VU Research Portal

Spike train distances and neuronal coding

Satuvuori, E.

2018

document version
Publisher's PDF, also known as Version of record

Link to publication in VU Research Portal

citation for published version (APA)
Satuvuori, E. (2018). *Spike train distances and neuronal coding.*

9

SUMMARY

A neuron transmits information by a moving potential difference over the cell membrane called action potential and by releasing neurotransmitters to other neurons as a result. The time discrete representation of an action potential is called a *spike*. The series of spikes fired by the same neuron over time are called *spike trains*

A lot of neuroscience research is concentrated on finding response patterns in spike trains of neurons. For a given stimulus a neuron that codes for something in the stimulus is expected to elicit similar response each time. This is called *reliability* of response. However, defining when a response pattern is reliable or not is not a simple task. One may argue that similarity itself is not a well-defined concept. In a way this boils down to defining which properties of the response are considered to describe similarity. And, of course, selecting them will affect the results of any comparison. The properties considered for assessing similarity of responses of neurons must come from the way the neurons transmit information or *code*.

Similarity can be self-similarity in response to the same stimulus at different times, or it can be similarity between responses of two different neurons. There are many ways to assess if spike patterns are similar or not that use different assumptions.

Over the years different measures have been developed in order to quantify similarities between two or more spike trains. The two most popular time scale parametric measures, the Victor-Purpura and the van Rossum distance, describe spike train (dis)similarity based on user-defined time scales to which the measures are mainly sensitive. A drawback of these measures is the fixed time scale, since it sets a boundary between rate and time coding for the recording. However, for real data which typically contain many time scales (such as regular spiking and bursts), this is difficult to detect with a measure that is mainly sensitive to only one of them.

The problem of having to choose one time scale has been eliminated in the three time-resolved and time scale independent measures ISI-dis-

tance, SPIKE-distance and SPIKE-synchronization. Since they always adapt to the local firing rate, all three of these measures are time scale free. While they correctly identify the relative firing rate differences, they have no concept of actual time scales and treat all time scales as equally important. As a the consequence, for very small time scales even minor deviations from perfect synchrony lead to very high values of dissimilarity. For real data the smallest time scales are often not very relevant and any dissimilarities there can mostly be disregarded. Thus the measures' focus on the local time scales results in a (spurious) amplification of dissimilarities which compared to the global time scales are rather negligible.

In Chapter 2 I addressed the problem by proposing generalizations to the three measures called adaptive ISI-distance (A-ISI-distance), adaptive SPIKE-distance (A-SPIKE-distance) and adaptive SPIKE-synchronization (A-SPIKE-synchronization). These generalized definitions add a notion of the relative importance of local differences compared to the global time scales. In particular, they start to gradually ignore differences between spike trains for ISIs that are smaller than a minimum relevant time scale. The adaptive generalizations allow to disregard spike time differences that are not relevant on a more global scale.

The methods were developed for use in neuroscience, but in fact they work the same with any point process. In neuroscience typical time scales are in the range of milliseconds or sometimes seconds and any time scales below this will not be considered relevant. In fields such as meteorology the respective time scales could be hours and days or even months and years. The relevant time scale clearly depends on the system under consideration. Setting the minimum relevant time scale for a given dataset might not be a simple task. To address this, I proposed a method to extract a threshold value from the spike trains, that is based on the proportions of the different time scales present in the data. The extended methods are intended to be used exclusively when there are multiple time scales in the data.

In Chapter 3 I explained a complementary method for identifying similarity between spike trains which quantifies the consistency of the leader-follower relationships within a spike train set. The framework consists of two directional measures (*SPIKE-Order* and *Spike Train Order*) that allows to define a value termed *Synfire Indicator* which quantifies the consistency of the leader-follower relationships in a rigorous and automated manner, and also to sort multiple spike trains from leader to follower. The SPIKE-Order profile was used for color-coding and visualizing local spike leaders and followers and Spike Train Order functioned as an overall indicator of leader-follower consistency. A set of spike trains ex-

hibiting perfectly consistent repetitions of the same global propagation pattern is called a *synfire pattern*. The synfire patterns were evaluated using spike-to-spike matching and expecting that the system where the spikes were recorded from functions with robust single spike timing.

In Chapters 2 and 3 I explained new methods for point processes. In Chapter 4 I took a step back and had a look on how the similarity descriptions in some of the most commonly used spike train distances actually match to what is expected of them. The analysis was based on the two main approaches to neuronal coding. Since a neuron does not discriminate where the stimulation it receives actually came from, it seems reasonable to assume that high enough firing rate in the downstream neurons will eventually build up and elicit spikes in the next one. This assumption is called *rate coding*. An alternative approach considers a more compact coding, where each neuron time their spikes carefully together in order to achieve the wanted spike upstream at a certain time. This referred to as *time coding*. Both the assumptions are perfectly reasonable, yet exclusive from the neuron's point of view. As quoted in Chapter 4:

- *"It is generally accepted that a rate encoding scheme is one in which the relevant information encoded about the stimulus is correlated only with the number of elicited spikes within the encoding window and is not correlated with any aspect of the temporal pattern of the spikes within the encoding window."*

- *"In a temporal encoding scheme, the relevant information is correlated with the timing of the spikes within the encoding window, over and above any information that might be correlated with the number of spikes within the window."*

This definition of temporal coding does not require rate correlation actually to be present. Also it is not limited to single spike correlations but applies to any correlations in spike patterns that would not be expected due to rate alone.

In Chapter 4 I used that definition of time coding as correlations beyond rate to investigate how the sensitivity of the different spike train distances to rate and time coding depends on the rate of the spike trains. In this study I asked two questions: How does the sensitivity of the different spike train distances to rate and time coding depend on the rates of the two processes and how high a rate is needed in order to obtain reliable estimates of timings in the data? The analysis was conducted using independent steady rate Poisson spike trains as surrogates for random spike trains with fixed rate and no timing information.

The first finding is that the spike-resolved Victor-Purpura distance compares the spike trains spike-for-spike and thus they are always sensitive to differences in spike counts even for parameter values seemingly indicating time coding. For large spike count differences the spike-resolved distances do not obtain the ability to assess timing information beyond spike pairs and thus in many cases most of the distance comes from mismatch in spike counts rather than timings, independently of the time scale parameter. As a result, for the Victor-Purpura distance timing information is only available for spike trains with almost identical rates. Since the behaviour of the van Rossum distance in response to rate differences closely resembles that of the Victor-Purpura distance, it also has the same problem (in addition to its normalization issues for different tau-values).

The second finding is that the time-resolved measures perform better in assessing timings in the normal case of reasonably high rates. These measures can also provide a meaningful instantaneous similarity profile within the coding window. Since they assess similarity in time, the exact spike count becomes less important and the actual timing of events becomes more relevant. However, they suffer from artefacts when the rates of the spike generation processes is so low that the floor effect takes place. For spike trains with only a few spikes one should use the spike-resolved Victor-Purpura or van Rossum distance, since they assess first similarity in spike count and then apply timing information assessment only for pairs of spikes.

The nervous system is believed to employ large populations of neurons to code and broadcast information. Population coding can be considered less vulnerable and, hence, a more reliable and robust manner than coding via single neurons. In Chapter 5 I evaluated existing approaches and introduced new algorithms for identifying the most discriminative subpopulations from a population of recorded neurons.

In neuronal recordings population coding can appear in two ways. First, all the neurons in the recorded population contribute equally. Patterns of activity within the population are irrelevant for coding as all that matters is whether or not any of the neurons fires. There, the information being conveyed is that of a single spike train generated by the population as a whole. In contrast to this so-called *summed population (SP) hypothesis*, each neuron may have a unique and distinguishable role. In this case, the population is best decoded neuron-by-neuron, which is referred to as the *labeled line (LL) hypothesis*. The coding via individual neurons and the summation of an entire population are the extreme case in a broad spectrum of possibilities. In fact, recent evidence points at some interme-

diate scenario in which a comparably small number encodes information not only in a robust but also very efficient way.

The search for an optimal coding population requires fine-tuned analyses under both the SP- and the LL-hypothesis. For these two cases I showed how to separate relevant from irrelevant subpopulations by identifying the subpopulation of neurons amongst all possible ones that discriminates best a given set of stimuli.

For the SP case, I compared three fundamentally different algorithms for finding the subpopulation that is able to most efficiently discriminate between a set of stimuli. First SP algorithm is a simple brute force solution that goes through every possible permutation. This is computationally heavy and even for relatively small data sets is not feasible. Computationally lighter gradient algorithms were also tested and I ran two simulations that are constructed such that each time one of the two variants of the gradient algorithm did not find the best subpopulation since it got trapped in a local maximum. Since gradient algorithms are much faster than the brute force approach and successful under idealized conditions, they can be used for first testing. However, our examples illustrate that they can generally not be relied upon in more realistic settings. I also constructed a third algorithm using simulated annealing that provides a recovery mechanism that considerably reduces the likelihood of getting stuck in a local maximum and providing correct result more consistently even if slightly increasing computation load.

For the LL case, I introduced a novel algorithm for identifying the most discriminative LL population by evaluating every neuron separately. First, for every individual stimulus pair the algorithm identifies the discriminative neurons and selects the best one. These best neurons are then combined to form the optimal LL-population. The algorithm can handle quite involved coding scenarios, even though its computational complexity is much lower than in the SP case. Moreover, one is guaranteed to find the best subpopulation since this time no search in a very high-dimensional subpopulation space is needed. I also identified that there is a mismatch between definitions of discrimination and coding, even if the two concepts seem to be linked. In reality one may have coding without discrimination and vice versa.

In this thesis I have introduced new methods and evaluated existing ones. The most important thing to consider when applying methods is to understand what they are designed for. What we think they should do is sometimes not what they are actually doing. Having a clear understanding of how the measures we use evaluate the data is crucial for proper interpretation of the results. My thesis was devoted to this challenge. My comparative studies and newly proposed methods provide

an encompassing approach to the analysis of spike trains and general point processes, and with that, hopefully paved the way to application in neuroscience and beyond.