
Spotting Neural Spike Patterns via Decomposition Into Types

Itay Gat **Naftali Tishby***
Institute of Computer Science and
Center for Neural Computation
The Hebrew University
Jerusalem, 91904 Israel
email: {itay,tishby}@cs.huji.ac.il

Abstract

Detecting a specific stochastic pattern embedded in an unknown background noise is a difficult pattern recognition problem. A similar problem appears when trying to detect a multi-neural spike pattern in a single electrical recording, embedded in the complex cortical activity of a behaving animal. The technical difficulty of this detection is due to the lack of a good statistical model for the background activity, which is rapidly changing with the recording conditions and activity of the animal. By employing a “type-based” information theoretic approach we introduce a new method for detecting patterns of non-homogeneous Poisson processes. The main idea is a decomposition of the log-likelihood distribution according to “types” of the empirical counts. This amounts to the assumption that the background activity has equal probability to contain all permutations of the required pattern. We demonstrate the application of this method for detection the “reward” patterns in the Basal-Ganglia of behaving monkeys, yielding some unexpected biological results.

1 Introduction

The association between stimuli and neural response is one of the fundamental questions regarding the neural code employed by the brain. In order to find this association one can either average many responses to the *same* stimulus and check for significance patterns emerging, or try to spot patterns in the neuronal data directly. The more established way of searching for such association is the first method, and this fact is mainly due to the noisy and sparse nature of the spike trains holding the neuronal data. In non primary areas of the brain the rate of these cells is on the order of few tens of spikes per second on one hand, and the response can be varying quite dramatically in a short period on the other hand. A manifestation for this difficulty can be seen in the following example. Suppose that a neuron is responding to a given stimulus, this response may be well contained within a period of 600 ms, where the average firing rate of the cell is 10 spikes-per-second. The outcome of this combination is that the response is manifested within 6 spikes on the average. If the response is made out of changes in the firing rate of the neuron - it may well be that a change in a single spike may signal the transition from high firing rate to low firing rate.

Although that the average method overcome the sparse nature of the spike trains, and therefore enabled us to extend our knowledge about the association described above, it suffers from several inherent limitations. It assumes that every time the animal is being presented with the same stimulus the same computation procedure is taking place. However, we are unable to produce *exactly* the same conditions for the animal trial after trial, therefore the averaging is done over *different* responses. Another weak point of this method is the blindness to neural responses which are not tightly locked to a stimulus. If a response is presented in a random manner in a specified segment of time, the averaging method would not be able to observe that. The last severe limitation lies in the fact that we are searching just in the vicinity of stimuli and not in all the continuous data. One may find for example that a response is also associated with a stimulus that was not observed by the researcher but still influenced the animal quite a lot. This can be observed only if the response was to be searched in all of the data.

In this work we try to take the road less traveled by of spotting the neuronal template on the continuous unmarked data. Only after doing that we would set out attention to the location of the spotting, and how are they related to the

external stimulus presented to the animal. To our aid come the fact that modern recoding allows us to record several cells simultaneously - so instead of averaging over time we may do that over many cells.

2 Statistical description of the data

The framework used for this work is as follows. We have a continuous unmarked stream of neuronal data. This data consists of the simultaneous spike trains, and we are searching for a specific response in this data. We take the assumption that the spikes within a short time (a bin) are obeying a Poisson distribution. A pattern to be searched in such data is therefore parameterized by a sequence of M firing rates $(\lambda_1, \lambda_2, \dots, \lambda_M) = \underline{\lambda}$, which for a non-homogeneous Poisson process[2].

The search itself was carried out by shifting a window whose length is the same as the template over the data stream. Each such shift would define a frame of data, and the spikes in this frame would be counted in bins whose width is equal to the bins of the template. The frame of data is therefore transformed into a vector of size M whose element are the counts $(n_1, n_2, \dots, n_M) = \underline{n}$. See the example of figure 1.

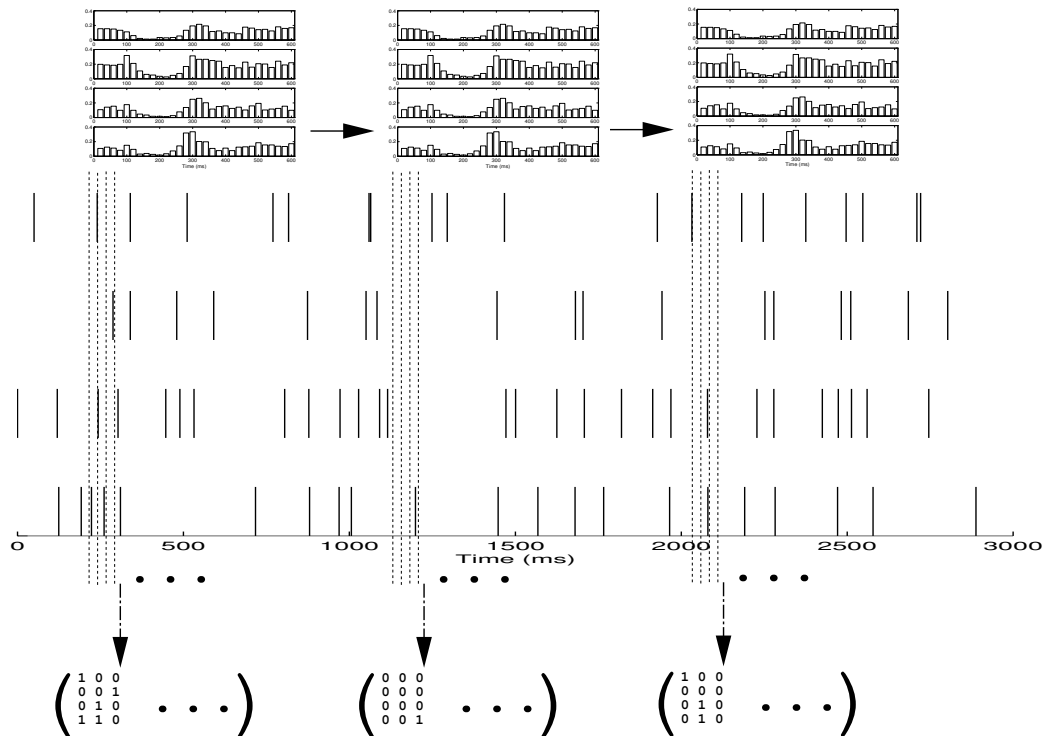


Figure 1: An example of a pattern to be searched for on continuous data. The pattern is being slid on the spike trains and the spike trains are binned according to the same bin size of the pattern. Each frame is thereafter associated with a probability to be the outcome of the pattern. Pay attention to the low firing rate and to its high variability.

Using the two vectors $\underline{\lambda}$ and \underline{n} the probability of the frame to be the outcome of the template would be computed. A specific pre-defined percentage of the frames with the highest probabilities would be considered as the spotting of the template over the data - *spotted frames*.

The basic statistical relationship between the data points and the pattern is the likelihood, which is the statistical value that define how likely is a pattern given a data point. Given a variable rate Poisson pattern with M bins $\underline{\lambda}$ and observed sample sequence of counts \underline{n}

$$p_{\underline{\lambda}}(\underline{n}) = \prod_{i=1}^M e^{-\lambda_i} \frac{\lambda_i^{n_i}}{n_i!}. \quad (1)$$

The decision theoretic problem of spotting a pattern of sequenced data in a “noisy background” is classically considered as a binary statistical hypothesis test, where the two possible sources are a statistical model of the pattern and a background statistical model. Given a complete description of the two sources there is a fully understood statistical approach, the Likelihood-Ratio-Test (LRT), first derived by Neyman and Pearson (see e.g. [4]). Unfortunately as can be seen clearly for figure 1 a background model is not known. Our main contribution in this work is to supply a reasonable background model, for such problems.

3 Decomposition of the log-likelihood distribution into sample “types”

Our background model assumption is as follows: we assume that we are spotting the template $\underline{\lambda}$ against “adversary background” which has the same likelihood to contain all the permutations of the required pattern. Detection against this adversary background is precisely what we obtain by decomposing the likelihood distribution according to the *sample type class*. The idea of types stem from information theory (see e.g. [1] Ch. 12).

The type class correspond to the set of all samples with the same empirical distribution. For this case it is clear that the probability of the type class is given by the product of the likelihood of a single sample in this class (all have the same probability) by the cardinality of the class, i.e. the type multinomial factor.

Since our observations are the bin spike counts, the sample type class is defined as the class of all spike patterns that have the same empirical distribution of bin-counts, i.e., c_0 empty bins, c_1 bins with one spike, c_2 with 2 spikes, etc. The type class is completely characterized by the vector of numbers c_0, c_1, c_2, \dots .

Denoting by $T(\underline{n})$ the type class of the count vector \underline{n} and by $c_i(\underline{n})$ the number of bins with i spike counts in \underline{n} . Clearly, for patterns with $M = \sum_i c_i$ bins the size of $T(\underline{n})$ is given by:

$$|T(\underline{n})| = \frac{M!}{\prod_{i=1}^M c_i(\underline{n})!}. \quad (2)$$

Since sample type classes are disjoint sets the sample log-likelihood distribution is uniquely decomposed into sum of the log-likelihood distributions for each type. That is,

$$P_{\underline{\lambda}}(\log p_{\underline{\lambda}}(\underline{n})) = \sum_{T(\underline{n})} P_{\underline{\lambda}}(T(\underline{n})) \times P_{\underline{\lambda}}(\log p_{\underline{\lambda}}(\underline{n})|T(\underline{n})). \quad (3)$$

Since given a sample its type is known, the idea is to use only one term in this sum for the detection, rather than all the sum. In other words, the sample type serves as a sufficient statistic, together with the sample likelihood, for each term in this mixture distribution. It remains to evaluate the different distributions in this decomposition.

3.1 The type class probability

The probability of the type class $T(\underline{n})$ under the model $\underline{\lambda}$ can be determined through

$$P_{\underline{\lambda}}(T(\underline{n})) = \sum_{n' \in T(\underline{n})} p_{\underline{\lambda}}(\underline{n}') = |T(\underline{n})| \langle p_{\underline{\lambda}}(\underline{n}) \rangle_{T(\underline{n})} \approx |T(\underline{n})| e^{(\log p_{\underline{\lambda}}(\underline{n}))_{T(\underline{n})}}, \quad (4)$$

where $\langle \dots \rangle_{T(\underline{n})} = \frac{1}{|T(\underline{n})|} \sum_{n' \in T(\underline{n})} \dots$ denotes an average over all the members of the type class $T(\underline{n})$, i.e., an average over all the $|T(\underline{n})|$ permutations of bins.

The last exponential approximation in Eq.(4) is based on the fact that $p_{\underline{\lambda}}(\underline{n})$ is a product over many positive statistically independent terms (each with a well behaving distribution) and thus distributed by a *log – normal* distribution for large enough M (its log “self-average”).

The type probabilities, $P_{\underline{\lambda}}(T(\underline{n}))$, should be normalized to 1, so each term can be considered as the type prior.

Next we evaluate the parameters of the normal distribution of $\log p_{\underline{\lambda}}(\underline{n})$ in each type, $P_{\underline{\lambda}}(\log p_{\underline{\lambda}}(\underline{n}))$. As can be seen from the histogram of the log-likelihood in the type, figure 2, the central limit theorem even better applies to this distribution, which can is thus considered as normal.

$$\mu_{T(\underline{n})} = \langle \log p_{\underline{\lambda}}(\underline{n}) \rangle_{T(\underline{n})} = -\Lambda - L(\underline{n}) + \left\langle \sum_{i=1}^M n_{\pi(i)} \log(\lambda_i) \right\rangle_{\pi \in T(\underline{n})} \quad \sum_{i=1}^M \lambda_i = \Lambda. \quad (5)$$

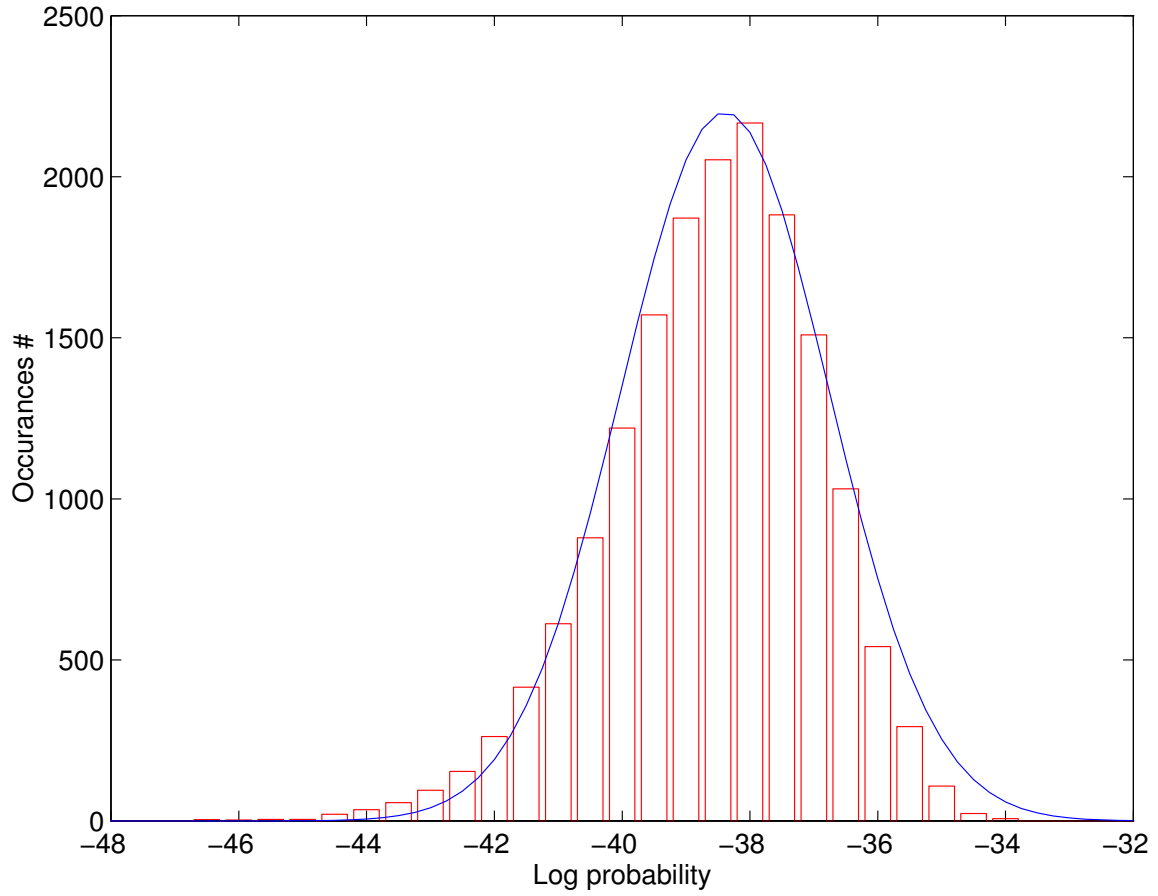


Figure 2: Histogram of the model log-likelihood in one type with the Gaussian approximation. The histogram was computed for one type out of a given $\underline{\lambda}$ template. The histogram was computed from simulated data, and the solid line shows the normal distribution with mean and variance which were analytically computed.

$$\sigma_{T(\underline{n})}^2 = \frac{1}{M-1} \left(\sum_i n_i^2 - \frac{N^2}{M} \right) \left(\sum_i \log^2(\lambda_i) - \frac{l(\underline{\lambda})^2}{M} \right) \quad \sum_{i=1}^M n_i = N, \quad (6)$$

or, through the variances of the counts and log rates, as

$$\sigma_{T(\underline{n})}^2 = \frac{M^2}{M-1} \text{Var}(\{n_i\}) \text{Var}(\{\log \lambda_i\}). \quad (7)$$

The last equation clearly shows that when either the $\{n_i\}$ or the $\{\lambda_i\}$ are uniform the variance of type log-likelihood vanishes and the log-likelihood is sharply centered on the type mean.

Using the above expressions for the mean and variance for the type log-likelihood distribution we obtain the type decomposition of $P_{\underline{\lambda}}$, Eq.(3), as:

$$P(\log p_{\underline{\lambda}}(\underline{n})) = \sum_{T(\underline{n})} P_{\lambda}(T(\underline{n})) P(\log p_{\underline{\lambda}}(\underline{n}) | T(\underline{n})) \quad (8)$$

$$= \sum_{T(\underline{n})} e^{(\log |T(\underline{n})| + \mu_{T(\underline{n})})} e^{\left(-\frac{(\log p(\underline{n}) - \mu_{T(\underline{n})})^2}{2\sigma_{T(\underline{n})}^2} - \frac{1}{2} \log(2\pi\sigma_{T(\underline{n})}^2) \right)} \quad (9)$$

The last term is just the log normal probability we have just computed the mean and the variance of, and the first part is:

$$|T(\underline{n})| \exp(\langle \log(p(\underline{n})) \rangle_{T(\underline{n})}),$$

or alternatively:

$$\exp(\log(|T(\underline{n})|) + \langle \log(p(\underline{n})) \rangle_{T(\underline{n})}).$$

For a given spike sample both the type and the model log-likelihood are known and our final probabilistic score would be:

$$S_{\lambda}(\underline{n}) = \log(|T(\underline{n})|) + \mu_{T(\underline{n})} - \frac{(\log p(\underline{n}) - \mu_{T(\underline{n})})^2}{2\sigma_{T(\underline{n})}^2} - \frac{1}{2} \log(2\pi\sigma_{T(\underline{n})}^2),$$

Where both $\mu_{T(\underline{n})}$ and $\sigma_{T(\underline{n})}^2$ depends both on the type $T(\underline{n})$, and the model λ .

4 Application

The application of the new detection method was shown on the neural response of the tonically active neurons (TANs) in the striatum which is the entry station of the Basal-Ganglia[8]. This type of cells shows a strong and robust response to cues predicting future rewards[5]. This response would henceforth be referred to as the “reward signature” of the TANs. The reward signature is characterized by a reduction in firing rate (pause), often flanked by brief elevation of the firing rate[8]. An example of such signature is presented in figure 3. It was shown that the reward signature was found in TANs recorded in different areas of the striatum, and as an outcome of different reward stimuli.




Figure 3: A reward signature of TANs. The peri-stimulus-time histogram shows an averaged response of cells to a reward. The averaging was performed using recording from one animal (monkey H), consisted of 23 different TANs taken from 10 different recording sessions. The bin size used in the computation of the histogram was 20 ms.

The data used for this work were the spike trains recorded from two behaving monkeys (referred to as H and I) in the higher brain function laboratory of Haddassah. The monkeys were trained to perform a delayed response task and upon success completion of the task were rewarded with a drop of juice. Before the animal was actually being rewarded it was given a cue for the future reward. This cue is referred to as the reward cue. While the animals were performing the task the simultaneous activity of up to 4 TANs was recorded. Using these recordings we estimated the reward signature by averaging the response of 23 TANs taken from monkey H. This estimation is presented in figure 3 and would be used henceforth as the pattern to be searched upon the continuous recording. A more detailed description of the results of this data analysis and its meaning is given elsewhere[6].

We used the newly developed spotting mechanism to perform this spotting task. The data used were 5 recording sessions which held at least 3 cells recorded simultaneously. From the five sessions used 3 were recorded from monkey H, and two were taken from monkey I. The outcome of the spotting mechanism supplied us with a list of segment which had the highest probability to be the outcome of a reward signature. These segments would be henceforth be referred to as the “reward signal” of our data.

4.1 The prototypical template is spotted significantly

We spotted the prototypical template in the continuous data and in order to examine the significance of this spotting we compared the number of reward signal with the same number found in surrogate data. The surrogate data consisted of the same spike trains in which the different cells were randomly shifted in time one against to the other. As can be seen from table 1 second column the spotting was highly significance for all recording sessions. In order to further validate this procedure we used 3 other null hypothesis templates. As can be seen from the third fourth and fifth columns in table 1 none of these templates produced any significance detection compared to the surrogate data. Pay attention to the fact that although the prototypical template was computed using the reward signature of animal H, it was highly effective also for the other animal (monkey I).

Table 1: The significance of spotting the neural reward signal. We used several templates to verify the significance of spotting: the prototypical template, a flat template (all bins are equal), a template computed from segments aligned 2 seconds before the reward cue, and a shuffled prototypical template. This significance level was the difference between the number of detections in the original data and the average number of detections in the surrogate data. This difference was measured in standard deviation (SD) of the shuffled data.

Session	Reward signature	Flat PSTH	Before reward	Random shuffle
H-1	15.3	-0.25	-0.91	-1.1
H-2	9.4	0.44	0.43	-0.59
H-3	9.9	-1.45	2.10	2.1
I-1	18.2	-0.88	1.70	-1.58
I-2	11.8	0.96	-0.53	0.8

4.2 The prototypical template association with reward

Based on averaging many reward cue epochs it was previously hypothesized that the reward signature is a binary signal which can be found only confined to the reward cue. We tested this hypothesis directly and found that the template is spotted in all periods of the trial. Naturally, the frequency of spotted events is maximal in the epoch locked to the reward cue. However, the template is significantly spotted in all periods of the trial and only 11% (range over recording sessions equal 6-23%) of all detections are time locked to the reward cue.

The frequency of the spontaneously emitted reward signal was not constant over all behavioral periods. As expected, this frequency reach a peak at the reward cue period, but remains high in the 2-4 seconds following the reward cue. An interesting tendency toward reduction in this frequency is observed following erroneous trials (trials in which the animal did not perform the needed task and therefore was not rewarded). It was seen that in *all* 5 recording sessions the amount of spotted frames after erroneous trials was lower than average and the number of spotted frames 2 seconds after the reward cue was higher than average. These results are summarized in table 2.

5 Discussion

There are several possible extensions to the work presented here. The most strait-forward one is to extend the types used here into ones containing information on higher correlation of the data. In our application we work against an adversary environment that consists of all the shuffles of the given pattern in single bins. Second order correlation would be to have an adversary environment that keeps also the order within pairs of bins. In order to perform such spotting we need to recompute the mean and variance inside each such new type. In general we can work with each order of correlation, and thus would be able to cope with even harder spotting cases. When we extend the correlation order the background is getting more and more similar to the original pattern under investigation.

In general we believe that the idea of decomposing the likelihood into sub-parts could be found useful. This would

Table 2: The deviation from average level in correct trials and erroneous ones. The change from average level of reward signal is given both for the segment of 1 second starting 2 seconds after correct trials, and the same segment in erroneous ones. The change is given in percentage of change from average level.

Session	Change in correct trials	Change in erroneous trials
H-1	61.7	-100
H-2	22.3	-60.6
H-3	28.7	-57.4
I-1	44.7	-12.7
I-2	12.5	-54.2

allow us to work in environments in which we can not assume the usual flat distribution of the background and without building a specific model for it we can perform reliable spotting. In our case the decomposition into types seems sufficient and produced the required results.

We have also examined a related solution to the spotting problem based on the method of types [3]. Direct application of this technique for detecting multi-neural spike patterns can be done by treating the empirical type of every bin in a template model (PST histogram) of the training data [7]. In this case one simply estimate the probability that the training data and the test sample are taken from the same source for every bin or combination of bins in the template model. It is shown elsewhere [9] that for our neuronal data this method suffer from insufficient training data and therefore performed rather poorly on our data, although in principal this method can overcome the difficulties related to this kind of pattern spotting.

Acknowledgments

Enlightening discussions with William Bialek, Moshe Abeles, and Shai Fine are greatly appreciated. We thank Hagai Bergman and Aeyal Raz for providing us with the TAN-cells recordings. This work was partially supported by a grant from the Israeli Science Foundation

References

- [1] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, Mew York, 1991.
- [2] M. Abeles, H. Bergman, I. Gat, I. Meilijson, E. Seidman, N. Tishby, E. Vaadia. "Cortical activity flips among quasi-stationary states" *Proc. Natl. Acad. Sci. USA* 92:8616–8620, 1995.
- [3] M. Gutman. "Asymptotically optimal classification for multiple tests with empirically observed statistics". *IEEE Trans. Info. Th.* 35:401–408, 1989.
- [4] F.L. Leheman. *Testing Statistical Hypotheses*. Wiley, New York, 1959.
- [5] M. Kimura, "The role of primate putamen neurons in the association of sensory stimuli with movement." *Neurosci.Res.* 3:436-443, 1986.
- [6] I. Gat, A. Raz, M. Abeles, N. Tishby, H. Bergman, "The Neural Signature of Reward is Emitted Continuously in the Striatum of Normal Primates and is Absent in the Parkinsonian State" *Preprint*, 1999.
- [7] D. H. Johnson, C.M. Gruner, K. Baggerly, C. Seshagiri, "Information -Theoretic Analysis of the Neural Code", *Preprint*, 1998.
- [8] A. M. Graybiel, T. Aosaki, A. W. Flaherty, M. Kimura, *Science* 265, 1826 (1994).
- [9] I. Gat, N. Tishby "The adversary background model in spotting neural spike patterns". *in preparation*.