



Article Estimating Mutual Information for Spike Trains: A Bird Song Example

Jake Witter * D and Conor Houghton

Faculty of Engineering, University of Bristol, Bristol BS8 1TR, UK; conor.houghton@bristol.ac.uk

* Correspondence: jake.witter@bristol.ac.uk

Abstract: Zebra finches are a model animal used in the study of audition. They are adept at recognizing zebra finch songs, and the neural pathway involved in song recognition is well studied. Here, this example is used to illustrate the estimation of mutual information between stimuli and responses using a Kozachenko–Leonenko estimator. The challenge in calculating mutual information for spike trains is that there are no obvious coordinates for the data. The Kozachenko–Leonenko estimator does not require coordinates; it relies only on the distance between data points. In the case of bird songs, estimating the mutual information demonstrates that the information content of spiking does not diminish as the song progresses.

Keywords: mutual information; spike trains; zebra finch; spike train metric

1. Introduction

The mutual information between two random variables *X* and *Y* is often conveniently described using a diagram like this:



where the whole rectangle represents the entropy H(X, Y) of the joint variable (X, Y). This is, in general, less than the sum of H(X) and H(Y) because X and Y are not independent. In this diagram, the purple and green regions together are intended to represent H(X), and the green and yellow regions are intended to represent H(Y). The purple region on its own represents H(X|Y): the entropy remaining, on average, when the value of Y is known. In the same way, the yellow region represents H(Y|X). Now, the mutual information is represented by the green section:



$$I(X,Y) = H(X) - H(X|Y) = H(Y) - H(Y|X)$$
(1)



Citation: Witter, J.; Houghton, C. Estimating Mutual Information for Spike Trains: A Bird Song Example. *Entropy* **2023**, *25*, 1413. https:// doi.org/10.3390/e25101413

Academic Editors: Arno Onken and Nina Kudryashova

Received: 4 May 2023 Revised: 31 August 2023 Accepted: 19 September 2023 Published: 3 October 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/).

It is

or, by substitution,

$$I(X,Y) = \mathbb{E}\log_2\left[\frac{p_{X|Y}(x|y)}{p_X(x)}\right] = \mathbb{E}\log_2\left[\frac{p_{Y|X}(y|x)}{p_Y(y)}\right]$$
(2)

Here, for illustrative purposes, mutual information is described relative to a specific example: the neural response of cells in the zebra finch auditory pathway to zebra finch songs. This is both an interesting neuroscientific example and an example which is typical of a broad set of neuroscience problems.

The zebra finch is a model animal used to study both auditory processing and learning; the male finch sings, and he has a single song which begins with a series of introductory notes, followed by two or three repetitions of the motif: a series of complex frequency stacks known as syllables, separated by pauses. Syllables are about 50 ms long, with songs lasting about two seconds. The songs have a very rich structure, and both male and female zebra finches can distinguish one zebra finch song from another.

Here, we use a data set consisting of spike trains recorded while the bird is listening to a set of songs, and we provide an estimate for the mutual information between the song identity and spike trains recorded from cells in the auditory pathway. This is an interesting and non-trivial problem. Generally, calculating mutual information is costly in terms of data because it requires the estimation of probabilities such as $p_Y(y)$ and $p_{Y|X}(y|x)$. For this reason, some measure of correlation is often used when quantifying the relationship between two random variables. However, not all data types have a correlation: calculating the correlation assumes algebraic properties of the data that are not universal. As an example, calculating the correlation between X and Y requires the calculation of $\mathbb{E}[XY]$, which in turn assumes that it makes sense to multiply the *x* and *y* values. This is not the case for the typical neuroscience example considered here, where the set of outcomes for X is the song identities and that for Y is the spike trains. To circumvent this, spike trains are often replaced with something else, spike counts for example. However, this involves an implicit assumption about how information is coded. This is likely to be inappropriate in many cases. Indeed, the approach taken to calculating mutual information can involve making very strong assumptions about information coding, the very thing that is being studied.

The purpose of this review paper is to demonstrate a different approach: there is a metric-space version of the Kozachenko–Leonenko estimator [1,2] introduced in [3–5] and inspired by [6]. This approach has been tested on simulated data, for example in [5], and this shows it to be promising. However, it is important to also test it on real data. Here, it is applied in the zebra finch example.

2. Materials and Methods

Let

$$\mathcal{D} = \{ (x_1, y_1), (x_2, y_2), \dots, (x_n, y_n) \}$$
(3)

be a data set. In our case, the x_i are the labels for songs in the set of stimuli, with each $x_i \in \{1, ..., n_s\}$; n_s is the number of different songs. For a given trial, y_i is the spiking response. This will be a point in "the space of spike trains". What exactly is meant by the space of spike trains is less clear, but for our purposes here, the important point is that this can be regarded as a metric space, with a metric that gives a distance between any two spike trains; see [7,8], or, for a review, [9].

Given the data, the mutual information is estimated by

$$I(X,Y) \approx \frac{1}{n} \sum_{i=1}^{n} \log_2 \left[\frac{p_{Y|X}(y_i|x_i)}{p_Y(y_i)} \right]$$

$$\tag{4}$$

where the particular choice of which conditional probability to use, $p_{Y|X}$ rather than $p_{X|Y}$, has been made for later convenience. Thus, the problem of estimating mutual information is one of estimating the probability mass functions $p_{Y|X}$ and p_Y at the data points in \mathcal{D} . In

our example, there is no challenge to estimating p_X , since each song is presented an equal number of times during the experiment $p_X(x_i) = 1/n_s$ for all x_i and, in general $p_X(x_i)$ is known from the experiment design. However, estimating $p_{Y|X}$ and p_Y is more difficult.

In a Kozachenko–Leonenko approach, this is performed by first noting that for a small volume R_i containing the point y_i ,

$$p_Y(y_i) \approx \frac{1}{\operatorname{vol}(R_i)} \int_{R_i} p_Y(y) \, dy \tag{5}$$

with the estimate becoming more and more exact for smaller regions R_i . If the volume of R_i were reduced towards zero, $p_Y(y)$ would be constant in the resulting tiny region. Here, $vol(R_i)$ denotes the volume of R_i . Now, the integral $\int_{R_i} p_Y(y) dy$ is just the probability mass contained in R_i and so it is approximated by the number of points in \mathcal{D} that are in R_i :

$$\int_{R_i} p_Y(y) \, dy \approx \frac{|\{y_j \in R_i\}|}{n}.$$
(6)

It should be noted at this point that this approximation becomes more and more exact as R_i becomes bigger. Using the notation

$$k_i = |\{y_j \in R_i\}| \tag{7}$$

this means

$$p_Y(y_i) \approx \frac{k_i}{n \operatorname{vol}(R_i)}.$$
 (8)

This formula provides an estimate for $p_Y(y_i)$ provided that a strategy is given for choosing the small regions R_i around each point y_i . As will be seen, a similar formula can be derived for $p_{Y|X}(y_i|x_i)$, essentially by restricting the points to $\mathcal{D}_i = \{(x_i, y_i) \in \mathcal{D} | x_i = x_i\}$:

$$p_{Y|X}(y_i|x_i) \approx \frac{h_i}{n_c \operatorname{vol}(R_i)}$$
(9)

where h_i is the number of points in R_i with label x_i and n_c is the total number of points with label x_i . In the example here, $n_c = n/n_s$. Once the probability mass functions are estimated, it is easy to estimate the mutual information. However, there is a problem: the estimates also require the volume of R_i . In general, a metric space does not have a volume measure. Furthermore while many everyday metric spaces also have coordinates providing a volume measure, this measure it not always appropriate since the coordinates are not related to the way the data are distributed. However, the space that the y_i 's belong to is not simply a metric space, it is also a space with a probability density, $p_Y(y)$. This provides a measure of volume:

$$\operatorname{vol}(R_i) = \int_{R_i} p_Y(y) dy \tag{10}$$

In short, the volume of a region can be measured as the amount of probability mass it contains. This is useful because this quantity can in turn be estimated from data, as before, by counting points:

$$\operatorname{vol}(R_i) \approx \frac{k_i}{n}.$$
 (11)

The problem with this, though, is that it gives a trivial estimate of the probability. Substituting back into the estimate for $p_Y(y_i)$, Equation (8) gives $p_Y(y_i) = 1$ for all points y_i . This is not as surprising as it might at first seem. Probability density is a volume-measure dependent quantity; that is what is meant by calling it a density and is the reason that entropy is not well defined on continuous spaces. There is always a choice of coordinate that trivializes the density.

However, it is not the entropy that is being estimated here. It is the mutual information and this is well defined: its value does not change when the volume measure is changed. The mutual information uses more than one of the probability densities on the space; in addition to $p_Y(y_i)$, it involves the conditional probabilities $p_{Y|X}(y|x)$. Using the measure defined by $p_Y(y)$ does not make these conditional probability densities trivial. The idea behind the metric space estimator is to use $p_Y(y)$ to estimate volumes. This trivializes the estimates for $p_Y(y_i)$, but it does allow us to estimate $p_{Y|X}(y|x)$ and use this to calculate an estimate of the mutual information.

In this way, the volume of R_i is estimated from the probability that a data point is in R_i , and this, in turn, is estimated by counting points. Thus, to fix the volume $vol(R_i)$, a number h of data points is specified, and for each point, the h - 1 nearest data points are identified, giving h points in all when the "seed point" is included. This is equivalent to expanding a ball around y_i until it has an estimated volume of h/n. This defines the small region R_i . The conditional probability is then estimated by counting how many points in R_i are points with label x_i , that is, are points in \mathcal{D}_i . In fact, this just means counting how many of the h points that have been identified are in \mathcal{D}_i , or, put another way, it means counting how many of the h - 1 nearest points to the original seed point are from the same stimulus as the seed point. In summary, the small region consists of h points. To estimate $p_{Y|X}(y_i|x_i)$, the number of points in the small region corresponding to label x_i is counted; this is referred to as h_i so

$$h_i = |\{y_i \in R_i \mid x_i = x_i\}| = |R_i \cap \mathcal{D}_i|.$$
(12)

This is substituted into the formula for the density estimator, Equation (6), to obtain

$$p_{Y|X}(y_i|x_i) \approx \frac{n}{n_c} \frac{h_i}{h}$$
(13)

where, as before, n_c is the total number of trials for each song. It is assumed that each song is presented the same number of times. It would be easy to change this to allow for different numbers of trials for each song, but this assumption is maintained here for notational convenience. Substituting back into the formula for the estimated mutual information, Equation (4) gives

$$I_0 = \frac{1}{n} \sum_{i=1}^n \log_2 \frac{n_s h_i}{h}$$
(14)

The calculation of I_0 is illustrated in Figure 1. The subscript zero has been added in order to preserve the unadorned I for the information itself and \tilde{I} for the de-biased version of the estimator; this is discussed below.

This estimate is biased, and it gives a non-zero value even if *X* and *Y* are independent. This is a common problem with estimators of mutual information. One advantage of the Kozachenko–Leonenko estimator described here is that the bias at zero mutual information can be calculated exactly. Basically, for the estimator to give a value of zero $h_i = h/n_s$ would be required for every *i*. In fact, while this is the expected value if *X* and *Y* are independent, h_i has a probability distribution which can be calculated as a sort of an urn problem. As detailed in [10], performing this calculation gives the de-biased estimator:

$$I \approx \tilde{I} = I_0 - I_b \tag{15}$$

where I_b , the bias, is

$$I_b = \sum_{r=1}^h \sum_{c=1}^{n_s} \frac{n_c}{n} u(r-1; n_c-1, h-1, n-n_c) \log_2 \frac{n_c r}{h}$$
(16)

and *u* is the probability for the hypergeometric distribution using the parameterization used by the distributions.jl Julia library.

$$u(k;s,h,f) = {\binom{s}{k}} {\binom{f}{m-k}} / {\binom{s+f}{m}} \equiv \text{Hypergeometric}(s,m,f)$$
(17)

Obviously, the estimator relies on the choice of smoothing parameter h. Recall that for a small h, the counting estimates for the number of points in the small region and for the volume of the small regions are noisy. For a large h, the assumption that the probability density is constant in the small region is poor. These two countervailing points of approximation affect I_0 and I_b differently. It seems that a good strategy in picking h for real data is to maximize $\tilde{I}(h)$ over h. This is the approach that will be adopted here.



Figure 1. The calculation of *I* and the spiking data. (**A**) illustrates how the estimator is calculated. The circles and triangle are data points, and red and blue represent two different labels. The dashed line is the small region around the seed point in the center marked by a triangle \blacktriangle while the other, non-seed points are circles: • and •. Here, h = 7, so the ball has been expanded until it includes seven points. It contains four red points, the colour of the central point, so $h_{\blacktriangle} = 4$. For illustration, the points have been drawn in a two-dimensional space, but this can be any metric space. (**B**) describes the data. The spiking responses of a typical neuron to each presentation of a song is plotted as a raster plot, with a mark for each spike. The trials are grouped by song, so the ten responses in each group correspond to repeated presentations of a single stimulus. Stimulus onset is aligned at 0, with the shortest song lasting 1.65 s.

As an example, we will use a data set recorded from zebra finches and made available on the Collaborative Research in Computational Neuroscience data sharing website [11]. This data set contains a large number of recordings from neurons in different parts of the zebra finch auditory pathway. The original analysis of these data are described in [12,13]. The data set includes different auditory stimuli; here, though only the responses to zebra finch song are considered. There are 20 songs, so $n_s = 20$, and each song is presented ten times, $n_c = 10$, giving n = 200. The zebra finch auditory pathway is complex and certainly does not follow a single track, but for our purposes, it looks like

auditory nerve
$$\rightarrow CN \rightarrow MLd \rightarrow OV \rightarrow Field L \rightarrow HVc$$
 (18)

where CN is the cochlear nuclei; MLd is the mesencephalicus lateralis pars dorsalis, analogous to mammalian inferior colliculus; OV is the nucleus ovoidalis; Field L is the primary auditory pallium, analogous to mammalian A1; and, finally, HVc is regarded as the locus of song recognition. The mapping of the auditory pathway and our current understanding of how to best associate features of this pathway to features of the mammalian brain is derived from, for example [13–17].

In the data set, there are 49 cells from each of MLd and Field L, and here, the entropy is calculated for all 98 of these cells.

3. Results

Our interest in considering the mutual information for bird songs was to check whether the early part of the spike train was more informative about the song identity. It seemed possible that the amount of information later in the spike train would be less than in the earlier portion. This does not seem to be the case.

There are a number of spike train metrics than could be used. Although these differ markedly in the mechanics of how they calculate a distance, it does appear that the more successful among them are equally good at capturing the information content. In Figure 2A, the total mutual information between song identity and spike train is plotted. Here, the Victor–Purpura (VP) metric [7], the spike count, the Earth mover distance (EMD) [18], and the van Rossum metric [8] are considered. The Victor–Purpura metric and van Rossum metric both include a parameter which can be tuned, roughly corresponding to the precision of spike timing. Here, the optimal value for each case has been used, chosen to maximize the average information. These values are q = 32.5 Hz for the VP metric and $\tau = 15$ ms for the vR metric. The mutual information estimator uses the metric to order the points, and each small region contains the h - 1 points nearest the seed point so the estimator does not depend on the distances themselves, just the order. Indeed, the estimated mutual information is not very sensitive to the choice of q or τ . This is demonstrated in Figure 2B, where the mutual information is calculated as a function of q, the parameter for the VP metric.



Figure 2. Information content according to different distances. (**A**) shows mean mutual information (MI) among the 98 neurons from both regions according to different distance metrics, the Victor–Purpura metric, the firing rate, the Earth mover distance, and the van Rossum metric. To calculate the mutual information, 1.65 s of spike train is used, corresponding to the length of the short song. (**B**) shows how that mean MI varies according to the *q* parameter for the Victor–Purpura metric. In both cases, blue corresponds to MLd and red corresponds to Field L. In (**B**), the translucent band corresponds to the middle 20% of data points; there is substantial variability in information across cells.

The Victor–Purpura metric and van Rossum metric clearly have the highest mutual information and are very similar to each other. This indicates that the estimator is not sensitive to the choice of metric, provided the metric is one that can capture features of the spike timing as well as the overall rate. The spike count does a poor job, again indicating that there is information contained in spike timing as well as the firing rate. Similar results were seen in [9,19], though a different approach to evaluating the performance of the metrics was used there.

The cells from MLd have higher mutual information, on average, than the cells from Field L. Since Field L is further removed from the auditory nerve than MLd, this is to be expected from the information processing inequality. This inequality stipulates that away from the source of information, information can only be lost, not created.

In Figure 3, the information content of the spike trains as a function of time is considered. To achieve this, the spike trains are sliced into 100 ms slices and the information is calculated for each slice. The songs have variable lengths, so the mutual information becomes harder to interpret after the end of the shortest song, marked by a dashed line. Nonetheless, it is clear that the rate of information and the information per spike are largely unchanged through the song.



Figure 3. Information content per time. These figures show the time-resolved mutual information by calculating the mutual information for spiking response over 0.1 s slices; the centers of which, *T*, are plotted against the mean mutual information. (**A**) shows how this varies over time, with a vertical line showing the ending of the shortest stimulus. (**B**) shows the mean information per spike; although (**A**) shows a small decrease, (**B**) seems to indicate that this corresponds to a reduction in firing rate, not in the information contained in each spike. In both cases, the metric is the VP metric with q = 30 Hz.

4. Discussion

As well as demonstrating the use of the estimator for mutual information, we were motivated here by an interest in the nature of coding in spike trains in a sensory pathway. It is clear that the neurons in MLd and Field L are not "grandmother" neurons, responding only to a specific song and only through the overall firing rate. The firing rate contains considerably less information than was measured using the spike metrics. The spike metrics, in turn, give very similar values for the mutual information; this appears to indicate that the crucial requirement of a spike train metric is a "fuzzy" sensitivity to spike timing. This demonstrates the need for an estimator such as the Kozachenko Leonenko estimator used here.

Approaches that do not incorporate spike timings underestimate the mutual information, but histogram methods, which do include timings, are computational impractical for modest amounts of data. A pioneering paper, [19], also examines mutual information for zebra finch songs, but using a histogram approach. The substantial conclusion there was similar to the conclusion here: there was evidence that spike timings are important. However, it seems likely that this early paper was constrained in its estimates by the size of the data set. This is suggested by the way the amount of information measured increased monotonically as the bin-width in the temporal discretization was reduced, a signature of a data-constrained estimate.

Finally, it is observed that it is not the case that the precision of spiking diminishes as the song continues. Since that song can often be identified from the first few spikes of the response, it might be expected that the neuronal firing would become less precise. Precision is metabolically costly. However, although the firing rate falls slightly, the information remains constant on a per-spike basis.

Author Contributions: Both authors contributed to conceptualization, methodology, and writing. All authors have read and agreed to the published version of the manuscript.

Funding: J.W. is supported by EPSRC DTP (EP/T517872/1). C.H. is a Leverhulme Research Fellow (RF-2021-533).

Data Availability Statement: Not applicable.

Acknowledgments: We are very grateful to Theunissen, F.E., Gill, P., Noopur, A., Zhang, J., Woolley, S.M.N. and Fremouw, T. for making their data available at the Collaborative Research in Computational Neuroscience—Data sharing website.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

- 1. Kozachenko, L.; Leonenko, N.N. Sample estimate of the entropy of a random vector. *Probl. Peredachi Informatsii* 1987, 23, 9–16.
- 2. Kraskov, A.; Stögbauer, H.; Grassberger, P. Estimating mutual information. Phys. Rev. E 2004, 69, 066138. [CrossRef] [PubMed]
- 3. Tobin, R.J.; Houghton, C.J. A kernel-based calculation of information on a metric space. *Entropy* 2013, 15, 4540–4552. [CrossRef]
- Houghton, C. Calculating mutual information for spike trains and other data with distances but no coordinates. *R. Soc. Open Sci.* 2015, 2, 140391. [CrossRef] [PubMed]
- 5. Houghton, C. Calculating the mutual information between two spike trains. *Neural Comput.* **2019**, *31*, 330–343. [CrossRef] [PubMed]
- 6. Victor, J.D. Binless strategies for estimation of information from neural data. Phys. Rev. E 2002, 66, 051903. [CrossRef] [PubMed]
- Victor, J.D.; Purpura, K.P. Nature and precision of temporal coding in visual cortex: A metric-space analysis. *J. Neurophysiol.* 1996, 76, 1310–1326. [CrossRef] [PubMed]
- 8. van Rossum, M.C. A novel spike distance. Neural Comput. 2001, 13, 751–763. [CrossRef] [PubMed]
- Houghton, C.; Victor, J. Measuring representational distances-the spike-train metrics approach. In *Visual Population Codes-Toward* a Common Multivariate Framework for Cell Recording and Functional Imaging; MIT Press:Cambridge, MA, USA, 2010; pp. 391–416.
 With a Mathematical Action of the spike train of the spike train of the spike train of the spike train. No. 2005, 2005 (2007)
- 10. Witter, J.; Houghton, C. A note on the unbiased estimation of mutual information. *arXiv* **2021**, arXiv:2105.08682.
- 11. Theunissen, F.E.; Gill, P.; Noopur, A.; Zhang, J.; Woolley, S.M.N.; Fremouw, T. Single-unit recordings from multiple auditory areas in male zebra finches. *CRCNS.org* **2011**. [CrossRef]
- 12. Gill, P.; Zhang, J.; Woolley, S.M.; Fremouw, T.; E. Theunissen, F. Sound representation methods for spectro-temporal receptive field estimation. *J. Comput. Neurosci.* 2006, 21, 5–20. [CrossRef] [PubMed]
- Amin, N.; Gill, P.; Theunissen, F.E. Role of the zebra finch auditory thalamus in generating complex representations for natural sounds. J. Neurophysiol. 2010, 104, 784–798. [CrossRef] [PubMed]
- 14. Kelley, D.B.; Nottebohm, F. Projections of a telencephalic auditory nucleus–Field L–in the canary. *J. Comp. Neurol.* **1979**, 183, 455–469. [CrossRef] [PubMed]
- 15. Vates, G.E.; Broome, B.M.; Mello, C.V.; Nottebohm, F. Auditory pathways of caudal telencephalon and their relation to the song system of adult male zebra finches (*Taenopygia guttata*). J. Comp. Neurol. **1996**, 366, 613–642. [CrossRef]
- Nagel, K.I.; Doupe, A.J. Organizing principles of spectro-temporal encoding in the avian primary auditory area Field L. *Neuron* 2008, 58, 938–955. [CrossRef] [PubMed]
- 17. Woolley, S.M.; Gill, P.R.; Fremouw, T.; Theunissen, F.E. Functional groups in the avian auditory system. *J. Neurosci.* 2009, 29, 2780–2793. [CrossRef] [PubMed]
- 18. Sihn, D.; Kim, S.P. A spike train distance robust to firing rate changes based on the earth mover's distance. *Front. Comput. Neurosci.* **2019**, *13*, 82. [CrossRef] [PubMed]
- 19. Wright, B.; Sen, K.; Bialek, W.; Doupe, A. Spike timing and the coding of naturalistic sounds in a central auditory area of songbirds. In Proceedings of the 2001 Neural Information Processing Systems (NIPS) Conference, Vancouver, BC, Canada, 3–8 December 2001.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.