RESEARCH ARTICLE

Inferring information flow in spike-train data sets using a trial-shuffle method

Benjamin L. Walker, Katherine A. Newhall *

Department of Mathematics, University of North Carolina at Chapel Hill, Chapel Hill, NC, United States of America

* knewhall@unc.edu

Abstract

Understanding information processing in the brain requires the ability to determine the functional connectivity between the different regions of the brain. We present a method using transfer entropy to extract this flow of information between brain regions from spike-train data commonly obtained in neurological experiments. Transfer entropy is a statistical measure based in information theory that attempts to quantify the information flow from one process to another, and has been applied to find connectivity in simulated spike-train data. Due to statistical error in the estimator, inferring functional connectivity requires a method for determining significance in the transfer entropy values. We discuss the issues with numerical estimation of transfer entropy and resulting challenges in determining significance before presenting the trial-shuffle method as a viable option. The trial-shuffle method, for spike-train data that is split into multiple trials, determines significant transfer entropy values independently for each individual pair of neurons by comparing to a created baseline distribution using a rigorous statistical test. This is in contrast to either globally comparing all neuron transfer entropy values or comparing pairwise values to a single baseline value. In establishing the viability of this method by comparison to several alternative approaches in the literature, we find evidence that preserving the inter-spike-interval timing is important. We then use the trial-shuffle method to investigate information flow within a model network as we vary model parameters. This includes investigating the global flow of information within a connectivity network divided into two well-connected subnetworks, going beyond local transfer of information between pairs of neurons.

OPEN ACCESS

Citation: Walker BL, Newhall KA (2018) Inferring information flow in spike-train data sets using a trial-shuffle method. PLoS ONE 13(11): e0206977. <https://doi.org/10.1371/journal.pone.0206977>

Editor: Kevin Scott Brown, Oregon Health and Science University - West Campus, UNITED STATES

Received: February 5, 2018

Accepted: October 23, 2018

Published: November 7, 2018

Copyright: © 2018 Walker, Newhall. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The source code for the integrate-and-fire model used to generate the spike-train data is available on github https://github.com/bwalker1/iaf_simulate.

Funding: BLW was supported by the National Institute of Health Big Data to Knowledge grant number 1T32 CA201159-01. BLW and KAN were supported by the University of North Carolina at Chapel Hill Junior Faculty Development Award. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Introduction

Understanding the flow of information in the brain is a key step in determining how the brain processes information. Methods for observing this flow of information include non-invasive methods such as MEG or calcium imaging that produce a pseudo-continuous image of activity in a brain region [1, 2] but lack high spatial resolution. Another approach, multi-unit recording, is an invasive method that inserts electrode arrays into the brain to detect when nearby neurons fire [3, 4], producing a detailed measurement of the activity of a relatively small number of neurons. This data can be processed to produce a list of spike-times for each observed

Competing interests: The authors have declared that no competing interests exist.

neuron known as spike-train data [5]. Analyzing this spike-train data can provide a window into how the neurons in the brain interact by identifying causal influences instead of just correlations in average activity.

Early methods of analyzing spike-train data include cross-correlation of spike activity and mutual information [5–7]. These methods seek to measure the similarity of two spike-trains, under the assumption that a causal relation between them will lead to a higher degree of similarity. However, these methods are symmetric, and so even if they do reveal a strong similarity they cannot indicate what the direction of causality is. Granger causality is a non-symmetric method, but requires the assumption of linear interaction [8, 9]. The desire to reveal the causal relations in data with a non-parametric approach has led to recent investigations into the use of transfer entropy (TE) [10–12].

Transfer entropy measures information transfer from a source process to a target process by quantifying the improvement in the ability to predict the future of the target process given knowledge of the past of the source process, over just knowledge of the past of the target. Uncorrelated processes will have an exact TE value of 0, and higher TE values generally correspond to greater transfers of information. Transfer entropy is defined on a pair of processes, but through pairwise evaluation can be extended to an entire network of processes such as a social network [13] or neuronal network [2, 14].

When computationally estimating TE values, statistical noise and bias in the estimator often causes positive TE values when no causal relationship is present. A natural way to determine whether the computed TE value is significant above the noise level is to choose a cutoff that represents the minimum TE value that corresponds to a significant transfer of information. However, TE values depend on the amount of information present, so neurons with higher firing rates will naturally have higher TE values. Because of this, an effort must be made to normalize the TE values onto the same scale before they can be compared, which previous papers have attempted [14, 15]. Our goal is to present a statistical method that circumvents the need to scale TE values before determining significance.

Our analysis is motivated by experimental data collected on ferrets performing an attention task [16]. This data consists of sets of 64 12-second trials performed consecutively, with recordings taken from two multi-electrode recorders, one each in the frontal lobe and the parietal lobe. Inspired by this, we consider analysis of models where the network is divided into two densely connected clusters with sparse interconnection. Also, based on the fact that the experimental task has a fixed duration, we create synthetic data consisting of many short trials instead of one long recording as used previously [17] and base our analysis on this data structure. Because the data consists of recordings of only a handful of neurons from two distant brain regions, it is unlikely that any two observed neurons will have a direct, anatomical connection. We must instead be able to detect indirect connections between the neurons, where the signal passes through one or more intermediary neurons, if we want to be able to find any flow of information present between the regions (i.e., functional connections).

Prior literature has found the creation of surrogate data to be an effective way to determine whether transfer entropy values are significant, using methods such as jittering [18], shuffling of trials [2, 19], and shuffling of inter-spike-intervals (ISIs) [20]. We have found that effective methods for determining statistical significance in transfer entropy values have not been well explored on spike-train data, which poses different challenges than continuous data. We propose the following four guiding criteria for a method to determine flow of information:

1. Strong ability to find weak/indirect connections in a network
2. Robust handling of heterogenous networks

3. No human input into the significance determination process
4. No knowledge of underlying network topology required

The reasons for these conditions are as follows: The first condition ensures that not only strong and direct connections are detected, but also weak and indirect connections, important for investigating functional connectivity. The second addresses the previously mentioned fact that neurons with higher firing rates can have falsely inflated TE values, and this could lead to false detection if the method does not adequately handle realistic variability in firing rates and coupling strengths. The third condition, pertaining to the method's design, is that the method should not involve human decision-making at any point in the process. With an increased focus on big data in computational neuroscience, human analysis falls flat in terms of scalability. Our last condition is motivated by the experimental data which consists only of recorded spike-trains, and not the structural connectivity.

We primarily consider a trial-shuffle method that meets these criteria, based on the trial shuffle for surrogate data introduced in [2, 19]. The trial-shuffle method makes use of the large number of separate trials, typical in experimental data, to independently assess significance in each pair of observed processes by statistically comparing the distribution of TE values to a computed baseline distribution, which we establish by estimating TE values for the same pair of neurons across different trials. This bypasses the problem of normalization [14] and prevents focus on only the largest TE values within the network. This baseline would also include any correlations induced by the experimental task, rather than coupled neurons, thus information flow beyond these low level correlations would be detected. We then consider how the pairwise analysis can be used to test for coarse-grain flow of information between two network regions, and the effects of partial observation where not all neurons are recorded. Furthermore, our method uses neither human intervention (e.g. [14, 15]) nor knowledge of the underlying network topology (e.g. [17]) to determine significance. The latter is especially important as it is a requirement for analyzing experimental data.

The remainder of the paper is organized as follows. In the methods section, we present our network model, based on the integrate-and-fire neuron model and then describe transfer entropy and the proposed distribution-based comparison method for determining significance. In the results section, we begin by looking at how TE estimation is influenced by neuron firing rate, coupling strength, trajectory length, and indirect connections. Next, we compare how creating the baseline distribution by shuffling trials, shuffling inter-spike-intervals, jittering, and time-shuffling behave relative to each other and the false positive method which does not fulfill the above criteria. Then, having established the validity of the trial-shuffle method, we use it to understand how the parameters in the network model affect the ability to detect transfer of information. Last, we present our discussions on these results.

Methods

Neuronal network model

We create synthetic data consisting of sequences of spike-times by numerically simulating networks of integrate-and-fire model neurons. This allows us to control the network structure and therefore the flow of information. Specifically, in a network of n neurons, the neuronal membrane potentials V_i obey the stochastic differential equations

$$dV_i = \alpha_i(\mu_i - V_i)dt + \beta_i dW_i + \sum_{j,k} c_{ji} \delta(t - \tau_{jk} - d_{ji})dt \quad i = 1, 2, \dots, n \quad (1)$$

where the μ_i are the constant driving input, the $\alpha_i > 0$ are the time-scale coefficient controlling

how fast the voltages approach μ_i , the $\beta_i > 0$ are the noise coefficients modeling variability in the neuron's input, and the $W_i(t)$ are independent standard Brownian motion processes. In addition to Eq (1) there is a reset mechanism: whenever a potential V_i reaches the threshold value of V_T , the neuron is said to fire a spike and its potential is reset to V_R . The potentials of the other neurons i are altered by a value c_{ji} after a delay d_{ji} . The times τ_{jk} are the prior spike times of neuron j .

For our simulations, we use non-dimensional variables and parameters and solve Eq (1) using a modified Euler-Maruyama method from time $t = 0$ up to time $t = t_{\max}$. We take $V_T = 1$ and $V_R = 0$, and create the spike train by binning the spike times into bins of length 0.001. Additionally, for the sake of simplicity, we will reduce the number of parameters by using $\alpha_i = \alpha = 1$, $\beta_i = \beta$, and $d_{ji} = 0.01$ corresponding to 10 time-bins. Except where otherwise noted, we additionally use $\beta = 1$.

To construct the network connections, we point out that neuronal networks are generally sparse, so most values of c_{ji} are 0, and all connections emanating from a neuron are either excitatory ($c_{ji} = w_E > 0$) or inhibitory ($c_{ji} = w_I < 0$). We consider a few classes of network topology. For investigation of basic properties, we consider several results from directed chain networks, with all excitatory connections. A directed chain network simply consists of n neurons with connections from neuron k to neuron $k + 1$ for $k = 1, 2, \dots, n - 1$. We then construct a more realistic model based on experimental data that focuses on the interaction of two regions of the brain. To model this in our simulated data, we use a directed stochastic block model [21] with 2 communities, which we describe next. In this paper, the neurons are divided equally between these two communities. We require a large number of connections within each cluster and a small number of connections between the regions. Depending on the parameters of the simulation, the connections between regions may be allowed in one, both, or neither direction. All possible connections, taken to be from neuron i to neuron j , are initialized with an independent probability of

$$p_{\text{inside}} = \frac{2k_1}{n-1}$$

if the neurons are in the same cluster. For a pair of neurons which are in different regions, the probability is either

$$p_{A \rightarrow B} = \frac{4k_A}{n^2} \quad \text{or} \quad p_{B \rightarrow A} = \frac{4k_B}{n^2}$$

depending on the direction of the connection. The parameter k_1 is the average number of outgoing connections a neuron has to other neurons within its region, k_A is the average number of outgoing connections in total from region A to region B, and vice versa for k_B . For example, for a two-region network in which there are only connections from region A to region B, one would set $k_A > 0$, $k_B = 0$. In this two-region network we randomly assign neurons to be inhibitory with probability 0.25; otherwise the neuron is excitatory.

We also consider the possibility that not all of the neurons are being observed. In this case we have n_{obs} neurons observed where $n_{\text{obs}} < n$, where the observed neurons are still evenly split between the regions. Spike trains are then only recorded and analyzed for those neurons that were observed. In this paper, $n_{\text{obs}} = n$ unless otherwise noted.

Definition of transfer entropy

Transfer entropy quantifies the improved predictability about the future of a given temporal sequence Y when including the history of another temporal sequence X , over just the history

of Y . Thereby, it attempts to measure the influence of X on Y , and in the context of neuronal networks, measure the functional connectivity from X to Y . Specifically, the transfer entropy from the discrete sequence of random variables X_t to Y_t is expressed as [22]

$$TE_{X \rightarrow Y} = H(Y_t | Y^-) - H(Y_t | Y^-, X^-) \quad (2)$$

where $H(A|B)$ denotes the conditional Shannon entropy

$$H(A|B) = \sum_{i,j} p_{A,B}(a_i, b_j) \log \frac{p_B(b_j)}{p_{A,B}(a_i, b_j)}.$$

Here, we write the formula for discrete random variables A and B with probability mass functions $p_A(a)$, $p_B(b)$ and joint probability mass function $p_{A,B}(a, b)$. In Eq (2), Y_t is the value of sequence Y at time t , and X^- and Y^- refer to the past states of the sequences X_t and Y_t .

In practice, we do not know the true distributions of the random variables, only a collection of observations. The probability distributions must therefore be estimated from these observations. We use a MATLAB toolbox [17] to numerically estimate the discrete transfer entropy values. The code takes as input a collection of sequences of sorted spike times. For each pair, it counts occurrences of spikes in both processes separated by a specified delay, and uses these counts to estimate the probability distributions that appear in the TE formula, resulting in a TE value for each pair of processes. This toolbox takes parameters for the interaction delay, and the history length for the source and target neurons. In this paper, we use a history length of 5 for the source neuron and 3 for the target neuron. To determine the interaction delay, we compute TE values for a range of possible delays and for each pair independently take the maximum TE value over this range. For the sake of computational efficiency, we consider delays that are multiples of 10 time-bins (0.01), up to a maximum of 200 time-bins (0.2).

Note that with this estimation method, even totally independent sequences will likely produce positive transfer entropy values due to statistical noise. This means that we can only claim that a positive transfer entropy value indicates information transfer if it is large relative to the bias of the estimator, which depends on the input series. There is further discussion of this problem in Properties of Transfer Entropy on our Model section in Results.

Throughout the paper we only consider detecting information transfer when the source neuron is excitatory. Although an inhibitory connection will theoretically produce a non-zero TE value, for relatively infrequent firing regimes, the absence of a firing event contains much less information than the presence of a firing event induced by an excitatory connection. Here, infrequent firing is relative to the time-binning of the data; a small fraction of time bins contain a firing event.

Distribution-based comparisons

In order to automatically handle the effects of varying neuron firing rate and other factors on the estimated TE values, we consider the usage of pairwise methods that determine the significance of the TE value between each pair of neurons in isolation from all other pairs. This prevents the possibility of neurons being overshadowed by other neurons with higher estimated TE values.

Prior literature has used methods in which a statistical baseline is computed using some sort of Monte Carlo method, and then the TE values are compared to this baseline distribution [23]. Using a baseline distribution can account for fluctuations in properties like firing rate as long as the assumption that each observation comes from the same underlying distribution. Inspired by typical experimental data consisting of multiple repeated trials, we create a baseline for each pair of neurons by computing the TE from the source neuron in one trial to the target

neuron in a different trial. Due to the time separation between trials, there cannot be actual transfer of information between the neurons, and any positive TE values must therefore be insignificant. The signed-rank test is used to determine if the distribution of estimated TE values from each of the trials is significantly greater than the distribution of TE values taken from time-shuffled comparisons. We next detail the trial-shuffle method, other methods for creating this baseline, the false-positive method, and then describe how we determine overall connectivity between regions.

Trial shuffle. The primary method considered in this paper is the trial-shuffle method, in which a number of “replaced” trials are assembled, where the data for each neuron is taken from a different original trial. For example, one might combine the data for neuron 1 from trial 1, neuron 2 from trial 2, etc. into a single replaced trial. By assembling the data this way, we can compute all the pairwise TE values for a shuffled trial and produce baseline values for all pairs simultaneously. Since typically fewer neurons are observed than trials exist, it is possible to construct data in this way, with each spike train in a shuffled trial coming from a different experimental trial. This method is also not limited to having fewer observed neurons than trials. In such a case, not all the baselines could be computed simultaneously, but could still be computed independently.

From the shuffled trials, we obtain a distribution of insignificant TE for each individual pair of neurons to serve as the baseline distributions for comparison. At the same time, we compute all of the pairwise TE values for each of the experimental trials, without shuffling. These values are collected to create the experimental distribution of TE values for each pair of neurons. At this point, we deem the connection to exhibit a significant transfer of information if the values in the experimental distribution for that pair of neurons are statistically significantly greater than those in the respective baseline distribution.

ISI shuffle. For the ISI shuffle [20], the baseline spike trains are created by shuffling the relative offsets between spike times. In other words, if spikes occur at times $t_1, t_2, t_3, \dots$ and have offsets between them of $o_1 = t_1, o_2 = t_2 - t_1, o_3 = t_3 - t_2, \dots$ we shuffle the set of $\{o_i\}$ and reconstruct the spike times accordingly.

Jittering. For jittering [18], surrogate baseline data is produced by moving each spike time in an existing spike train by a certain random offset. In this paper, we generate the offsets using a uniform random distribution from -20 time-bins to 20 time-bins, with resampling if the new location bin already contains a spike.

Time shuffle. For the time shuffle, a pairwise baseline is created by randomizing the spike times of each trial, such that each spike time is moved to a new time that is uniformly randomly chosen within the total range of time, with resampling if shuffled spike times fall in the same time-bin.

False positive method. The false positive method [17], unlike the other methods presented above, is not a pairwise comparison. In [17], all computed TE values are compared to a single cutoff chosen to enforce a certain false positive rate. To define a false positive, we classify a directed pair of neurons (a, b) as such:

1. Direct Connection—Neuron a has a direct structural connection to neuron b ($c_{ab} \neq 0$)
2. Indirect Connection—A sequence of multiple structural connections connect neuron a to neuron b .
3. No Connection—It is impossible to reach neuron b from neuron a traveling through the network.

We define a false positive as a connection of type 3, where no transfer of information is possible, that is detected as having a significant transfer of information. To extend this approach

to the multiple trial data format, we take the mean TE value across all of the trials for each pair of neurons, and then apply the standard false positive cutoff approach by choosing the cutoff to be the value such that only 5% of unconnected pairs of neurons have a mean TE value above the cutoff. A pair of neurons whose mean TE is above the cutoff is then deemed significant.

Assessing inter-region connectivity. We are interested in being able to determine, from the pairwise connection results, whether there is evidence of information flow between the two regions in our model. We test the number of significant connections between regions against the null hypothesis that all pairwise connections have an independent 5% chance of being detected. This significance level is chosen because our tests are done at a 5% significance level. The 95th percentile of the binomial distribution $\text{Binomial}(n_{\text{obs}}^2/4, 0.05)$ when observing 50 neurons, $n_{\text{obs}} = 50$, gives the critical value of 40 connections.

Results

Properties of transfer entropy on our model

We will begin by illustrating how the transfer entropy estimation is affected by the input. We show that the estimated TE values can drastically vary independent of the actual causal relation between the neurons. This demonstrates the need for careful determination of the significance of TE values to infer functional connectivity.

Relation between transfer entropy bias and trajectory length. While theoretically, only correlated sequences would have non-zero transfer entropy values, in practice, the estimation of distributions leads to a positive bias in the estimated TE values. By the law of large numbers, the bias decreases as the observation length of the processes increases and more data is used to estimate the distributions.

To get an intuitive idea of how this error should behave, consider first the simpler case of trying to estimate entropy of a stationary process. Consider a stationary process $Y_t \in \{0, 1\}$ consisting of a sequence of independent, identically distributed Bernoulli random variables with mean p . Clearly the entropy of such a process will be

$$H(Y) = -(p \log p + (1 - p) \log (1 - p)) = \theta(p).$$

If we do not know p , and only know observations of Y_t for $t = 1, 2, \dots, m$, then we can't compute $\theta(p)$. Instead, we can let $\hat{p}$ be an estimator of p and then estimate the entropy as $\hat{H}(Y) = \theta(\hat{p})$. The question arises: what will be the error in this estimate?

Clearly $\theta(p)$ is a concave function, which tells us that for any imperfect estimator $\hat{p}$,

$$\mathbb{E}_{\hat{p}} \theta(\hat{p}) < \theta(\mathbb{E}_{\hat{p}} \hat{p})$$

by Jensen's inequality. We will specifically consider the estimator $\hat{p} = \frac{1}{m} \sum_{t=1}^m Y_t$ which is unbiased ($\mathbb{E}_{\hat{p}} \hat{p} = p$) so from the above equation it will always produce an entropy estimate less than the true value. The estimator $\hat{p}$ is distributed as $\frac{1}{m} \text{Binomial}(m, p)$, and by taking a Taylor expansion about the peak p of the distribution of $\hat{p}$,

$$\mathbb{E}_{\hat{p}} \theta(\hat{p}) \approx \theta(p) + \theta'(p) \mathbb{E}_{\hat{p}} [\hat{p} - p] + \frac{1}{2} \theta''(p) \mathbb{E}_{\hat{p}} [(\hat{p} - p)^2] = \theta(p) + \frac{\theta''(p)}{2} \text{Var}(\hat{p}).$$

We find that the bias in this entropy estimator is approximately

$$\frac{1}{2} \theta''(p) \text{Var}(\hat{p}) = \frac{1}{2} \left(-\frac{1}{1-p} - \frac{1}{p} \right) \frac{p(1-p)}{n} = -\frac{1}{2n}.$$

Given this inverse relationship with the sample length in the simple entropy case, we might expect to see the same inverse relationship when estimating transfer entropy of our spike-train data. To demonstrate this relationship, we simulated a network of 50 neurons with no connections, and estimated the TE for various observation lengths. These TE values, averaged over the population, are shown in Fig 1(a). We can see that the estimated TE between independent processes does decay with increased observation time, and for sufficiently long time it follows an inverse relationship.

Firing rate vs. transfer entropy. When determining significance, one might attempt to compare computed TE values across all pairs in the network to obtain a single cutoff value, above which indicates a transfer of information. However, since factors like firing rate affect the magnitude of the TE values, pairs with greater TE values do not necessarily have a stronger flow of information or even a connection at all. Intuitively, this increase in TE value with firing rate can be understood by considering that spike-train data consists of a sequence of correlated Bernoulli random variables which maximize their entropy at $p = 0.5$. Because the neuron firing rate per bin is much less than 0.5, an increase in firing rate increases the entropy of the process and correspondingly increases the upper bound for the transfer entropy. Due to the nature of the estimator, we would also expect that independent processes have higher estimated TE values when firing rates are higher.

Given this, we would like to understand the effect that the firing rate of neurons has on the estimated TE between them. We expect that higher firing rate neurons will have inflated TE values, independent of whether or not the pairs are actually connected. Fig 1(b) shows the sharp increase in TE values for higher firing rate neurons, for a collection of neurons with no connections between them. This demonstrates why some sort of normalization is required if one wishes to compare estimated TE values from different pairs of neurons with different firing rates, especially for a collection of heterogeneous neurons as one might find *in vivo*. This normalization is difficult even for just the firing rate due to its asymmetric non-linear dependence shown in Fig 1(b). Furthermore, firing rate is not the only factor influencing TE values.

Coupling strength vs. transfer entropy. Naturally, neurons connected with a higher coupling strength will generally have higher TE values observed between them. This bears relevance on our ability to detect weaker connections, which are harder to detect when the signal-to-noise ratio is low. We illustrate this effect with a pair of neurons by varying their (excitatory) coupling strength.

The data, as shown in Fig 1(c), suggests that for small coupling strength the estimated TE values are roughly proportional to the square of the coupling strength. As might be expected, once the coupling strength is comparable to the distance between the reset value and the firing threshold of the model neurons, the target neuron always fires after the source neuron does and therefore further increasing the coupling strength does not actually increase the transfer of information from the source neuron to the target neuron, causing the TE value to plateau. Notice how a low-firing rate, strongly coupled neuron pair and a high-firing rate, weakly coupled neuron pair could therefore both produce the same TE value.

Indirect connections. Another issue in detecting functional connections in experimental data is that typically only a fraction of neurons are observed, so it is unlikely that one would always observe pairs of directly connected neurons. In order to detect functional connectivity, we must therefore be able to detect indirect connections. Here, we review how estimated transfer entropy values are affected by increasingly indirect connections. Fig 1(d) shows the transfer entropy computed down the length of three chains of varying coupling strength. All three chains show that the direct connection from the head neuron to the first neuron down the chain is very significant, but they vary significantly in how the transfer entropy values decay further down the chain. In the weakest coupling strength chain, the values decay almost

Fig 1. (a) Black line: Relationship between observation length and estimated TE averaged over a simulated network of disconnected neurons. Red line: sample slope of -1. For this simulation $n = 50$ and $\mu = 1$. (b) Relationship between neuron firing rate and calculated TE values showing an asymmetric non-linear dependence. The simulation is performed with $n = 200$ unconnected neurons. μ values range linearly from 0 to 16, and $t_{max} = 10$. (c) Black line: Relationship between coupling strength and TE values showing an initial quadratic dependence before saturating due to the target neuron always firing after the source neuron. Red dashed line: Quadratic model fit to data for w_E from 0 to 1, with $R^2 = 0.99552$ (fit parameters in Appendix A). For this simulation, $n = 2$, $t_{max} = 10$, $\mu = 0.25$ and w_E is given on the x-axis. (d) Computed TE values for indirectly connected neurons down a chain with three

different coupling strengths $w_E = 0.5, 0.8, 1.5$ showing that higher w_E causes significant TE values further down the chain. For this simulation $n = 20, \mu = 0.8, t_{max} = 100$.

<https://doi.org/10.1371/journal.pone.0206977.g001>

immediately to noise levels, meaning detecting more than one or two neurons down the chain would be difficult. In the medium coupling strength chain, there appears to be an exponential decay, falling to noise levels after about 5 neurons. However, in the strongest coupling strength chain, the transfer entropy remains a significant fraction of the directly-connected neurons' TE value for the entire length of the chain. This is not entirely surprising as the large value of $w_E = 1.5$, relative to $V_T - V_R = 1$ (the distance from reset to threshold voltage), allows for a wave of firing to frequently move from the head neuron all the way down the chain. This wave of firing is only interrupted if the noise, which here has strength $\beta = 1$, has driven the voltage of a neuron below -0.5 when the previous neuron fires. Given this is fairly unlikely to happen, high TE values arise for even the furthest neuron down the chain.

Validation of method

We establish the validity of using the trial-shuffle method to further investigate flow of information by making a qualitative comparison to the other methods presented in Methods. We begin by creating synthetic spike-train data according to the two-region network model as explained previously in Methods. We simulate 64 trials and use this same data for each of the significance testing methods to identify connections. To parallel previous work [17] we will report the fraction of connections that were identified for various lengths of connection, and we will also compare the number of false positives from each method.

Fig 2 shows the percentage of connections present in the network that are identified as significant by each method for three different sets of model parameters. Recall, we only consider detecting connections starting from excitatory neurons. Connections of length above 5 are not shown as they are not present in great enough quantity to produce meaningful information. The network behavior of the neurons for each set of parameters tested in Fig 2 is shown in Fig 3.

While we are primarily interested in the four pairwise distribution comparison methods, we also show the false positive method for comparison. We reiterate that it requires as input the full network topology, which includes the true solution to the detection problem, and so is not a valid comparison. However, the parameters and randomly initialized structure of the network significantly affect how easy it is to detect connections in the network, and so the performance of the false positive method provides a simple reference point for how easy it is to detect connections in a given network. We stress that comparisons should not be made between the false positive method and the other methods.

Throughout the parameter space (three typical cases shown) we see a general pattern of the trial-shuffle method and the ISI-shuffle method detecting more connections, with the jittering method in the middle and the time-shuffle method clearly detecting the least. These results indicate that both the trial-shuffle method and the ISI-shuffle method would be reasonable options for further investigation of transfer entropy behavior.

Of course, the detection must also be judged in light of the false positive rate, as increased detection is not valuable if it is not restricted to those connections that are actually present. In Table 1 we present the false positive rates associated with each of the four pairwise distribution comparison methods we consider. Recall that all statistical tests were performed at a 5% significance level.

We break up the false positive rates into the overall false positive rate and the $B \rightarrow A$ false positive rate. Recall that the network contains two regions, A and B , for which inter-region

Fig 2. Comparison of the false positive and distribution-based methods in terms of fraction of detected connections vs. path length in the network. For all frames $n = k_A = 50$, $k_B = 0$, $t_{max} = 10$. (a) $w_E = 0.4$, $w_I = -0.5$, $\mu = 10$ and $k_1 = 2$ (b) $w_E = 0.27$, $w_I = -0.4$, $\mu = 3$ and $k_1 = 2$ (c) $w_E = 0.35$, $w_I = -0.45$, $\mu = 2$ and $k_1 = 3$.

<https://doi.org/10.1371/journal.pone.0206977.g002>

connections exist only from A to B . Because our research is motivated by the search for macro-scale information flow, we are particularly interested in avoiding those false positives that indicate fictitious transfer of informations between regions, i.e. $B \rightarrow A$.

We see that the trial-shuffle method has lower false positive rates than the ISI-shuffle method which had similar levels of detection. The jittering method has similar false positive rates. The time-shuffle method has very low false positive rates in panels (a) and (b), which when combined with its low true detection rate indicate that this method of creating surrogate data leads to difficulty finding any connections at all.

From this data we can conclude that the trial-shuffle method does appear to be a reasonable way to create surrogate data to test for transfer entropy significance. We will now investigate how the model and simulation parameters affect the ability of the trial-shuffle method to find connections in a two-region network.

Inferring information transfer

We will now show the detection of information flow as a function of network parameters coupling strength and connectivity. Specifically, we will use the trial-shuffle method to determine global causal links between regions in the two-region network model.

A natural question is the relationship between the coupling strength of two neurons and our ability to detect a connection between them. Obviously, we would expect a stronger

Fig 3. Raster plots and inter-spike-interval (ISI) distribution for the two-region network model in various parameters regimes. (a) through (c) correspond to (a) through (c) in Fig 2. We try to avoid the dynamics shown in panel (d) with $k_1 = 10$, $w_E = 0.5$, $w_I = -2$, $\mu = 1$. This case has no interesting transfer of information due to the highly synchronous behavior.

<https://doi.org/10.1371/journal.pone.0206977.g003>

Table 1. False positive rates corresponding to Fig 2. “All” refers to connections found relative to all pairs without any type of structural connection. “B→A” refers to just connections from region B to region A for which no connections of any kind exist.

method	panel a		panel b		panel c	
	all	B→A	all	B→A	all	B→A
Trial Shuffle	7.7	4.2	3.7	2.5	3.6	2.5
ISI Shuffle	10.0	6.4	4.5	4.2	6.9	6.4
Jittering	4.1	1.7	4.9	4.1	3.8	2.4
Time Shuffle	1.4	0.16	2.2	1.6	5.4	4.9

<https://doi.org/10.1371/journal.pone.0206977.t001>

excitatory connection to be more likely to be detected as significant. To investigate this, we created a directed chain network of 20 neurons, and investigated the detected connections originating from the neuron at the head of the chain. In theory, this neuron is connected to all 19 neurons down the chain from it; however, when the coupling strength is low these connections will likely not all be able to be detected. For an ensemble of 1000 trials, we observed how many neurons down the chain were detected as connected to the head neuron, and these results are shown in Fig 4. We therefore see, depending on the coupling strength, how long a sample must be taken to reliably detect direct connections or some indirect connections.

Especially visible in the $t_{max} = 100$ case is the nonlinear structure of the data. As the coupling strength increases, the probability of detecting a connection to the first neuron increases quadratically. This intuitively aligns with what we previously found in Fig 1(c), that the transfer entropy itself increases quadratically with coupling strength. After the coupling strength is high enough that the first connection can be very consistently detected, the second connection does not immediately become detectable. Instead, there is a region in which increasing the coupling strength does not increase the length of connection detected.

The most important thing we want to understand is how the degree of connectivity in the two-region network relates to the ability of our method to correctly identify the flow of information. We will be again considering the two-region network with monodirectional flow, e.g.

Fig 4. For a range of connection strengths w_E and three total simulation times t_{max} , how many neurons down the chain are detected as connected to the head, on average. The leftmost quadratic model is fit on x from 0 to 0.09, with $R^2 = 0.99109$. The middle quadratic model is fit on x from 0.2 to 0.32, with $R^2 = 0.98718$. The rightmost quadratic model is fit on x from 0.35 to 0.5, with $R^2 = 0.98325$. (fit parameters in Appendix A) For this figure we use $n = 20$, $\mu = 1$. All connections are excitatory with coupling strength as given on the x -axis, and t_{max} is specified in the legend.

<https://doi.org/10.1371/journal.pone.0206977.g004>

from region $A \rightarrow B$ but not from $B \rightarrow A$. The connectivity can be influenced both by the edge probability parameters k_1 and k_A and the number of neurons N in the network. We would like to find a way to quantify the amount of inter-region connectivity across networks with variation in all of these parameters. To this end, we introduce a variation of the global efficiency [24], which is the average of the inverse shortest path length between all pairs of nodes. Our variation measures the directed efficiency from region A to B by considering only those paths from a neuron in region A to a neuron in region B, and then averaging the inverse shortest path length.

We generated a large number of networks of varying sizes from $n = 50$ to 1000, only observe $n_{\text{obs}} = 50$ neurons, $k_1 = 2$ or 3, and $k_A \in \{n/2, n/2 + 1, \dots, 2n\}$, leading to a wide range of directed efficiencies. For each of these networks we run 10 sessions of 64 trials to get a quantification of the spread in potential detection. Fig 5 shows the number of detected connections for each network against the directed efficiency. This confirms the expected, that more connected networks have a higher number of connections detected. At an efficiency of approximately 0.28, the method begins to always detect the regions as being connected, although it frequently still does so for lower values. Also of interest is the fairly narrow spread in detection across the substantial variation in parameters for the lower efficiency range. This suggests that our method is stable across a variety of networks.

We would also like to see how the ability of our method to detect connections improves using trials of increasing time length. Recall Fig 1(a), which shows that the noise in the TE estimator is inversely proportional to the length of time in the trial. This suggests that a longer trial should have a higher signal-to-noise ratio and correspondingly better detection. Additionally, since neurons are actually connected in this network, the longer time means that there are more chances for a neuron to affect another, meaning there is more information transfer for the method to find. Fig 6 shows the number of significant connections detected between regions in each direction, for a network that only contained $A \rightarrow B$ inter-region connections with 50 neurons. In line with our expectations, we see a steady increase in the number of connections we detect as the length of the simulation increases.

Fig 5. $A \rightarrow B$ efficiency vs detection. The red line indicates the significant level of connections indicating transfer of information between regions. Here we use $t_{\text{max}} = 10$, $\mu = 1$, $w_E = 0.5$, $w_I = -1$. The networks are randomly initialized with n ranging from 50 to 1000 with $n_{\text{obs}} = 50$, $k_1 \in \{2, 3\}$, and k_A ranging from $n/2$ to $2n$ with $k_B = 0$.

<https://doi.org/10.1371/journal.pone.0206977.g005>

Fig 6. Number of detected connections vs. simulation time for each trial. Blue = $A \rightarrow B$, red = $B \rightarrow A$. The black line indicates the level above which the number of connections is statistically significant ($p < 0.05$). This simulation uses a two-region network with $n = 50$, $k_1 = 3$, $k_A = 50$, $\mu = 3$, $w_E = 0.27$, $w_I = -0.3$, and t_{max} given on the x-axis.

<https://doi.org/10.1371/journal.pone.0206977.g006>

Discussion

We use a pairwise comparison method to determine the significance of transfer entropy values in an objective way for spike-train data comprised of multiple trials. By shuffling trials, we create surrogate data to serve as a baseline for insignificant TE values and then performing a statistical test to see if the true TE values are greater than the baseline TE values. We also compare to several other options for creating such surrogate data: shuffling the spike times, shuffling the inter-spike-intervals, and jittering the spike times. Our analysis showed that the trial-shuffle and ISI-shuffle methods had the highest detection, and the trial-shuffle method had slightly lower false positive rates, which justified our decision to look further into the behavior specifically of the trial-shuffle method. We hypothesize that the reason why the trial-shuffle and ISI-shuffle methods have the highest detection is that while all of the methods preserve the firing rate in the spike trains, only the trial shuffle and ISI shuffle preserve the temporal structure of the data.

This trial shuffling method focused on data separated into trials, to allow for a coarse-grained shuffling in time. The time shuffling method is another extreme of this shuffling in which the data is shuffled on the order of a single observation time step. The fact that the time-shuffle method fails to accurately determine significance whereas the trial-shuffle method does not suggests that shuffling is only valid when done on a sufficiently coarse timescale. There is presumably some “minimum time window” that could be shuffled, such that the temporal structure of the spike trains is still preserved. We hypothesize that this time window must be greater than the timescale at which information flows through the network. Breaking up data on such a timescale and then shuffling it will allow the trial-shuffle method to be extended to spike-train data consisting only of a single long recording, instead of a number of separate recorded trials, as long as analysis confirmed a robust choice of shuffling timescale.

We observed interesting features of information flow on a network. Looking at how far down a directed chain of connected neurons transfer entropy detects significant information flow revealed a quadratic dependence of the distance down the chain on the coupling strength. This aligns with the observed quadratic dependence of the estimated transfer entropy on the coupling strength. We also observed the global flow of information from one network region to the other as a function of network connectivity. While increased detection of connections is

expected as the connectivity of the network increases, we note the fairly narrow spread in detection across substantial variation in parameters. One might expect these general trends to remain when considering different transfer entropy estimating methods, or more realistic models for producing spike-train data.

A modification to the neuronal network model could allow for testing on a more organized network in future work. Various methods exist to organize information flow in a neuronal network, including Spike-Timing Dependent Plasticity (STDP) and other extensions of the Hebbian learning philosophy. A meaningful external drive could also be applied to create a source of non-random information. These methods could be used to create a network that has more concentrated connections that direct information flow instead of diffusing it as in a random network. Computational advances could also help bridge the wide gap between the number of neurons we used in simulated networks, up to 1000, and the many billions of neurons actually present in the human brain. This would allow us to see the effect that the vastly sparser observation has on the results, when coupled with the increase in organization as described above.

In this paper we chose to use a basic model for simulating spike-train data. However, more sophisticated models such as the Izhikevich model [25] or a Hodgkin-Huxley model with adaptation [26] could lead to different sorts of spiking behaviors in addition to greater historical dependence of the current behavior. These effects would certainly affect the problem of detecting TE values, and so would provide an interesting followup to the work presented in this paper.

While we chose one transfer entropy estimator in our analysis, we expect many of the features we observed in the behavior of information detection to remain regardless of the TE estimator. Many variations and extensions of transfer entropy have been explored in the literature. For example, Ito et al. [17] suggest that first order TE might be better than higher order TE unless the orders are tuned. Multivariate extensions of transfer entropy, which in addition to controlling for information contained in the past of the target process, can control for information present in other processes besides the source, provide transfer entropy values that reflect the information coming from the source neuron and not from any of the other neurons. These tools, combined with an effective method to determine significance such as the trial-shuffle method, could be used to create functional connectivity graphs that show how information is actually moving between the neurons based on spike-train data. This would then allow a more in-depth understanding of how the neuronal networks in the brain actually processes information.

Real-world data brings many additional challenges when looking for functional connectivity. For one, the trial-shuffle method assumes each trial is an independent sample from the same distribution, thus cannot account for adaptation or depression across multiple trials. Also, we expect connections to be in general weaker, and the number of neurons is many times greater than used in our simulated tests. This is balanced to a degree by the fact that neurons in the brain are extremely well organized to facilitate the flow of information, in contrast to the random connections in our simulated network. This may lead to greater sensitivity of the method, as detection could depend heavily on whether an organized pathway is actually picked up by the probes.

We might try to mitigate the real-world challenge of detection by obtaining more data, as we have shown the ability to detect connections scales with the amount of observed data. Depending on the nature of the experiment, it may or may not be possible to increase the length of the observed trials. Instead, more data could be obtained by increasing the number of trials. However, the analysis requires that the same neurons be recorded consistently over all of the trials, and small movements of the probes over time can potentially cause neurons to

be lost in the recording. These challenges in obtaining large amounts of experimental data highlight the importance of a statistical analysis that is as powerful as possible.

Appendix A: Quadratic fit parameters

The parabola fit in Fig 1(c) is $y = -8.6573 \times 10^{-6} + 0.0030409x + 0.010546x^2$.

The parabola fits in Fig 4 are $y = 0.043545 - 5.2788x + 166.67x^2$ for the leftmost, $y = 4.0181 - 28.291x + 67.208x^2$ for the middle, and $y = 9.2834 - 38.99x + 52.425x^2$ for the rightmost.

Author Contributions

Conceptualization: Katherine A. Newhall.

Data curation: Benjamin L. Walker.

Formal analysis: Benjamin L. Walker.

Funding acquisition: Katherine A. Newhall.

Investigation: Benjamin L. Walker, Katherine A. Newhall.

Methodology: Benjamin L. Walker, Katherine A. Newhall.

Software: Benjamin L. Walker.

Validation: Benjamin L. Walker, Katherine A. Newhall.

Visualization: Benjamin L. Walker.

Writing – original draft: Benjamin L. Walker.

Writing – review & editing: Benjamin L. Walker, Katherine A. Newhall.

References

1. Orlandi JG, Stetter O, Soriano J, Geisel T, Battaglia D. Transfer entropy reconstruction and labeling of neuronal connections from simulated calcium imaging. *PloS one*. 2014; 9(6):e98842. <https://doi.org/10.1371/journal.pone.0098842> PMID: 24905689
2. Vicente R, Wibral M, Lindner M, Pipa G. Transfer entropy: a model-free measure of effective connectivity for the neurosciences. *Journal of computational neuroscience*. 2011; 30(1):45–67. <https://doi.org/10.1007/s10827-010-0262-3> PMID: 20706781
3. Green J. A simple microelectrode for recording from the central nervous system. *Nature*. 1958; 182(4640):962–962. <https://doi.org/10.1038/182962a0> PMID: 13590200
4. Humphrey DR, Schmidt EM. In: Boulton AA, Baker GB, Vanderwolf CH, editors. *Extracellular Single-Unit Recording Methods*. Totowa, NJ: Humana Press; 1990. p. 1–64. Available from: <https://doi.org/10.1385/0-89603-185-3:1>.
5. Brown EN, Kass RE, Mitra PP. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature neuroscience*. 2004; 7(3):456–461. <https://doi.org/10.1038/nn1228> PMID: 15114358
6. Borst A, Theunissen FE. Information theory and neural coding. *Nature neuroscience*. 1999; 2(11):947–957. <https://doi.org/10.1038/14731> PMID: 10526332
7. Nirenberg S, Carcieri SM, Jacobs AL, Latham PE. Retinal ganglion cells act largely as independent encoders. *Nature*. 2001; 411(6838):698–701. <https://doi.org/10.1038/35079612> PMID: 11395773
8. Granger CW. Some recent development in a concept of causality. *Journal of econometrics*. 1988; 39(1-2):199–211. [https://doi.org/10.1016/0304-4076\(88\)90045-0](https://doi.org/10.1016/0304-4076(88)90045-0)
9. Kamiński M, Ding M, Truccolo WA, Bressler SL. Evaluating causal relations in neural systems: Granger causality, directed transfer function and statistical assessment of significance. *Biological cybernetics*. 2001; 85(2):145–157. <https://doi.org/10.1007/s004220000235> PMID: 11508777
10. Schreiber T. Measuring information transfer. *Physical review letters*. 2000; 85(2):461–464. <https://doi.org/10.1103/PhysRevLett.85.461> PMID: 10991308
11. Kaiser A, Schreiber T. Information transfer in continuous processes. *Physica D: Nonlinear Phenomena*. 2002; 166(1):43–62. [https://doi.org/10.1016/S0167-2789\(02\)00432-3](https://doi.org/10.1016/S0167-2789(02)00432-3)

12. Timme NM, Ito S, Myroshnychenko M, Nigam S, Shimono M, Yeh FC, et al. High-degree neurons feed cortical computations. *PLoS computational biology*. 2016; 12(5):e1004858. <https://doi.org/10.1371/journal.pcbi.1004858> PMID: 27159884
13. Ver Steeg G, Galstyan A. Information transfer in social media. In: *Proceedings of the 21st international conference on World Wide Web*. ACM; 2012. p. 509–518.
14. Gourévitch B, Eggermont JJ. Evaluating information transfer between auditory cortical neurons. *Journal of Neurophysiology*. 2007; 97(3):2533–2543. <https://doi.org/10.1152/jn.01106.2006> PMID: 17202243
15. Li Z, Li X. Estimating temporal causal interaction between spike trains with permutation and transfer entropy. *PloS one*. 2013; 8(8):e70894. <https://doi.org/10.1371/journal.pone.0070894> PMID: 23940662
16. Sellers KK, Yu C, Zhou ZC, Stitt I, Li Y, Radtke-Schuller S, et al. Oscillatory dynamics in the frontoparietal attention network during sustained attention in the ferret. *Cell reports*. 2016; 16(11):2864–2874. <https://doi.org/10.1016/j.celrep.2016.08.055> PMID: 27626658
17. Ito S, Hansen M, et al. Extending Transfer Entropy Improves Identification of Effective Connectivity in a Spiking Cortical Network Model. *PloS one*. 2011; 6(11):e27431. <https://doi.org/10.1371/journal.pone.0027431> PMID: 22102894
18. Nigam S, Shimono M, Ito S, Yeh FC, Timme N, Myroshnychenko M, et al. Rich-club organization in effective connectivity among cortical neurons. *Journal of Neuroscience*. 2016; 36(3):670–684. <https://doi.org/10.1523/JNEUROSCI.2177-15.2016> PMID: 26791200
19. Lindner M, Vicente R, Priesemann V, Wibral M. TRENTOOL: A Matlab open source toolbox to analyse information flow in time series data with transfer entropy. *BMC neuroscience*. 2011; 12(1):119. <https://doi.org/10.1186/1471-2202-12-119> PMID: 22098775
20. Rivlin-Etzion M, Ritov Y, Heimer G, Bergman H, Bar-Gad I. Local shuffling of spike trains boosts the accuracy of spike train spectral analysis. *Journal of neurophysiology*. 2006; 95(5):3245–3256. <https://doi.org/10.1152/jn.00055.2005> PMID: 16407432
21. Wang YJ, Wong GY. Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association*. 1987; 82(397):8–19. <https://doi.org/10.1080/01621459.1987.10478385>
22. Wibral M, Vicente R, Lindner M. *Transfer Entropy in Neuroscience*. In: *Directed Information Measures in Neuroscience*. Springer, Berlin, Heidelberg; 2014.
23. Bauer M, Cox JW, Caveness MH, Downs JJ, Thornhill NF. Finding the direction of disturbance propagation in a chemical process using transfer entropy. *IEEE transactions on control systems technology*. 2007; 15(1):12–21. <https://doi.org/10.1109/TCST.2006.883234>
24. Latora V, Marchiori M. Efficient behavior of small-world networks. *Physical review letters*. 2001; 87(19):198701. <https://doi.org/10.1103/PhysRevLett.87.198701> PMID: 11690461
25. Izhikevich EM. Simple model of spiking neurons. *IEEE transactions on neural networks*. 2003; 14 6: 1569–72. <https://doi.org/10.1109/TNN.2003.820440> PMID: 18244602
26. Woo J, Miller CA, Abbas PJ. Simulation of the electrically stimulated cochlear neuron: modeling adaptation to trains of electric pulses. *IEEE Transactions on Biomedical Engineering*. 2009; 56(5):1348–1359. <https://doi.org/10.1109/TBME.2008.2005782> PMID: 19473930