between two multi-neuron spiking patterns is conventionally measured in two steps: (1) By computing the number of spikes/sec (firing rates) for each neuron in a certain time window; and (2) computing e.g. the angle or Euclidean distance between multi-neuron rate vectors. Using this method, it has been shown for example that high-dimensional neural ensembles span a low-dimensional manifold that relates to stimulus or behavioral variables in a meaningful way [9,10]. However, the computation of firing rates disregards potentially rich information contained by the precise temporal order in which spikes are fired by multiple neurons (e.g. neuron i firing at time t and neuron j firing at $t + \tau$). For instance, we expect that any time-varying sensory stimulus (motion) or action sequence may be encoded by a unique multi-neuron temporal pattern of spiking. It has been shown that multi-neuron temporal sequences encode information about sensory stimuli or are required for the generation of complex motor patterns like bird songs [1–7,11]. Temporal sequences may also be critical for memory formation, because the plasticity rules in the brain are highly sensitive to the temporal order in which neurons fire spikes [12–14].

We have previously developed a dissimilarity measure between multi-neuron temporal spiking patterns called SPOTDis (Spike Pattern Optimal Transport Dissimilarity). This measure is defined as the minimum energy (optimal transport) that is needed to transform all pairwise cross-correlations of one epoch k into the pairwise cross-correlations of another epoch m . We showed that, due to employment of optimal transport, SPOTDis has several attractive features (see Results) [15], and can be combined with unsupervised dimensionality-reduction techniques like t-SNE or clustering techniques like HDBSCAN to characterize the state-space of multi-neuron temporal sequences [15]. However, the SPOTDis dissimilarity measure has two principal weaknesses: (1) Its computational complexity, for comparing two time epochs, is $\mathcal{O}(N^2)$, where N is the number of neurons. This becomes a major problem for computing an $M \times M$ dissimilarity matrix (for M time epochs) for thousands of neurons. (2) It relies exclusively on pairwise spike-timing relationships (i.e 2nd order correlations), which means that it might not be sensitive to higher-order correlations in spiking patterns; in fact we provide some examples demonstrating that SPOTDis does not form a proper geometry for spike sequences. This relates to the fact that SPOTDis does not solve the optimal transport problem for the entire spike pattern, but only for neuron pairs separately.

Here, we develop a dissimilarity measure for multi-neuron spiking patterns called SpikeShip, which has linear computational complexity of $\mathcal{O}(N)$. We achieve this by (1) computing the minimum transport cost of spikes for each spike train separately, and (2) discounting a global translation term in the transport flow across neurons. SpikeShip can be interpreted as the optimal transport cost to make all spike-timing relationships between two different multi-neuron spiking patterns identical; i.e. it solves the optimal transport problem for the entire spiking pattern. It is thus sensitive to the higher-order structure of spike trains.

Results

Our overall goal is to develop a dissimilarity measure between multi-neuron spiking patterns that is exclusively based on the temporal order of spiking (i.e. not on the spike

count). Suppose that we have measured the spikes from N neurons (“spike trains”). Spike trains are divided into M epochs of length T (measured in seconds or samples). Epochs could be defined by e.g. trials (e.g. stimulus presentations) or sliding windows [15]. The problem is to find a dissimilarity measure with the following properties [15]:

1. The measure should depend on the temporal order of firing across neurons, but not on the spike count, although the dissimilarity measure below can be generalized to incorporate spike count differences.
2. If two spike patterns are identical in terms of cross-neuron spike timing relationships (i.e. temporally translated version of one another), then the dissimilarity measure should equal zero.
3. The measure does not require binning or smoothing, and is based on the exact timing of the spikes.
4. It should allow for the discovery of more patterns than the number of neurons, when used in combination with clustering techniques like HDBSCAN.
5. It should measure dissimilarity in a continuous way and avoid the problem of “combinatorial explosion”. This problem occurs with methods that search for exact matches in spiking patterns, if the number of neurons is large.

Previous work: SPOTDis

[15] developed a dissimilarity measure called SPOTDis, which is based on optimal transport theory. We briefly summarize the computing steps of SPOTDis (see Methods):

1. Compute cross-correlation histograms for all N^2 pairs (i, j) of two neurons, $s_{ij}^{\text{norm}}(k)$. These cross-correlations capture the information about the spike-timing relationships across all pairs of neurons.
2. Normalize these histograms to unit mass, which is needed for the optimal transport measure, and removes the effect of spike count.
3. Compute, for each pair of neurons (i, j) and pair of epochs (m, k) , the minimum transport cost (Earth Mover’s Distance) to transform $s_{ij,k}^{\text{norm}}$ into $s_{ij,m}^{\text{norm}}$. This yields a dissimilarity measure denoted $D_{ij,km}$.
4. Compute the average of this dissimilarity measure across all pairs of neurons, yielding a dissimilarity measure D_{km} called SPOTDis.

Using simulations, we showed that SPOTDis can be used to find distinct clusters of multi-neuron spiking patterns in an unsupervised way [15].

As discussed in [15], SPOTDis satisfies the six requirements listed above. In particular, note that the EMD can be computed based on exact spike times and does not require binning or smoothing [15]. The use of the EMD for multi-neuron spike train dissimilarity has several additional advantages, as explained in [15]: The EMD is a symmetric and metric measure of similarity between two probability distributions (unlike the KL-divergence). In contrast to KL divergence, it does not quantify whether two distributions are overlapping in *the same bins*; rather, it examines how far distributions are shifted away from each *across bins* in a metric space (in this case: time). Different from KL, EMD is well behaved both when comparing continuous and discrete (sum of delta functions) distributions. Further, the EMD does not rely on the computation of a measure like the center of mass or the peak time of firing activation, but can deal with multi-modal spiking patterns [15]. In other words it does not require any prior specification of features. Finally, as shown in [15], SPOTDis is highly noise robust (i.e. to insertion of noise spikes or spike-time jitter) because of the use of EMD.

SpikeShip

There are two main problems with the construction of SPOTDis: First, its computational complexity scales as $\mathcal{O}((Nn)^2)$, where n is the average number of spikes per epoch. Note that this computation typically needs to be carried out M^2 times in order to create a dissimilarity matrix for all epoch-to-epoch comparisons. This means that in practice, computation of SPOTDis might not be feasible when N is on the order of thousands, which is nowadays achievable with modern recording technologies. Second, as discussed in [15], SPOTDis is based exclusively on the pairwise cross-correlations between neurons. Hence, it may be insensitive to the higher-order correlations in the spike trains.




Fig 1. Illustration of SpikeShip Example of two epochs with spike times $t_k = (10, 10, 10, 10, 10)$ and $t_m = (20, 30, 35, 45, 50, 60)$ (note only one spike per neuron in this example). The vector \vec{c} contains the differences of spike times t_k and t_m . The median of \vec{c} , $g_{\min} = 30$ is the optimal global shift such that $f_i = c_i - g_{\min}$. The neuron-specific shifts $\vec{f} = (-20, -10, -5, 5, 10, 20)$ contain all the information about the structure of distances between t_k and t_m . SpikeShip equals $F_{km} = \frac{1}{6} \sum_i^N |f_i| = \frac{70}{6}$.

To address these problems we introduce the SpikeShip measure (see Methods). SpikeShip yields a dissimilarity between multi-neuron spike train patterns that has similar characteristics as SPOTDis. But, it has the following desirable properties:

1. Its computational complexity grows linearly with the number of neurons and spikes, i.e. $\mathcal{O}(Nn)$.
2. By construction, it is sensitive to higher-order spike-timing relationships.
3. It allows for the decomposition of transport cost into a temporal rigid translation term and neuron-specific flow terms.

SpikeShip is computed in two steps (see Methods, Figure 1):

1. Compute the optimal transport flow to transform a spiking pattern in epoch k into the spiking pattern in epoch m . This flow vectors are computed for each neuron separately (i.e. in the spike raster plot, mass from spikes can be moved horizontally, but not vertically). This is performed in two steps: (i) Assign a mass to each spike, such that the total mass for the i th neuron in epoch k and m is equal; (ii) compute for each spike, how its mass should be shifted in order to make the spike trains in the window T completely identical.
2. By using an expression for a global temporal translation term (across neurons) and subtracting this from all the spike shifts, we minimize the overall transport

cost. This transport cost amounts to the amount of energy to shift the individual spikes such that all spike-timing relationships across neurons become identical.

When we compute the EMD in step 1, we can compute the total transport in discrete steps, shifting mass from a spike in pattern k to mass of a spike in pattern m . This can be algorithmically solved, for each neuron, in approximately order $\mathcal{O}(n)$. We denote the discrete steps $c_{i,u}$ with moved mass $w_{i,u}$, such that $\sum_u w_{i,u} = 1$. From these flows computed in Step 1, we can then subtract a global rigid translation term. This uniquely yields the minimum transport cost to transform the multi-neuron spiking pattern in epoch k such that its spike-timing relationships are identical to another pattern m . In the Methods section we state our main result, namely that the optimal transport flows are given by

$$f_{i,u} = c_{i,u} - g^{\min}. \quad (1)$$

Here g^{\min} is the weighted median across all the original flows $c_{i,u}$ with associated mass $w_{i,u}$. Thus we can decompose the transport flow in two terms: (1) an optimal global shift between two epochs, shared across all neurons; and (2) an optimal neuron-specific transport flow based on the EMD method. We then define SpikeShip (see Methods) as

$$F_{km} \equiv \frac{1}{|\mathcal{A}_{km}|} \sum_{i \in \mathcal{A}_{km}} w_{i,u} \sum_{u=1}^{n_i^*} |f_{i,u}|. \quad (2)$$

Here \mathcal{A}_{km} is the set of all neurons that fired a spike both in epoch k and m . The weight $w_{i,u} = \frac{1}{n_i^*}$, effectively assigns an equal weight to each neuron that contains at least one spike. The algorithm to compute SpikeShip has computational complexity $\mathcal{O}(Nn)$, because the weighted median has complexity $\mathcal{O}(N)$ (Figure 2). This result is, in a way, surprising: One would expect that for a dissimilarity measure that is based on the relative spike-timing between neurons and is sensitive to higher-order correlations, we would need a computational complexity of at least $\mathcal{O}(N^2)$.

Relationship between SpikeShip and SPOTDis

We show in the Methods section that there exists an analytical relationship between SpikeShip and SPOTDis for a specific class of sequences: Suppose that (1) in each temporal sequence, a neuron is active in only in a small part of the time window T (or fires only one spike), and (2) that the temporal sequence is uniformly spanned across neurons (i.e. a uniform distribution across neurons). The relationship between SPOTDis and SpikeShip (see Figure 2) then equals

$$D_{km} \approx \frac{7}{5} F_{km}. \quad (3)$$

Thus, for these kind of sequences, the total transport cost F_{km} , which equalizes all the higher-order correlations in the spike patterns, can be related to the transport cost over all the pairwise spike-timing relationships (SPOTDis). However, an analytical relationship may not exist for the case of multimodal spiking patterns or other types of distributions (see e.g. Example 1).

SpikeShip is a more sensible dissimilarity measure than SPOTDis because it is explicitly sensitive to higher-order structure in the spike train. In addition, it is $\mathcal{O}(Nn)$ faster. The difference between SpikeShip and SPOTDis can be further illustrated with the following simple examples.

Example 1: Suppose that we have a vector of spike times $\vec{t}_k = (0, 0, 0, 0)$ and $\vec{t}_m = (-20, 0, 0, +20)$. We now have cost $\vec{f} \equiv (-20, 0, 0, +20)$ and $F_{km} = \frac{40}{4} = 10$. The

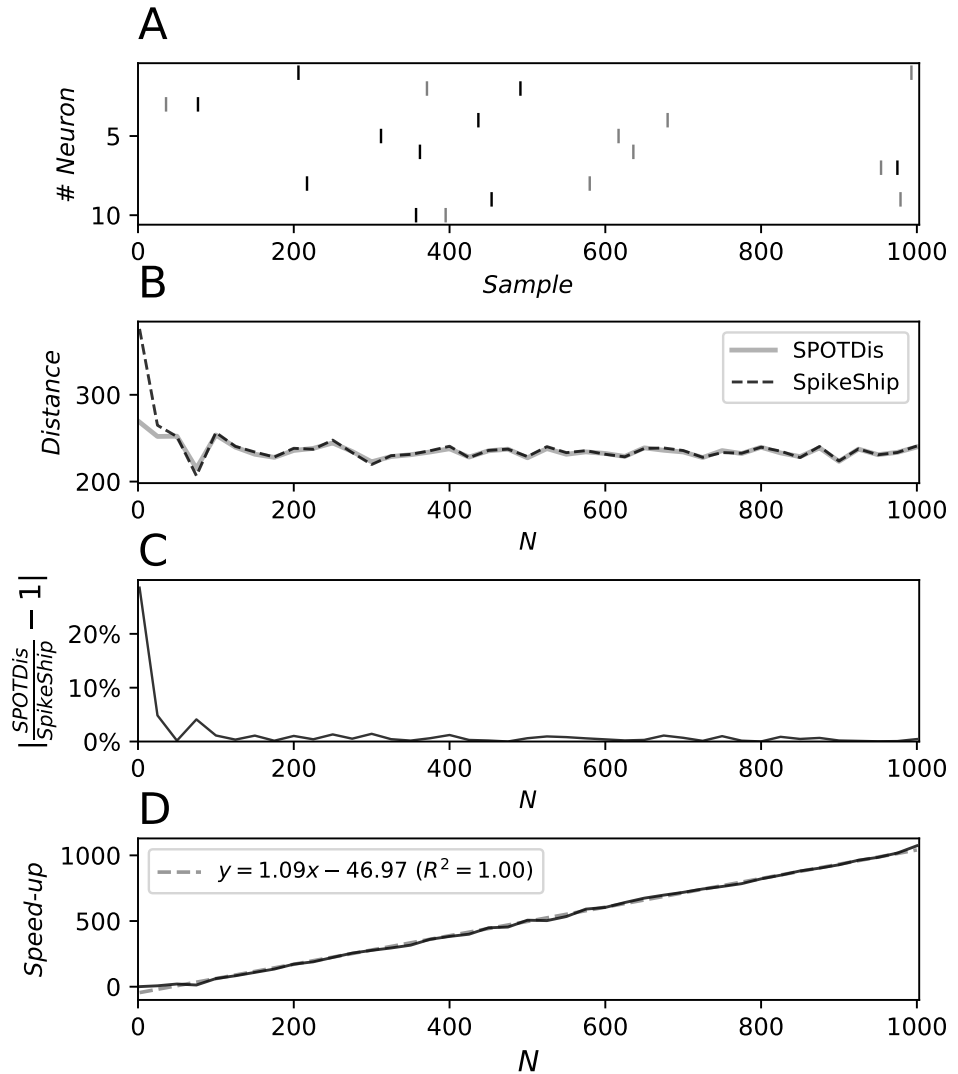


Fig 2. Accuracy and Speed-up comparison for single-spike patterns ($n = 1$).

A) Example of single spike trains for two epochs for 10 neurons. Patterns were generated as uniform sequences with 1 spike per neuron per epoch. B) Values of SPOTDis and SpikeShip as a function of the number of neurons N . The value of SpikeShip is scaled by a factor $\frac{7}{5}$ based on the analytical relationship for patterns with a single spike per neuron. Note that this simulation confirms the analytical relationship between SPOTDis and SpikeShip (see Methods). C) Difference between SpikeShip vs. SPOTDis (scaled by same factor as in (B)) expressed as percentage. D) Computational speed-up for SpikeShip vs. SPOTDis, which is approximately N when there is 1 spike per neuron.

SPOTDis in this case equals

$$\frac{2 \times (0 + 20 + 20 + 40) + 2 \times (0 + 0 + 20 + 20)}{16} = 15. \quad (4)$$

Consider now another vector of spike times $\vec{t}_k = (0, 0, 0, 0)$ and $\vec{t}_m = (-15, -15, +15, +15)$. We now have cost $\vec{f} \equiv (-15, -15, +15, +15)$ and, different from above, $F_{km} = \frac{60}{4} = 15$. The SPOTDis in this case equals

$$\frac{4 \times (0 + 0 + 30 + 30)}{16} = 15, \quad (5)$$

In this case, the mover costs on the 2nd order pairwise timing relationships (SPOTDis) are the same for both sequences; however, SpikeShip correctly reflects the increased mover cost that is needed to match the two patterns to the zero pattern. We can consider further the distance between the two sequences. In this case, SpikeShip equals $F_{km} = \frac{40}{4} = 10$. However, SPOTDis now equals $D_{km} = \frac{280}{16} = 17.5$. Thus, SPOTDis assigns a disproportionately large transport cost between the first and the second sequence as compared to the $(0, 0, 0, 0)$ sequence 3.




Fig 3. Example of SPOTDis and SpikeShip. Example of three single-spike patterns: $(-20, 0, 0, +20)$, $(0, 0, 0, 0)$, and $(-15, -15, +15, +15)$, from left to right. SpikeShip assigns a geometrically more appropriate transport cost between pattern 1 and 2 ($F_{1,2} = 10$) than SPOTDis ($D_{1,2} = 17.5$), considering their distance with pattern 0.

Dissimilarity matrix and cluster visualization

We compute SpikeShip to detect unique, recurring spiking patterns generated as in [15] (Figure 4). The dataset contains $P = 7$ patterns across $N = 100$ neurons and $M = 360$ epochs. Low-dimensional embedding was performed with t-SNE as in [15] (Fig 4B). Clusters were derived using HDBSCAN and are plotted in a 2-dimensional space. These matrices were computed using a workstation with 12 CPU cores (DUAL Xeon E5646 @ 2.40GHz). The computation time for SpikeShip and SPOTDis dissimilarity matrices were 28.2 sec and 20186.2 secs (5 hours, 36 minutes, and 26.2 secs), respectively. This is

a speedup of about 700 times (which is higher than $N = 100$ due to presence of multiple spikes).

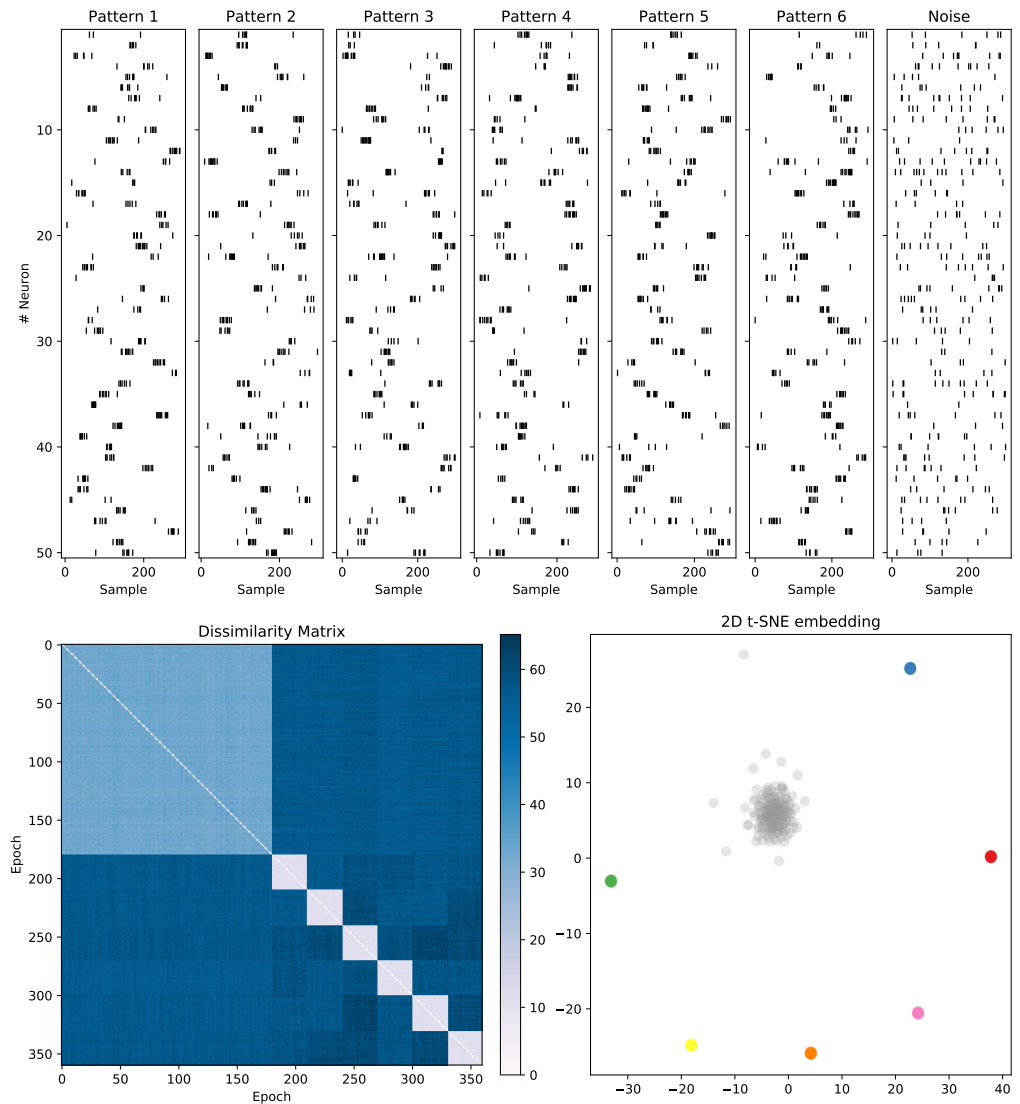


Fig 4. Dissimilarity matrices and clustering comparison. Top: Samples from simulated neuronal outputs according to an inhomogenous Poisson process (using scripts provided by [15]). Bottom: Dissimilarity matrix using SpikeShip dissimilarity measures (left), and a 2-dimensional t-SNE embedding using SpikeShip dissimilarity matrix (right).

Discussion

We studied the problem of measuring the dissimilarity between two multi-neuron spiking patterns based on the relative spike-timing relationships across neurons (i.e. firing order). We developed a new measure called *SpikeShip*. SpikeShip solves the problem of optimally transporting the spikes of individual neurons, such that the global pattern of spike-timing relationships becomes identical between two epochs. Intuitively, one would think that such a measure has computational complexity of at least $\mathcal{O}(N^2)$, but we show that it can be computed with $\mathcal{O}(N)$ computational complexity; this is a major improvement

over our previous work [15]. We show that a dissimilarity between two spiking patterns can be decomposed into neuron-specific flows and a temporal rigid translation term. Compared to our previous measure SPOTDis [15], SpikeShip is not restricted to the 2nd order correlations, but is based higher-order structure in the spike train.

Our technique can also be used to align temporal sequences in an unsupervised way if there is a global jitter between spike patterns. This achieves a similar goal as a recent study in Neuron [16], described as “warping”; however, the method in our paper does not require any parameter estimation; furthermore, it can also be used to align trials corresponding to different stimuli or behavioral conditions. For example, suppose that we want to analyze the sequence of neural activation in relationship to some event, e.g. freely moving your arm. Suppose that we do not know the exact onset of the neuronal sequence in relationship to moving your arm, and we record multiple realizations of this event. With SpikeShip, we can immediately decompose the transport cost between any two trials in terms of a global translation term (which is a non-linear computation) and a neuron-specific shift. This allows one to align trials that were recorded in different conditions, e.g. moving your arm left or right. And it yields, in essence, a state-space characterization of temporal trajectories (e.g. related to different movements), rather than a trajectory through a state-space of firing rate vectors.

SpikeShip opens new avenues to study temporal sequences in high-dimensional neural ensembles. Recent technological developments now allow for recordings of thousands of neurons simultaneously, either using electrical recordings or two-photon imaging [8]. The technique developed here is applicable to both kinds of data, due to linear computation time. Application of SpikeShip to such kind of data might generate important insights into the role of temporal sequences in sensory coding and memory consolidation.

Materials and methods

Previous work: SPOTDis

We previously defined a dissimilarity measure between multi-neuron spike train patterns called SPOTDis [15]. We give a brief overview of the definition of the SPOTDis measure; for a detailed account we refer to [15]. Suppose we record spike trains of N neurons in M separate epochs (e.g. trials); we refer to these as “spike patterns”. SPOTDis was defined as follows: We first compute, for all M epochs separately, a normalized cross-correlation function for all N^2 pairs of neurons,

$$s_{ij,k}^{\text{norm}}(\tau) \equiv \frac{s_{ij,k}(\tau)}{\sum_{\tau=-T}^T s_{ij,k}(\tau)}, \quad (6)$$

where $s_{ij,k}(\tau) \equiv \sum_{t=-T}^T s_{i,k}(t)s_{j,k}(t+\tau)$ is the cross-correlation function (or cross-covariance), and $s_{i,k}(t)$ and $s_{j,k}(t)$ are the spike trains of neurons i and j in the k th trial. The normalization of the cross-correlation function is necessary to avoid an influence of firing rate fluctuations, and to compute the EMD. Next, we compute the Earth Mover Distance (EMD) between the normalized cross-correlograms $s_{ij,k}^{\text{norm}}(\tau)$ and $s_{ij,m}^{\text{norm}}(\tau)$ for each (k, m) . This yields $M(M-1)/2 \times N(N-1)/2$ EMD values $D_{ij,km}$. In [15], we used the L1 norm to measure dissimilarity on the time axis, but the L2 norm could be used as well. After computing the EMDs between each pair of epochs for each neuron pair separately, we compute SPOTDis as the average EMD across neuron pairs,

$$\bar{D}_{km} \equiv \frac{1}{|\mathcal{A}_{km}|} \sum_{(i,j) \in \mathcal{A}_{km}} D_{ij,km}. \quad (7)$$

Here, $\mathcal{A}_{km} \equiv \{(i, j) : n_{i,m} > 0 \wedge n_{i,k} > 0 \wedge n_{j,m} > 0 \wedge n_{j,k} > 0\}$, i.e. the set of all pairs of neurons that are active in both epochs k and m . We ignored the other neuron pairs (i.e. not contained in this set) for computing the SPOTDis, because it avoids assigning an arbitrary value to the EMD when there is no information about the temporal relationship between the neurons (i.e. when we do not have any spikes for one neuron in one epoch).

Derivation of SpikeShip

A major problem with the computation of SPOTDis (for all pairs of epochs) is that it has computational complexity of $\mathcal{O}(N^2 M^2)$. Furthermore, because it is based on 2nd order correlations, it is not explicitly sensitive to higher-order correlations. Here we derive a new dissimilarity measure, called SpikeShip, which has computational complexity $\mathcal{O}(N)$ for one pair of epochs. To derive this measure, we first consider the simplified case where each i th neuron fires one spike for all M epochs.

SpikeShip for single spike per epoch

Let $\vec{f}_{km} = (f_{1,km}, \dots, f_{N,km})$ be a vector of flows for each neuron in epoch m . In other words, $f_{1,km}$ is the shift of the spike fired by the first neuron in the m th epoch. The total moving cost equals

$$F_{km} \equiv \frac{1}{N} \sum_{i=1}^N |f_{i,km}|. \quad (8)$$

The problem statement is to find a flow vector \vec{f}_{km} such that after moving the spikes, the resulting spike train patterns are identical between epoch k and epoch m , in terms of the full matrix of spike timing relationships. We wish to find the flow vector that satisfies this constraint with minimum cost F_{km} .

Example: Consider the following example: Suppose there are two epochs for $N = 6$ neurons with spike times $\vec{t}_k = (10, 10, 10, 10, 10, 10)$ and $\vec{t}_m = (20, 30, 35, 45, 50, 60)$. We will show that the flow vector with minimum cost, such that spike patterns have identical temporal structure, equals $\vec{f}_{km} = (-20, -10, -5, 5, 10, 20)$ (Figure 1).

More formally, let $t_{m,i}$ be the timing of the spike for the i th neuron in the m th epoch. We denote the spike times after (post) moving them in epoch m as

$$t_{m,i}^{\text{post}} \equiv t_{m,i} - f_i \quad (9)$$

$\forall i$, where we omitted the subscripts k, m from the variable $f_{i,km}$ for simplicity. The constraint that all the across-neuron spike timing relationships should be identical after moving implies that

$$t_{m,i}^{\text{post}} - t_{m,j}^{\text{post}} = t_{k,i} - t_{k,j} \quad \forall i, j. \quad (10)$$

In other words, the delay between two spikes from two different neurons (i, j) should be identical between pattern k and m after moving the spikes. Substituting based on Eq. 9, this can be expressed as

$$(t_{m,i} - f_i) - (t_{m,j} - f_j) = t_{k,i} - t_{k,j}. \quad (11)$$

Let c_i be the shift in spike timing for each neuron in epoch m , such that the spike train patterns become identical,

$$c_i \equiv t_{m,i} - t_{k,i}. \quad (12)$$

Note that with this definition of c_i , the equation

$$(t_{m,i} - c_i) - (t_{m,j} - c_j) = t_{k,i} - t_{k,j} \quad (13)$$

holds for all (i, j) . For all i , we can express f_i as a function of the shift c_i , such that

$$f_i \equiv c_i - g_i. \quad (14)$$

We wish to solve

$$\vec{g}_{\min} \equiv \arg \min_{\vec{g}} F_{km} = \arg \min_{\vec{g}} \sum_i^N |c_i - g_i| \quad (15)$$

under the constraint of Eq 11.

From Eqs 11, 13 and 14 it follows that

$$(t_{m,i} - c_i + g_i) - (t_{m,j} - c_j + g_j) = (t_{m,i} - c_i) - (t_{m,j} - c_j) \quad (16)$$

Hence the equation $g_i = g_j$ holds for all (i, j) . Thus, we can rewrite Eq. 27 to

$$g^{\min} \equiv \arg \min_g \sum_i^N |c_i - g| \quad (17)$$

to find a global shift that minimizes the L1-norm of the residuals. The solution to this equation is

$$g^{\min} = \text{Median}\{\vec{c}\}. \quad (18)$$

Thus, our main result is that the original shifts between the two spiking patterns can be written as the decomposition

$$c_i = g^{\min} + f_i, \quad (19)$$

for all i , i.e. the optimal transport in terms of neuron-specific shifts and a global temporal rigid translation term. Thus, the optimal transport between two spiking patterns (e.g. after stimulus onset) can be decomposed into the optimal transport in terms of neuron-specific shifts and a global temporal rigid translation term.

Global shift definition for multiple spikes

We now consider the case where in each epoch, every neuron fires a variable number of spikes. We will show that a similar derivation for SpikeShip can be made based on the weighted median. Let $n_{i,k}$ and $n_{i,m}$ be the number of spikes for the i -th neuron in epoch k and m . In the SPOTDis, we normalized the cross-correlations to unit mass, such that each spike effectively gets the same weight. For SpikeShip, we will also assign a weight to each spike, such that the total weight per neuron in the computation of SpikeShip is equal. In order to do so, we first find the smallest common multiple of the spike counts for each neuron, denoted n_i^* . For instance, if $n_{i,k} = 2$ and $n_{i,m} = 6$ then the least common multiple equals 6. We now replicate each spike $\frac{n_i^*}{n_{i,k}}$ and $\frac{n_i^*}{n_{i,m}}$ times. Each spike will now have a weight of $w_i \equiv \frac{1}{n_i^*}$. (Note that in the actual computation, we do not replicate the spikes in practice, but use an algorithm similar to the one detailed in [15]). Then, based on the optimal transport cost (EMD), we obtain shifts $c_{i,u}$ for the u -th spike of the i -th neuron, $u = (1, \dots, n_i^*)$. Note that computation of the EMD amounts to shifting the most left-ward spike (i.e. first spike) out of n_i^* spikes in epoch m to the most left-ward spike in epoch k [15].

We use a similar derivation as the one above. Let $t_{m,i,u}$ be the timing of the u -th replicated spike for neuron i in the m th epoch, for all (i, m, u) , $u = (1, \dots, n_i^*)$. We denote the spike times after (post) moving them in epoch m as

$$t_{m,i,u}^{\text{post}} \equiv t_{m,i,u} - f_{i,u} \quad (20)$$

$\forall(i, u)$, where we omitted the subscripts k, m from the variable $f_{i,k,m,u}$ for simplicity. The constraint that all the across-neuron spike timing relationships should be identical after moving implies that

$$t_{m,i,u}^{\text{post}} - t_{m,j,u}^{\text{post}} = t_{k,i,u} - t_{k,j,u} \quad (21)$$

In other words, the delay between two spikes from two different neurons (i, j) should be identical between pattern k and m after moving the spikes. Substituting based on Eq. 20, this can be expressed as

$$(t_{m,i,u} - f_{i,u}) - (t_{m,j,u} - f_{j,u}) = t_{k,i,u} - t_{k,j,u}. \quad (22)$$

There is one additional constraint, namely

$$(t_{m,i,u} - f_{i,u}) - (t_{m,i,v} - f_{i,v}) = t_{k,i,u} - t_{k,i,v}. \quad (23)$$

In other words, all the pairwise relationships after moving within the same neuron should be identical.

Let $c_{i,u}$ be the shift in spike timing for each neuron in epoch m , such that the spiking patterns in the window become identical,

$$c_{i,u} \equiv t_{m,i,u} - t_{k,i,u}. \quad (24)$$

Note that with this definition of $c_{i,u}$, the equation

$$(t_{m,i,u} - c_{i,u}) - (t_{m,j,u} - c_{j,u}) = t_{k,i,u} - t_{k,j,u} \quad (25)$$

holds for all (i, j, u) . For all i, u , we can express $f_{i,u}$ as a function of the shift $c_{i,u}$, such that

$$f_{i,u} \equiv c_{i,u} - g_{i,u}. \quad (26)$$

Define $\vec{g} \equiv ((g_{i,1}, \dots, g_{i,n_i^*}), \dots, (g_{N,1}, \dots, g_{N,n_N^*}))$.

We wish to solve

$$\vec{g}_{\min} \equiv \arg \min_{\vec{g}} \sum_i^N \sum_u^{n_i^*} w_{i,u} |c_{i,u} - g_{i,u}| \quad (27)$$

under the constraints of Eq 22, 23 and 25.

Given these two constraints, the equation $g_{i,u} = g_{j,u}$ holds for all (i, j, u) . Thus, we can rewrite Eq. 27 to

$$g^{\min} \equiv \arg \min_g \sum_i^N \sum_u^{n_i^*} w_{i,u} |c_{i,u} - g| \quad (28)$$

to find a global translation term that minimizes the L1-norm of the residuals. The solution to this this equation is the weighted median g ,

$$g^{\min} = \text{WeightedMedian}((c_{1,1}, w_{1,1}, \dots, c_{1,n_1^*}, w_{1,n_1^*}), \dots, (c_{N,1}, w_{N,1}, \dots, c_{N,n_N^*}, w_{N,n_N^*})).$$

Now average mover cost, *SpikeShip*, equals

$$F_{km} \equiv \frac{1}{|\mathcal{A}_{km}|} \sum_{i \in \mathcal{A}_{km}} \frac{1}{n_i^*} \sum_u^{n_i^*} |f_{i,u}|. \quad (29)$$

where $\mathcal{A}_{km} \equiv \{i : n_{i,m} > 0 \wedge n_{i,k} > 0\}$ is the set of neurons that are active in both epochs k and m . Thus, we derive the result that the original shifts between the two spiking patterns can be written as the decomposition

$$c_{i,u} = g^{\min} + f_{i,u}, \quad (30)$$

for all (i, u) , i.e. the optimal transport in terms of neuron-specific shifts and a global translation term. Thus, the optimal transport between two spiking patterns (e.g. after stimulus onset) can be decomposed into the optimal transport in terms of neuron-specific shifts and a global temporal rigid translation term. This computation has linear complexity $\mathcal{O}(N)$. We base the computation of the weighted median on [17, 18] and an adaptation of the *Robustats* [19] Python library. In practice, the computation of $c_{i,u}$ (optimal transport per neuron) can be efficiently performed not by replicating the spikes to the common multiple integer (used in derivation above), but by using a similar algorithm as in [15].

Relationship SpikeShip and SPOTDis

We now derive an analytical relationship between SpikeShip and SPOTDis, for the case of uniform sequences in which each neuron fires only at one time point of the sequence; note that this result does not generalize to more complex distributions, or the case of multiple spikes. We note that the pairwise distance among the cross-correlation for four spikes (neuron i and j , epoch k and m) equals $|(t_{m,i} - t_{m,j}) - (t_{k,i} - t_{k,j})| = |c_i - c_j|$. For simplicity we define SPOTDis over all N^2 elements (note the distance of an element with itself equals zero) and assume that all neurons fire a spike, without loss of generalization. From Eq. 31 we can rewrite the SPOTDis as

$$\begin{aligned} D_{km} &\equiv \frac{1}{N^2} \sum_{i,j} |c_i - c_j| \\ &= \frac{1}{N^2} \sum_{i,j} |f_i - f_j|. \end{aligned} \quad (31)$$

We have as expected value $\mathbb{E}\{D_{km}\} = \mathbb{E}\{|f_i - f_j|\}$. Because we defined $f_i = c_i - g^{\min}$, where $g^{\min} = \text{Median}\{\bar{c}\}$, we know that the median of our distribution is zero. We can thus assess the relationship between SpikeShip and SPOTDis for different distributions of f_i . We can divide the set of N (assuming N is even, without loss of generalization) neurons into two sets, namely \mathcal{I} and \mathcal{J} , with $\mathcal{I} \equiv \{i : f_i \leq 0\}$ and $\mathcal{J} \equiv \{i : f_i \geq 0\}$. For the combinations (i, j) for which $i \in \mathcal{I}$ and $j \in \mathcal{J}$ we obtain the contribution to SPOTDis as

$$\begin{aligned} D_{km}^{pos,neg} &= \frac{1}{2} \mathbb{E}\{|f_i - f_j|\} \\ &= \frac{1}{2} \mathbb{E}\{|f_j| + |f_i|\} \\ &= \mathbb{E}\{F_{km}\} \end{aligned} \quad (32)$$

where the factor $\frac{1}{2}$ reflects the probability of f_i and f_j having opposite signs. Thus, the term $D_{km}^{pos,neg}$ does not depend on the specific distribution of the spikes.

For the combinations (i, j) for which $i \in \mathcal{I}$ and $j \in \mathcal{J}$ the expected distance depends on the distribution of f_i . If we assume that the different spike patterns each uniformly cover a window length of $2T$, then we know that f_i will have a triangular probability distribution, i.e.

$$p(f_i) = \frac{T - |f_i|}{T^2}, \quad -T \leq f_i \leq T. \quad (33)$$

We now obtain

$$\mathbb{E}\{D_{km}^{pos,pos} + D_{km}^{neg,neg}\} = \frac{1}{2} \int_0^T \int_0^{f_i} 8 \frac{T - f_i}{T^2} \frac{T - f_j}{T^2} (f_i - f_j) df_j df_i = \frac{2T}{15} \quad (34)$$

Note that for this distribution we also have $\mathbb{E}\{|f_i|\} = \frac{T}{3}$. Hence,

$$\mathbb{E}\{D_{km}\} = \mathbb{E}\left\{\left(1 + \frac{2}{5}\right) F_{km}\right\}. \quad (35)$$

Thus, the average amount of transport on the individual spikes is smaller ($\frac{5}{7}$) than the average transport on the pairs of spikes.

Acknowledgements

This project was financed by the BMF (Bundesministerium fuer Bildung und Forschung), Computational Life Sciences, project BINDA (031L0167).

Authorship Contributions

Conception and problem statement: BSG, FPB, MV. Data analysis and simulations: BSG. Mathematical analysis: BSG and MV. Supervision: MV. Writing of main draft: BSG and MV.

References

1. Havenith MN, Yu S, Biederlack J, Chen NH, Singer W, Nikolic D. Synchrony makes neurons fire in sequence, and stimulus properties determine who is ahead. *J Neurosci*. 2011;31(23):8570–84.
2. Vinck M, Lima B, Womelsdorf T, Oostenveld R, Singer W, Neuenschwander S, et al. Gamma-phase shifting in awake monkey visual cortex. *Journal of Neuroscience*. 2010;30(4):1250–1257.
3. Luczak A, Bartho P, Harris KD. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*. 2009;62(3):413–25.
4. Lu T, Liang L, Wang X. Temporal and rate representations of time-varying signals in the auditory cortex of awake primates. *Nat Neurosci*. 2001;4(11):1131–8.
5. Pastalkova E, Itskov V, Amarasingham A, Buzsaki G. Internally generated cell assembly sequences in the rat hippocampus. *Science*. 2008;321(5894):1322–7.
6. Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, et al. Synfire chains and cortical songs: temporal modules of cortical activity. *Science*. 2004;304(5670):559–564.

7. Wehr M, Laurent G. Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature*. 1996;384(6605):162–6.
8. Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature*. 2017;551(7679):232–236.
9. Chang L, Tsao DY. The code for facial identity in the primate brain. *Cell*. 2017;169(6):1013–1028.
10. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, et al. Neural population dynamics during reaching. *Nature*. 2012;487(7405):51–56.
11. Dragoi G, Buzsaki G. Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron*. 2006;50(1):145–57.
12. Markram H, Lubke J, Frotscher M, Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*. 1997;275(5297):213–5.
13. Dan Y, Poo MM. Spike timing-dependent plasticity of neural circuits. *Neuron*. 2004;44(1):23–30.
14. Abbott LF, Blum KI. Functional significance of long-term potentiation for sequence learning and prediction. *Cereb Cortex*. 1996;6(3):406–16.
15. Grossberger L, Battaglia FP, Vinck M. Unsupervised clustering of temporal patterns in high-dimensional neuronal ensembles using a novel dissimilarity measure. *PLoS computational biology*. 2018;14(7):e1006283.
16. Williams AH, Poole B, Maheswaranathan N, Dhawale AK, Fisher T, Wilson CD, et al. Discovering precise temporal patterns in large-scale neural recordings through robust and interpretable time warping. *Neuron*. 2020;105(2):246–259.
17. Cormen TH, Leiserson CE, Rivest RL, Stein C. *Introduction to algorithms*. MIT press; 2009.
18. Bleich C, Overton ML. *A linear-time algorithm for the weighted median problem*. Courant Institute of Mathematical Sciences, New York University; 1983.
19. Bovo F. *Robustats*; 2020. <https://github.com/FilippoBovo/robustats>.