

## METHODS

# Inferring effective networks of spiking neurons using a continuous-time estimator of transfer entropy

David P. Shorten <sup>1,2</sup>, Viola Priesemann<sup>3\*</sup>, Michael Wibral <sup>4</sup>, Joseph T. Lizier<sup>5\*</sup>

**1** School of Computer and Mathematical Sciences, The University of Adelaide, Adelaide, South Australia, Australia, **2** Centre for Complex Systems, Faculty of Engineering, The University of Sydney, Sydney, New South Wales, Australia, **3** Max Planck Institute for Dynamics and Self-Organization, Göttingen, Germany, **4** Campus Institute for Dynamics of Biological Networks, Georg August University, Göttingen, Germany, **5** School of Computer Science and Centre for Complex Systems, Faculty of Engineering, The University of Sydney, Sydney, New South Wales, Australia

\* [viola.priesemann@ds.mpg.de](mailto:viola.priesemann@ds.mpg.de) (VP); [joseph.lizier@sydney.edu.au](mailto:joseph.lizier@sydney.edu.au) (JTL)



## OPEN ACCESS

**Citation:** Shorten DP, Priesemann V, Wibral M, Lizier JT (2025) Inferring effective networks of spiking neurons using a continuous-time estimator of transfer entropy. *PLoS Comput Biol* 21(10): e1013500.

<https://doi.org/10.1371/journal.pcbi.1013500>

**Editor:** Luca Faes, University of Palermo: Universita degli Studi di Palermo, ITALY

**Received:** November 26, 2024

**Accepted:** September 5, 2025

**Published:** October 22, 2025

**Copyright:** © 2025 Shorten et al. 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 cell culture data used here is openly available and can be downloaded from: <https://potterlab.bme.gatech.edu/development-data/html/index.html>. The network inference algorithms and transfer entropy estimator used in this work have been added to JIDT, which is freely available from <https://github.com/jlizier/jidt>. All code used to construct the network simulations are available at <https://zenodo.org/records/17169511>.

**Funding:** VP received funding from the Deutsche Forschungsgemeinschaft (DFG) in

## Abstract

When analysing high-dimensional time-series datasets, the inference of effective networks has proven to be a valuable modelling technique. This technique produces networks where each target node is associated with a set of source nodes that are capable of providing explanatory power for its dynamics. Multivariate Transfer Entropy (TE) has proven to be a popular and effective tool for inferring these networks. Recently, a continuous-time estimator of TE for event-based data such as spike trains has been developed which, in more efficiently representing event data in terms of inter-event intervals, is significantly more capable of measuring multivariate interactions. The new estimator thus presents an opportunity to more effectively use TE for the inference of effective networks from spike trains, and we demonstrate in this paper for the first time its efficacy at this task. Using data generated from models of spiking neurons—for which the ground-truth connectivity is known—we demonstrate the accuracy of this approach in various dynamical regimes. We further show that it exhibits far superior inference performance to a pairwise TE-based approach as well as a recently-proposed convolutional neural network approach. Moreover, comparison with Generalised Linear Models (GLMs), which are commonly applied to spike-train data, showed clear benefits, particularly in cases of high synchrony. Finally, we demonstrate its utility in revealing the patterns by which effective connections develop from recordings of developing neural cell cultures.

## Author summary

Network inference is a useful technique for the analysis of high-dimensional time series. It allows us to reduce the complexity of the raw data to a network summarising the

two grants: SFB1528 and EXC 2067/1-390729940. MW received funding from the Deutsche Forschungsgemeinschaft (DFG), grant number: SFB1528. JTL received funding from the Australian Research Council (ARC), grants: DE160100630 and DP240101295. None of the funders played any role in the study design, data collection and analysis, decision to publish or preparation of the manuscript.

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

relationships between the different elements of the time series. Effective networks in neuroscience perform this task by finding the smallest set of source elements which provide maximum explanatory power for the activity of each target node. A directed connection is then drawn from each parent to each target. Transfer Entropy (TE) is a popular tool for inferring these networks. However, the use of TE to infer effective networks from spike train data had previously been limited by the lack of a good estimator of TE for this class of data. This paper demonstrates that a recently-proposed continuous-time estimator of TE on spike trains, when combined with an existing greedy algorithm, is a powerful tool for inferring effective networks.

## 1. Introduction

For many of the complex systems that scientists are most interested in, our ability to record high-fidelity data from the numerous components of these systems is improving rapidly. For instance, the number of biological neurons that can be simultaneously recorded from is increasing exponentially, with a doubling rate of around six to seven years [1,2], while spatial resolution is increasing at a similar rate [3]. The process of drawing scientific insight from this flood of data is, however, not always straightforward [4].

The inference of effective network models [5] from high-dimensional time-series data has become a popular and productive technique for reducing the complexity of this class of data. Such datasets often consist of millions of (or far more) individual data points [6]. The inference of effective networks aims to produce a minimal system model from the data, by finding the smallest set of system source components capable of explaining the activity of each target component [7]. As such, it compresses the large number of data points down to a single directed network diagram describing the relationships between components of the system, thus facilitating the interrogation of the data at hand.

Effective network models should only incorporate sources in the model when they provide unique and/or synergistic information about the target with respect to the other sources in the model, but not included when the information they provide is merely redundant with other sources [7,8]. One use of such a minimal system model is that it can serve as a base upon which to examine these high-order effects across the sources and in the transfer entropy from them. As recently proposed by Stramaglia et al. [9], multiplets of sources can be searched for to identify unique, synergistic and redundant atoms. Other approaches have also been proposed for performing such information decompositions [10,11]. However, the focus of this work is on inferring the minimal set of sources for each target, as opposed to performing such a decomposition across the set.

In this work, we specifically focus on the inference of effective networks for time-series data of spiking neural activity, summarised by the timestamps of action potentials (spikes) at each node in the network. Much research has addressed network relationships using Transfer Entropy (TE) [12,13], a widely accepted measure of information transfer, suitable for mapping information flows in neural systems [14]. This has incorporated exploring [15–18] and applying TE-based methods to *in vitro* [19–25] and *in vivo* [26] recordings of spiking activity. All of this work has estimated TE in a pairwise fashion, producing what are often referred to as *directed functional* (as opposed to *effective*) networks [5].

The main contribution of this work is to marry a recently-proposed continuous-time estimator of TE for spike trains [27] with an existing greedy effective network inference algorithm. The greedy approach iteratively adds sources to a parent set for each target in a greedy

fashion, choosing the source with the highest conditional TE at each step until no candidate sources with (statistically significant) non-zero TE remain [7,8,28–30].

The main obstacle that has prevented the inference of effective networks from spike trains using TE has been the manner in which the traditional method of estimating TE on spike train data – which uses long sequences of time bins to represent temporal spiking patterns – causes a rapid increase in the dimensionality as we add conditioning processes [27].

The recent development of a continuous-time estimator of TE for event-based data [27] bypasses the dimensionality problems of the discrete-time estimator. Specifically, as it uses inter-spike intervals to efficiently represent the temporal spiking patterns, it is capable of capturing dependencies over relatively long periods of time (on the order of seconds [31]), with no loss of time precision. This makes the estimation of TE with sizeable conditioning sets feasible, supporting efficient inference of effective networks.

In this paper, we bring the greedy network inference algorithm and the continuous-time TE estimator together for the first time. We validate the efficacy of this combination on synthetic examples, where the underlying causal network is recovered by the model. We further compare its efficacy against generalised linear models [32] and a convolutional neural network based approach [33], finding its performance to be highly competitive. We finally demonstrate its ability to uncover biological insight by inferring the effective networks of developing cell cultures of dissociated cortical rat neurons [34].

## 2. Results

In this section we apply the greedy TE-based effective network inference algorithm [7] in conjunction with the continuous-time TE estimator for event-based data [27]. Please see [Sect 4.1](#) for details of the operation of the greedy algorithm, along with a description of a few minor changes that were made for the application to event-based data. [Sect 4.4](#) summarises the TE estimation approach used.

The first three subsections of this section focus on the inference of effective networks from simulated time-series of systems of spiking neurons for which the ground truth connectivity is known. We must emphasize, however, that in general we do not expect the effective networks inferred by TE to align with the causal structure [7]. Whilst the effective networks always provide a useful model for interpreting the directed relationships in the system, it is only under certain specific conditions that we expect them to match the causal structure, most importantly full observability of the nodes involved in the dynamics, and under certain assumptions such as faithfulness and the causal Markov property [35,36]. Therefore, we evaluate the performance of the network inference scheme by comparing the inferred network to this ground truth under these idealised conditions, since this provides an important validation of the output of the inference when this match can be expected. In order to measure the accuracy of the inference scheme, we make use of the commonly-employed classification metrics of *recall* and *precision*. They are defined as:

$$\text{recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (1)$$

and

$$\text{precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}. \quad (2)$$

Here, TP is the number of true positives, FP is the number of false positives and FN is the number of false negatives. In the context of network inference, recall can be interpreted as the proportion of true connections that were predicted by the algorithm. Precision, on the other hand, is the proportion of predicted edges that are true edges.

The final subsection of the results focuses on the application of the estimator to the spike times from recordings of cultures of dissociated rat cortical neurons. This provides a demonstration of the utility of this approach for extracting insights from biological data.

## 2.1. Inference at varying levels of synchrony

We constructed networks of Leaky-Integrate-and-Fire (LIF) [37] neurons with alpha synapses [38]. These networks were composed of 30 excitatory and 20 inhibitory neurons. Each neuron had exactly three excitatory and two inhibitory sources, where these sources were selected randomly from the respective sets. By varying the ratio of inhibitory to excitatory connection strength  $g$ , we could vary the level of synchrony within the networks. We ran simulations for three different levels of synchrony, which we refer to as “low” ( $g = 3$ ), “medium” ( $g = 1.5$ ) and “high” ( $g = 1$ ). Varying the relative strength of inhibitory connections across the excitation-inhibition balance threshold is a known method for adjusting the degree of synchrony in these networks [39]. Please see [4Sect .6](#) for full details on these network models. It is also worth noting that the level of synchrony present in these networks (even at “high” synchrony) is far lower than in the biological data examined in [Sect 2.4](#).

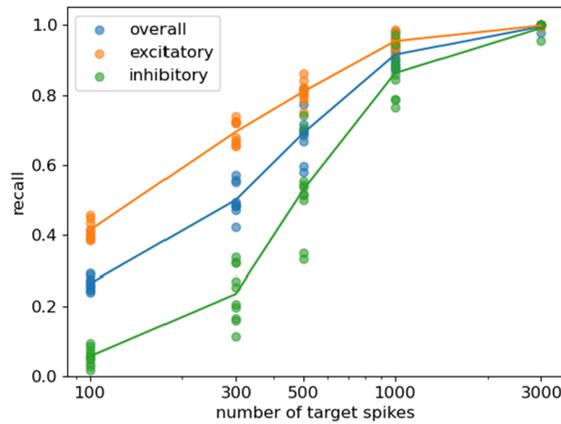
The combination of the greedy inference algorithm and the continuous-time estimator was applied to these networks, and the resulting inferred networks were compared against the ground truth for varying numbers of target spikes available to the estimator: 100, 300, 500, 1000, and 3000 (extra runs at 5000 target spikes were included for the high synchrony network, as the recall rose more slowly in this case). 10 independent simulations of the network model were run for each level of synchrony and the algorithm was applied to each run for each number of target spikes, although it was only applied to the first 5 simulations at 3000 spikes and the first 3 simulations at 5000 spikes due to the high computational requirements.

The precision and recall of the resulting inferences was calculated and is plotted in [Fig 1](#) for the different dynamical regimes and numbers of target spikes.

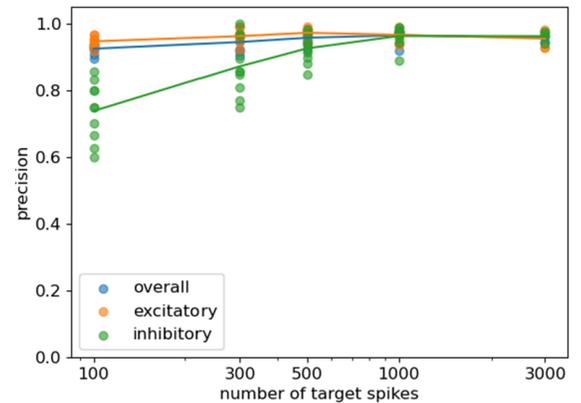
In the results shown in [Fig 1](#), we see that the algorithm exhibits high precision for all combinations of dynamical regime and number of target spikes. The precision only drops below 0.9 where the recall is very low, i.e. when few links are inferred. This demonstrates the high confidence with which it predicts links — a very low proportion of the predicted links turn out to be false positives — which can also be seen as a conservative approach.

In these plots, the recall begins low, but rises rapidly with the increase in the number of target spikes available. In the case of low network synchrony ([Fig 1A](#)), we observe that, by 3000 target spikes, the recall has risen to nearly one. As the precision is also nearly one at this number of target spikes, the networks are being inferred nearly perfectly. Taken in conjunction with the apparent trends towards converging on perfect inference as the number of spikes increases for the other regimes, this suggests evidence for validation that the approach provides a consistent inference of the underlying network under the aforementioned idealised conditions. Note that we do not expect the accuracy of our approach (in terms of providing a network model to explain the observed dynamics) to decrease outside of these conditions. However, outside of cases of full observability (and other assumptions outlined in the Introduction), the effective network in explaining the observed dynamics is not guaranteed to nor cannot be expected to align with the causal network.

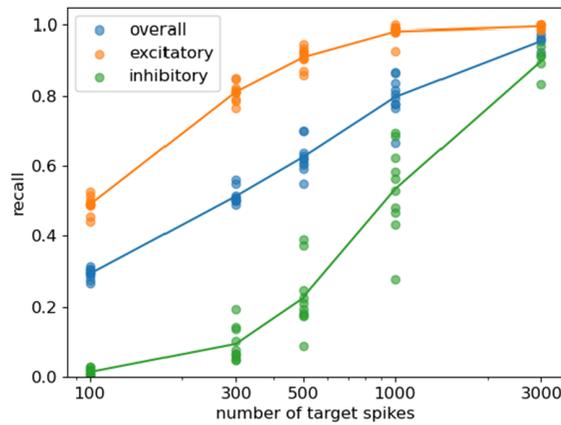
By comparing [Fig 1A](#), [1C](#) and [1E](#), we can observe that the achieved overall recall drops as the level of synchrony is increased. This is entirely driven by a drop in the recall on inhibitory connections. In fact, we observe a small increase in the recall on excitatory connections. This drop in recall is due to the increased complexity in the nature of the statistical relationship between the activity of a given target and an inhibitory source in the case of high synchrony.



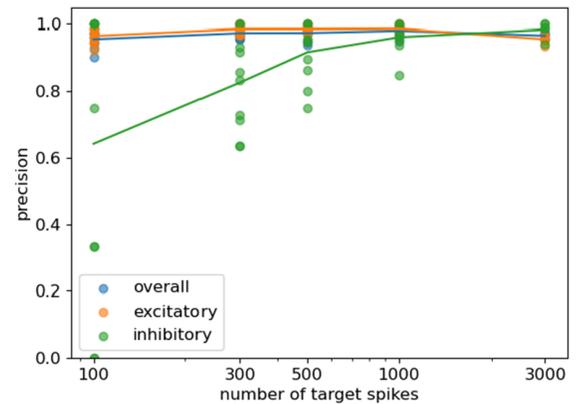
(A) Lower synchrony network, recall



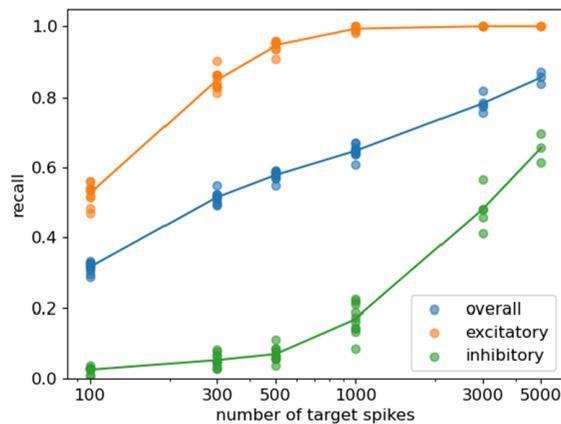
(B) Lower synchrony network, precision



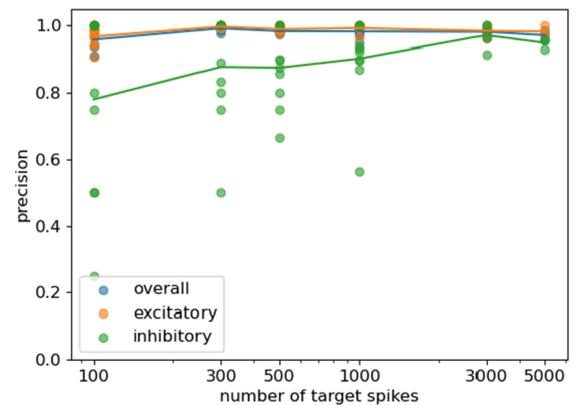
(C) Medium synchrony network, recall



(D) Medium synchrony network, precision



(E) High synchrony network, recall



(F) High synchrony network, precision

**Fig 1. Plots showing the resulting precision and recall from running the network-inference scheme on networks of LIF neurons composed of 30 excitatory neurons and 20 inhibitory neurons.** Each neuron has exactly 3 excitatory and 2 inhibitory inputs, chosen randomly. The ratio of the inhibitory to excitatory connection strength was varied in order to change the level of synchrony in the network. Plots are shown for three different synchrony levels. Each plot contains points for each experiment the precision and recall for the inhibitory and excitatory sources separately as well as for their overall weighted average. The lines pass through the means of these points.

<https://doi.org/10.1371/journal.pcbi.1013500.g001>

When the populations are highly synchronous, all of the cells spike close together, so the firing of an inhibitory source can become positively correlated with the firing of its target, when considering the purely pairwise relationship between the source and target. Note that the genuine relationship that we are seeking to infer here is the negative correlation in firing resulting from inhibition. It is only when conditioned on the target's excitatory sources (which becomes possible with more target spikes observed), for any given firing pattern of the excitatory sources, the firing of the inhibitory source is associated with a decrease in the probability of the firing of the target for that given pattern, allowing the inhibitory source to be identified. Crucially the precision remains high despite the spurious pairwise correlations that appear in the highly synchronous regime, due to the conditioning in the multivariate approach removing such redundant sources being included in the inference. Fig 3F shows an ROC curve for the use of a purely pairwise approach on this same high-synchrony example for 1000 target spikes, which will be discussed in Sect 2.3. We see that it cannot achieve a high true positive rate without a substantial increase in the false positive rate (and thus a decrease in the precision). This highlights the necessity of using a full multivariate approach when dealing with highly-synchronous neural populations.

In order to compare the performance of the proposed scheme with an existing network inference approach, we ran the recently-proposed CoCONNECT [33] algorithm on the spike times from the simulations. This approach makes use of pre-trained convolutional neural networks and has been demonstrated to be competitive when compared with other existing network inference algorithms for spike trains. In order to perform this inference, we made use of the associated web-app provided by the authors [40]. The resulting precision and recall plots are shown in S1 Fig. We consistently see that, for any given combination of number of target spikes and dynamical regime, the proposed approach is able to achieve both higher precision and higher recall. In particular, the precision of the results of the CoCONNECT inference is particularly low, being largely in the region 0.1-0.2.

We also compared the performance of the proposed approach with a Generalised Linear Model (GLM), which is a popular approach for modelling spiking neural data [32,41,42], including for inferring connectivity [43,44]. We closely followed previous work [43] which demonstrated the use of these models for connectivity inference, with a few minor differences, as specified in Sect 4.8. The resulting precision and recall plots are shown in S2 Fig. The GLM approach exhibits markedly lower precision than the proposed TE-based approach, except for precision on inhibitory sources for very low numbers of available spikes in the target spike trains, where our more conservative approach is inferring very few links. Moreover, in the case of high synchrony, the precision of the proposed approach is far superior, whereas the very low precision of the GLM approach results in almost half of all possible connections being inferred as present, rendering those inferred network models far less useful. For low numbers of target spikes, the GLM approach is able to achieve better recall. However, this is always at a cost of significantly lower precision, and any advantage in recall for the GLM approach disappears as the number of target spikes is increased above 1000 (except in the high-synchrony case). Moreover, these results were achieved by carefully tuning certain parameters of the GLM model setup in order to maximise performance on these specific examples. Importantly, this parameter tuning was done using the known ground-truth. In real-world inference applications we do not have access to such a ground truth. Using sub-optimal parameters can have a dramatic effect on the performance of this approach. S3 Fig shows how the precision deteriorates even further when a larger bin size  $\Delta t$  is used ((40 ms as opposed to 20 ms)). S4 Fig shows a similar deterioration in performance when we more closely mirror the original implementation by Song et al. [43] (i.e. by increasing the number of knots used in the spline basis functions to include all 16 knot locations in the original).

Interestingly, the better recall that our approach shows for excitatory versus inhibitory sources is reversed in the GLM approach.

## 2.2. Inference at varying levels of stimulus regularity

In [Sect 2.1](#) neurons were always provided with an independent Poisson stimulus. However, we can vary the properties of this stimulus in order to mimic different plausible inference scenarios. For instance, simulations constructed with a fully regular stimulation of the neurons provide us with an example of dynamics with no hidden sources of variability. That is, it is possible to perfectly predict the dynamics of the neuron based on its past and the past of those neurons connecting to it. Moreover, the system is completely deterministic. As we move towards the semi-regular and fully random Poisson stimuli, we are modelling increasing amounts of hidden activity or noise within the system, which prevents perfect predictability of the future state of the units. Full determinism can make the inference of networks from time series more difficult [35]. Conversely, very large amounts of noise can make it challenging to detect a comparatively weak relationship between two nodes. As such, it is important that both potentially problematic ends of this spectrum are tested.

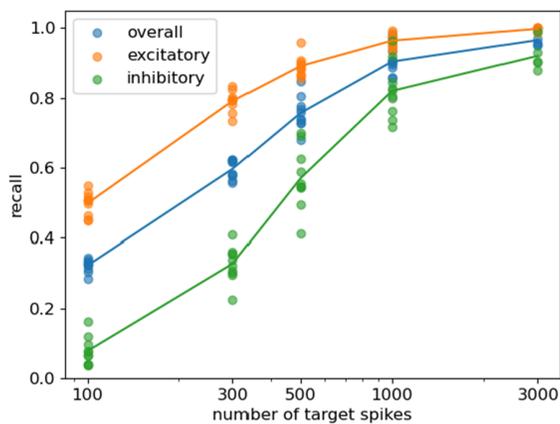
We constructed model networks as in [Sect 2.1](#), however, this was done at a single level of ratio of inhibitory to excitatory connection strength. The same ratio ( $g = 3$ ) used to produce the low synchrony runs was used. Instead, we varied the nature of the stimulus provided to each neuron, providing them with a regular, semi-regular or Poisson (fully random) stimulus. The regular stimulus was composed of spike times placed at a fixed interval. There was slight variation in this interval between neurons, in order to prevent the network settling into a simple, fixed, pattern. The semi-regular stimulus was similar to the regular stimulus, but with the addition of a small amount of Gaussian noise. Note that a few other minor simulation parameters had to be changed from the simulations used in [Sect 2.1](#) in order to ensure numerical stability. See [Sect 4.6](#) for a full specification of these differences.

[Fig 2](#) shows plots of precision and recall, for different numbers of observed spikes in the target, for these different levels of stimulus regularity. We observe that the recall increases slightly when moving from the regular to the semi-regular stimulus, but then drops as we move to the Poisson stimulus. By contrast, the precision exhibited a slight increase with increasing irregularity. This is likely because the increasing irregularity reduces the correlations between true and false sources, thereby decreasing the likelihood of false positives. In all three cases, good performance is achieved at 3000 target spikes, with overall precision and recall both being around 0.9.

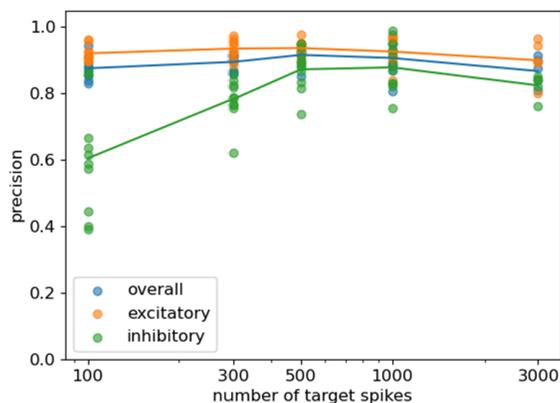
These results demonstrate that the proposed combination of estimator and network inference scheme is capable of successful inference at various levels of determinism or unobserved noise.

## 2.3. Comparing the greedy algorithm to pairwise inference

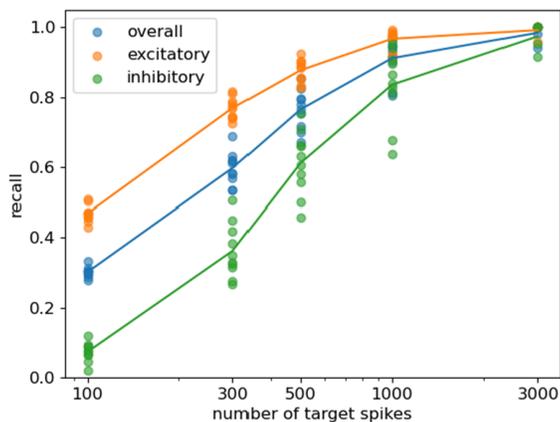
Recent work [45] has highlighted the improvements that can be gained when performing network inference using TE in its full multivariate sense, via the greedy algorithm, as opposed to the simpler pairwise approach. The pairwise approach operates by only checking for a statistically significant non-zero TE value between each source-target pair, without taking into account the other processes in the system. It is generally found that the full multivariate approach tends to exhibit much higher precision, as, among other reasons, it is able to distinguish true sources (which provide information about the target even when conditioned on all other system components) from spurious sources, which are merely correlated with the true



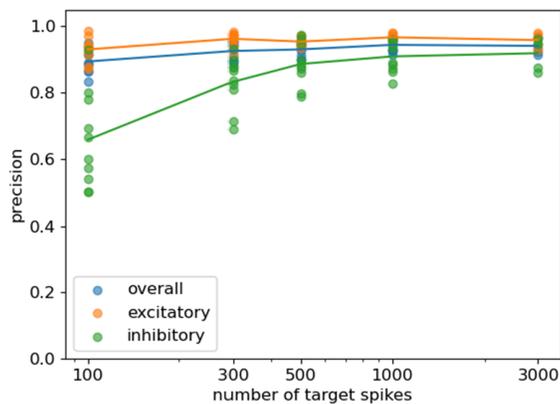
(A) Regular stimulus, recall



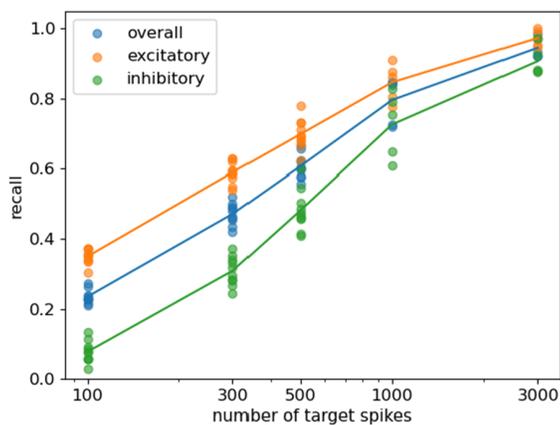
(B) Regular stimulus, precision



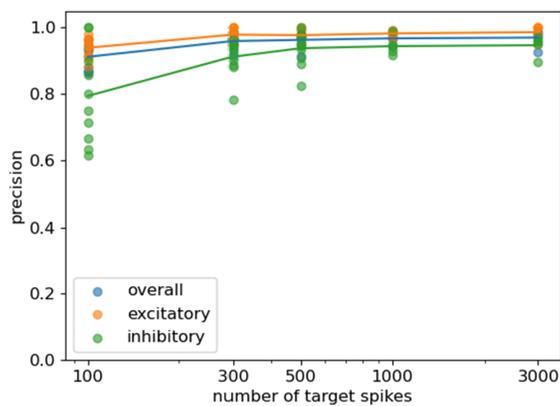
(C) Semi-regular stimulus, recall



(D) Semi-regular stimulus, precision



(E) Poisson stimulus, recall



(F) Poisson stimulus, precision

**Fig 2. Plots showing the resulting precision and recall from running the network-inference scheme on networks of LIF neurons composed of 30 excitatory neurons and 20 inhibitory neurons. The regularity of the stimulus provided to each neuron was varied.**

<https://doi.org/10.1371/journal.pcbi.1013500.g002>

sources but provide no additional information about the target when conditioning on these true sources.

The previous work [45] which compared multivariate TE and the greedy algorithm to the pairwise approach did so for standard time series of continuously varying signals sampled at a fixed interval. In this section, we verify that similar results hold when analysing event-based data using the continuous-time estimator of TE. Moreover, the analysis in this section will confirm the benefit of using the multivariate greedy approach, over the pairwise approach, when inferring networks from spike trains using TE.

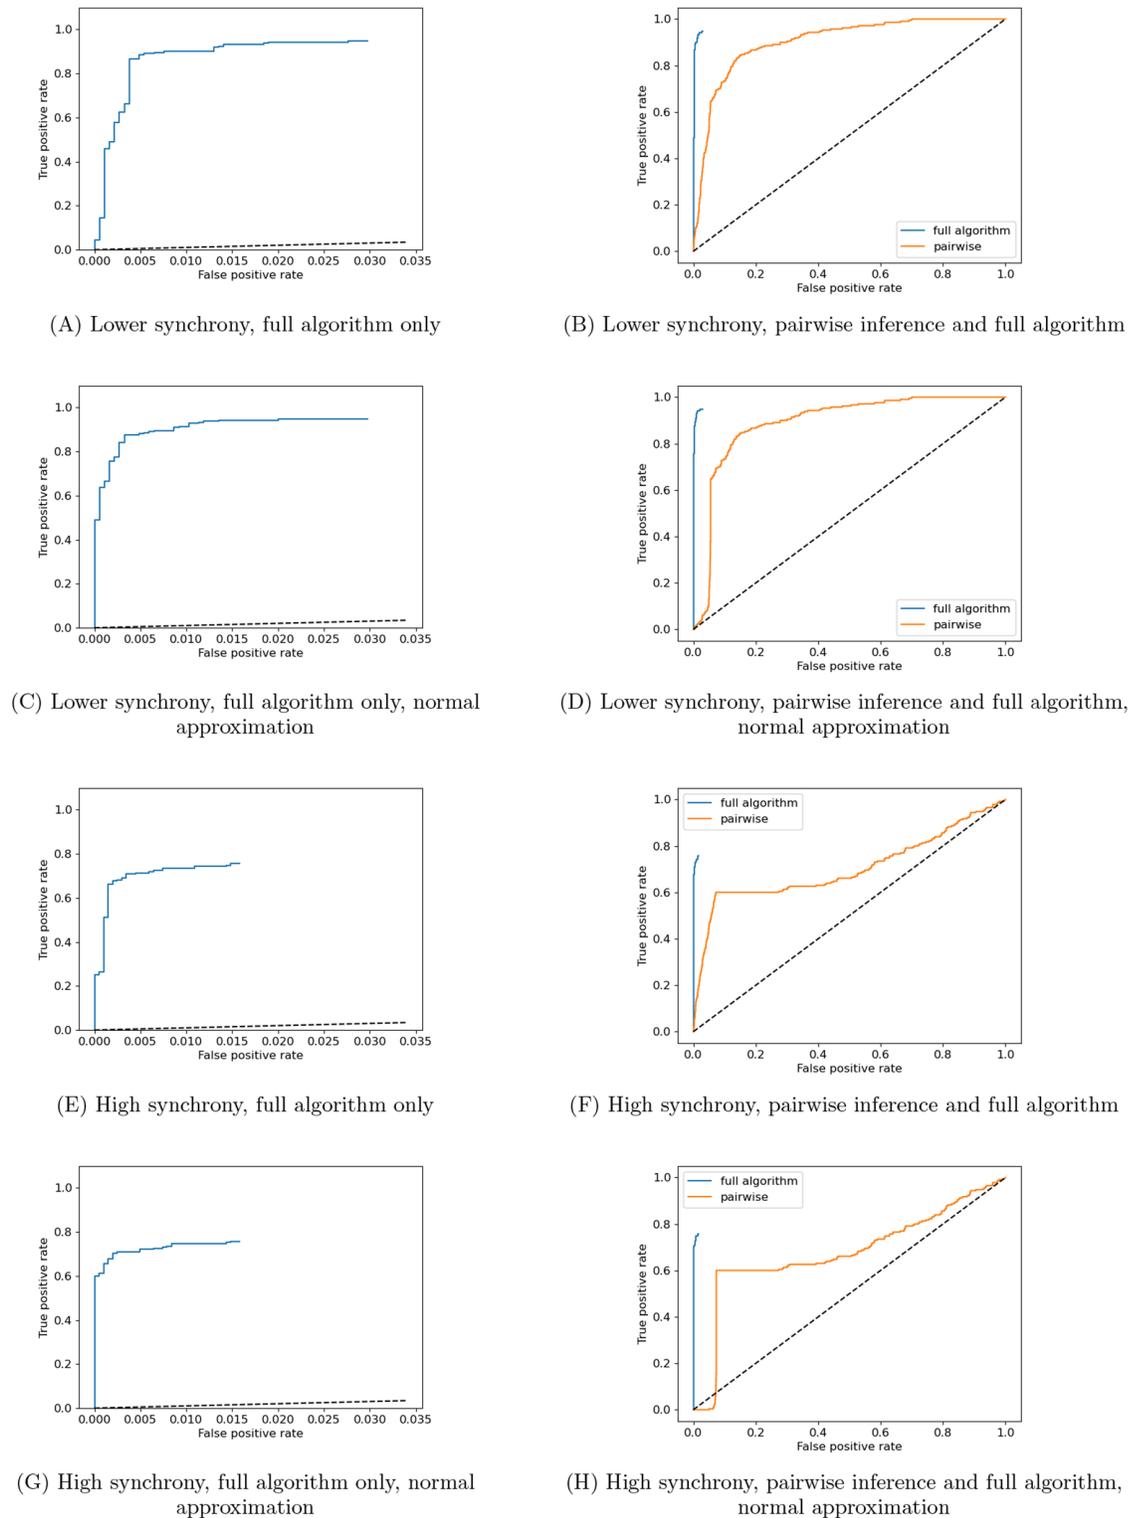
We make use of the higher and lower synchrony simulations presented in Sect 2.1, with 500 target spikes available. We applied a simple pairwise network-inference scheme to the resulting spike times from these simulations, which simply tested for statistically significant non-zero TE between each source-target pair. The resulting ROC curves are shown in Fig 3. These ROC curves are created by sweeping through the  $\alpha$  cutoff values (the threshold below which the  $p$  value must be for a link to be inferred) between 0 and 1 and recording the false-positive and true-positive rate observed at each  $\alpha$  value. Calculating the  $p$  values for the statistical significance tests of non-zero TE was done by both counting the proportion of empirical surrogates (see Sect 4.2 for a discussion of how these empirical surrogates are created) larger than the measured TE (Fig 3F and 3B), as well as via a fitting a normal distribution to the surrogate values (Fig 3H and 3D).

The purpose of also performing this normal approximation is that it allowed us to use much lower  $p$  value thresholds for a given number of surrogate calculations than is possible when evaluating  $p$  values by counting proportions of empirical surrogates, making it possible to efficiently gain more resolution on the far left of the ROC curves.

We also ran the full greedy algorithm with different  $\alpha$  cutoff values between 0 and 0.75, providing it with 1000 target spikes (the pairwise approach made use of the same number of target spikes). The final pruning step (see Sect 4.1) was, however, excluded, as this allowed for greater computational efficiency in a single bulk run. The resulting ROC curves are also plotted in Fig 3. Note that these ROC curves will not reach the point where the true positive rate and false positive rate both equal one, as we are unable to inspect  $p$  values larger than 0.75.

By inspecting the ROC curves in Fig 3A through 3D we can compare the performance of the two approaches on the networks with lower levels of correlation. The full multivariate approach is seen to very quickly arrive at a true positive rate of above 0.9, for very few false positives, which again underlines the effectiveness of this approach. In contrast, the true positive rate of the pairwise approach rises much more slowly; in other words it costs a substantially larger number of false positives to achieve the same true positive rate. Visually we see this in that the ROC curve of the multivariate approach is markedly above that of the pairwise approach.

We see an even starker difference in the performance of these two approaches when we look at the results of inference run on the networks exhibiting higher synchrony (Fig 3E through 3H). Here, we see that the true positive rate of the pairwise approach rises much more slowly than the multivariate approach as before, but its performance also saturates around a true positive rate of around 0.6 before false positives begin to strongly dominate further inference. In this regime the entire population has a tendency to be active together and also remain quiescent together. As the activity of all neurons are therefore correlated, the pairwise approach is unable to delineate which particular neurons are driving the activity of others. The multivariate approach is, by contrast, more robust to the higher synchrony and still able to achieve very high true positive rates at incredibly low false positive rates.



**Fig 3. ROC curves for both the presented network inference technique, as well as performing the inference using the continuous-time estimator in a simple pairwise fashion.** Plots are shown for the higher synchrony and lower synchrony examples from Fig 1, with 1000 target spikes available to the algorithms. We also show plots for both the presented surrogate testing method as well as when using a normal approximation fitted to the surrogate population in order to estimate the  $p$  value.

<https://doi.org/10.1371/journal.pcbi.1013500.g003>

These results demonstrate the substantial advantages in using multivariate TE estimation in conjunction with the greedy algorithm as opposed to a pairwise (functional network) approach.

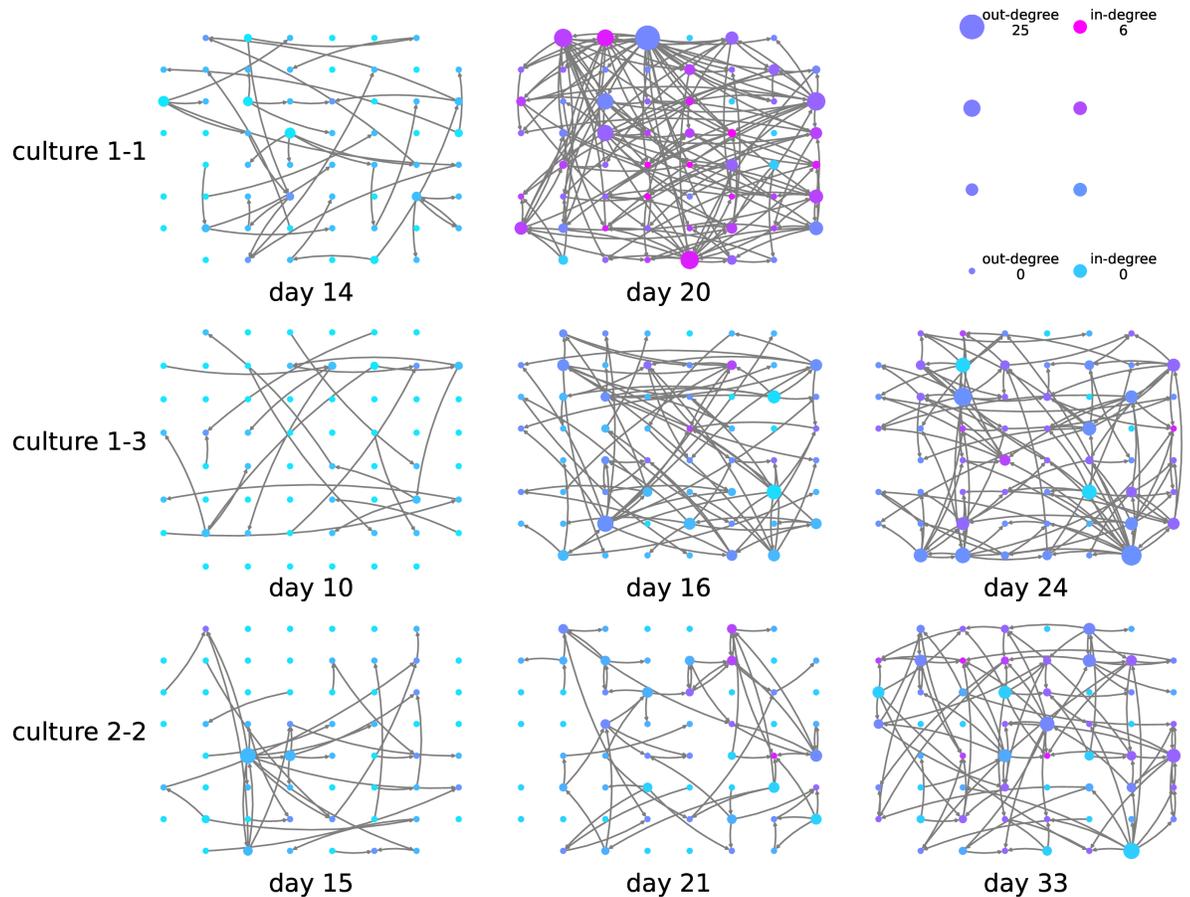
## 2.4. Inference of the effective networks of developing cell cultures

In order to demonstrate the utility of the application of this network inference scheme to biological data, we inferred the effective networks at various stages of development of cultures of dissociated cortical rat neurons. These recordings are part of a freely-available public dataset [34,46]. See Sect 4.7 for a summary of the nature of this dataset as well as details on how the network inference scheme was applied to it. In brief, cultures were allowed to develop over periods of around 30 days. On certain days, overnight recordings were performed. As these long overnight recordings contain sufficient numbers of spikes for effective application of information-theoretic estimators, they are eminently suitable for the application of our network inference approach. No spike sorting was performed, and so the networks are being inferred between the time series of the recording electrodes. This allows the nodes in the network to remain identifiable across different stages in development.

The results of applying the greedy algorithm along with the continuous-time estimator are displayed in Fig 4. Note that the first recording days of each culture are not included in the figure as hardly any links (less than 10) were inferred in any of these recordings. We observe that effective networks with a rich and complex structure emerge, beginning to appear around the tenth day in vitro or so and quickly becoming more dense. This path of development correlates with the authors' previous investigation [31] of these recordings, with simpler directed functional networks inferred using pairwise transfer entropy (via the same underlying estimator). The density of the networks inferred by the multivariate algorithm are lower than for the directed functional networks in [31], with the total number of edges in the last recording days declining from around 1000 to 2000 to around 100 to 200. This is as expected because the strongest action of the multivariate algorithm is to remove redundant sources [45].

Despite the difference in density, the effective network structures retain some of the interesting features observed for the directed functional networks in [31], such as containing pronounced inward and outward hubs (that is, nodes with particularly high in-degree or out-degree), and various features of the networks being locked in early in development. Specifically, in the directed functional networks in [31] characteristics such as the total inward or outward information flow for a given node exhibited high correlation between early and late days of development. Here, Fig 5 shows scatter plots of the out-degrees of the inferred effective networks on earlier and later days of development. We see in these plots that, as with the functional networks, in all cases, there is a positive correlation between the out-degree on the earlier and later days of development. Moreover, there are no statistically significant negative correlations. The positive correlation for the out-degree across the last two recording days is statistically significant for each culture, with some of these relationships, such as for culture 1-3, being particularly strong. Fisher's method [47] for combining  $p$ -values was applied to all the  $p$ -values of the correlation coefficients (a  $p$ -value of 1 was used in the two cases of negative correlation). This results in a meta-test of the hypothesis that there is a positive correlation between the out-degree on the earlier and later days of development across the cultures. The null hypothesis of no or negative correlation was rejected at the 0.01 significance level.

Fig 6 shows similar plots, but for the in-degrees of the nodes. Again, we see a positive correlation between the in-degree on earlier and later days in every case. Fisher's method for combining  $p$ -values was applied as above for the out-degree. The null hypothesis of no or negative correlation was rejected at the 0.01 significance level.

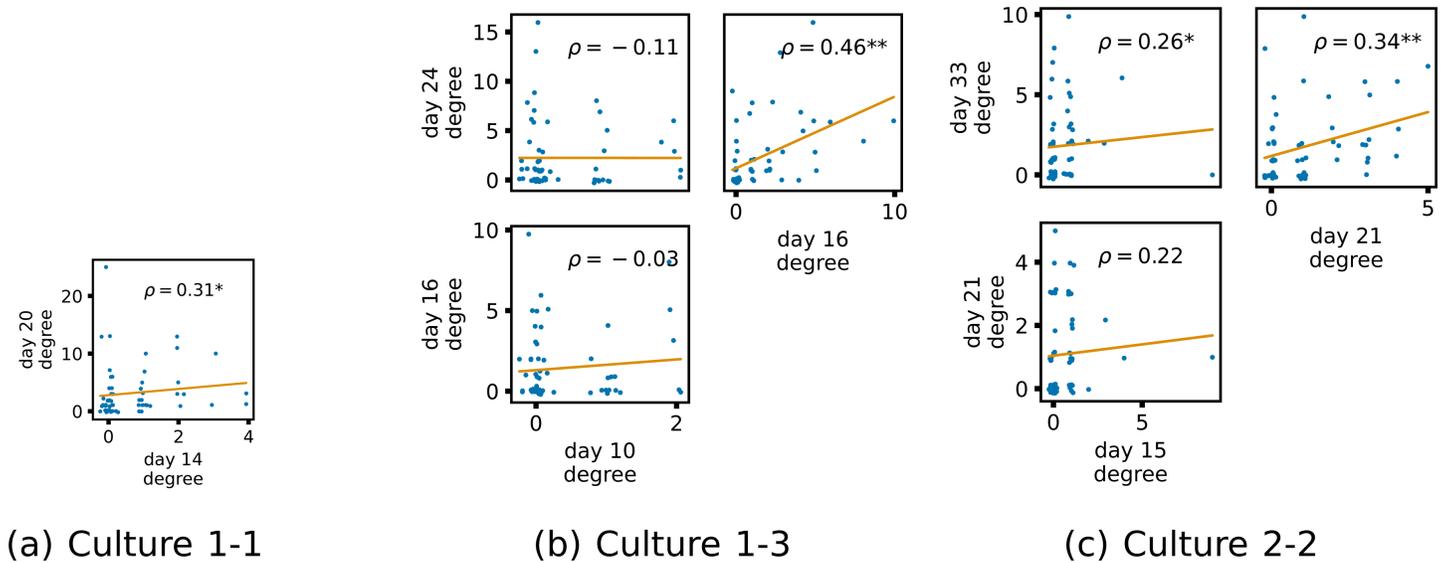


**Fig 4. Effective networks inferred using the presented approach between electrodes in developing cultures of dissociated cortical rat neurons.** Each node in the network visualisations is placed in the same relative spatial location that the corresponding electrode occupied in the recording apparatus. Networks were inferred at different stages of development (days *in vitro*). The recordings used are part of an openly-available public dataset [34,46]. Node colour and size is proportional to the in and out degrees (see the legend in the top right). The spacing between the electrodes is 200  $\mu\text{m}$  centre to centre [34,46].

<https://doi.org/10.1371/journal.pcbi.1013500.g004>

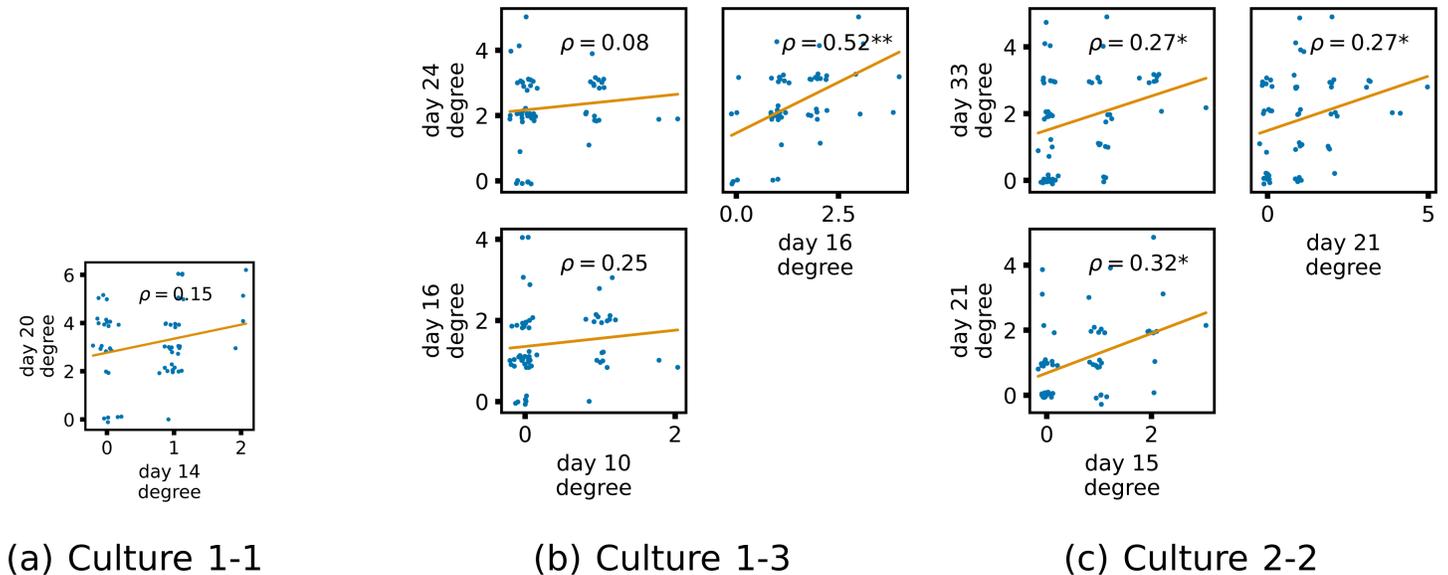
These results indicate that features of the effective networks, representing the multivariate information flows here, are being locked in early in development. This is particularly interesting since there are large structural differences between the networks inferred in these two studies. The effective networks inferred here are substantially sparser than the directed functional networks in [31]. This suggests that the lock-in effect is deeply ingrained in the system.

Fig 7 displays the proportion of possible links that are inferred in the networks at various physical distances between the nodes. Fig 7A does so for the effective networks inferred in this work and Fig 7B does so for the functional networks studied in the authors' previous work. These plots show that the effective networks inferred on this dataset exhibit a clear preference towards links between nodes that are physically close together. This preference appears to become stronger with developmental time. By contrast, the functional networks do not exhibit this preference. This highlights a further difference between the inferred effective and functional networks, making the maintenance of the lock-in effect more interesting.



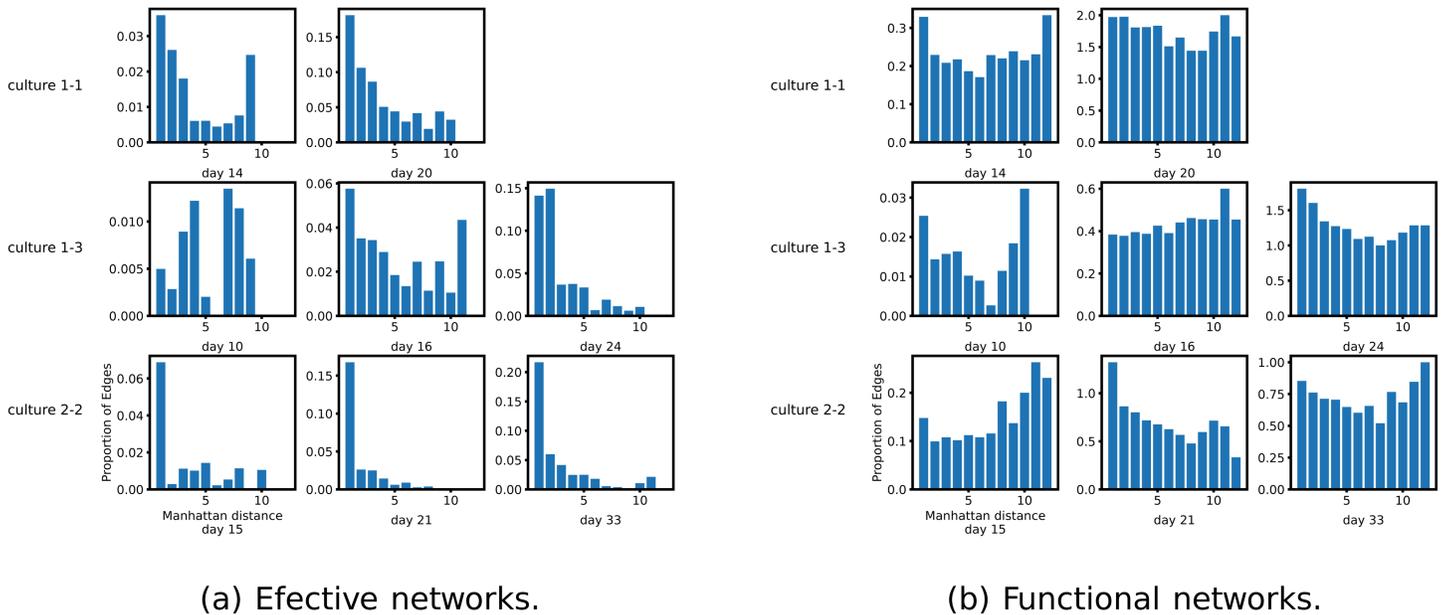
**Fig 5. Plots showing the relationship between the out-degree of a given node over different days of development.** Each group of plots shows scatter plots between all pairs of days for each culture analysed. Specifically, in each scatter plot, the x value of a given point is the out-degree of the associated node on an earlier day and the y value of that same point is the out-degree of the same node but on a later day. The days in question are shown on the bottom and sides of the grids of scatter plots. A small amount of Gaussian jitter ( $\sigma = 0.1$ ) is added to the points to aid the visualisation of repeated values. The orange line shows the ordinary least squares regression. The Spearman correlation ( $\rho$ ) between the out-degrees on the two days is displayed in each plot. Values of  $\rho$  significant at the 0.05 level are designated with an asterisk and those significant at the 0.01 level are designated with a double asterisk.

<https://doi.org/10.1371/journal.pcbi.1013500.g005>



**Fig 6. Plots showing the relationship between the in-degree of a given node over different days of development.** Each group of plots shows scatter plots between all pairs of days for each culture analysed. Specifically, in each scatter plot, the x value of a given point is the in-degree of the associated node on an earlier day and the y value of that same point is the in-degree of the same node but on a later day. The days in question are shown on the bottom and sides of the grids of scatter plots. A small amount of Gaussian jitter ( $\sigma = 0.1$ ) is added to the points to aid the visualisation of repeated values. The orange line shows the ordinary least squares regression. The Spearman correlation ( $\rho$ ) between the in-degrees on the two days is displayed in each plot. Values of  $\rho$  significant at the 0.05 level are designated with an asterisk and those significant at the 0.01 level are designated with a double asterisk.

<https://doi.org/10.1371/journal.pcbi.1013500.g006>



**Fig 7. Bar plots showing the proportion of possible edges that were inferred at different inter-node distances.** (A) shows the proportions for the networks inferred using the presented greedy algorithm and, whose diagrams are displayed in Fig 4. (B) shows the same proportions for the functional networks inferred on the same data in the authors' recent work [31]. The inference of such functional networks only considers pairwise relationships. The distances on the  $x$  axis are the Manhattan (cityblock) distances between electrodes. It is clear from the plots that, on this dataset, the effective network inference algorithm has a greater propensity to infer short distance links.

<https://doi.org/10.1371/journal.pcbi.1013500.g007>

### 3. Discussion

In this work, we have validated the efficacy of the combination of an existing greedy multivariate TE-based network inference algorithm with a recently-introduced continuous-time estimator of TE on event-based data. Although we have validated the approach on neural spiking data, it can be applied to any event-based data, such as the times of social-media posts or the times of stock-market trades.

As the inference of networks from the spike times of neurons is a common goal within neuroscience, we expect this particular task to be a core application of the presented approach. Indeed, there have been several previous studies which have proposed TE-based methods for inferring networks from the spike times of neurons and evaluated them against ground truth [15–18]. There have also been a number of studies which used TE to infer networks from *in vitro* [19–25] and *in vivo* [26] recordings of spiking activity. These networks were found to exhibit highly non-random structure [20], including rich-club topologies [19]. All of the aforementioned work has estimated the TE in a pairwise fashion, that is, without conditioning on other recorded processes. It has also performed estimation on discretised data. Networks inferred based on pairwise statistics are often referred to as *functional* (as opposed to *effective*) networks [5], since they are reporting pairwise relationships rather than a minimal multivariate directed model of the dynamics.

However, apart from a single very recent study [48], this previous work has always considered only the pairwise relationships between the activity on each node. As was demonstrated in Fig 3, even when using the highly-effective continuous-time TE estimator, this approach suffers substantial drawbacks. Perhaps most notably, when the entire population is highly correlated, the pairwise approach is unable to find a small subset of sources which can

collectively explain the activity of the target. Instead, it infers large numbers of sources due to the redundant information they hold with the true sources [45].

Here, by contrast, we have presented a multivariate approach to inferring *effective* networks from neural spike trains using TE. Unlike the pairwise approach, this strategy infers a set of parents for a target collectively rather than individually for each source. In doing so, the multivariate strategy considers the activity of other nodes within the network when determining the directed relationship between any two nodes. Specifically, as the multivariate strategy iteratively or greedily adds new candidate sources to the parent set for a target, it requires that each source provide statistically significant non-zero TE, when conditioning on all other current parents for the target. This is in contrast to the pairwise approach, which only requires a non-zero TE value between the source and target, without taking other processes into account. That iterative conditioning, along with a final associated pruning step, supports much more accurate inference because it eliminates redundant information from being spuriously attributed to other sources, and captures synergistic or collective interactions between multiple sources which jointly impact the target.

The term “accurate” here has specific meaning in the context in which we have evaluated the performance of the multivariate approach. Although we cannot and do not always expect the inferred effective networks to align with the underlying causal or structural network of the system being examined, under certain highly-specific idealised conditions (full observability etc., see Sect 1) we do indeed expect a minimal model explaining the dynamics of the variables (the effective network) to align with the causal structure in this way. As such, confirming such alignment – even though in highly-specific conditions – is an important validation of the performance of such an approach. Indeed, this validation is clear from our results, including in Figs 1 and 2, and by the improved performance of the multivariate approach compared to pairwise, using the same underlying TE estimator, in Fig 3. There, it was found that the multivariate approach was able to achieve a comparable true positive rate as the pairwise approach with a much lower corresponding false positive rate.

Maintaining a low false positive rate (or, equivalently, a high precision) is of utmost importance for network inference in a neuroscientific context. Zalesky states that “False positives are at least twice as detrimental as false negatives” [49]. As demonstrated in all of our results, the presented approach errs on the conservative side and consistently maintains high precision (a low false-positive rate) even in the challenging cases where the activity on the nodes is highly correlated (Fig 1F). This high precision is a result not only of the multivariate strategy, but also the local permutation surrogate generation method used for the significance testing of the individual TE estimates being nonzero. This method was developed in tandem with the recently-developed continuous-time TE estimator [27] that is used here and was demonstrated to have substantially lower false-positive rates when compared with using traditionally-used approaches for surrogate generation in conjunction with this estimator.

Our results also compared the proposed approach with two existing approaches for the inference of connectivity from spiking data (CoNNECT in S1 Fig and GLMs in S2 Fig), finding that it exhibited far superior precision in most cases, particularly so under conditions of high synchrony. The performance of the GLM approach degraded further when not using optimal parameters tuned using the ground truth (S3 and S4 Figs). This is particularly important when we reflect on the goal of effective network inference being to provide a “minimal model” that can explain the dynamics given the time-series data. Moreover, our approach, being grounded in information theory, has the unique advantage of inferring networks that are readily interpretable in terms of the fundamental computational operations of information storage, transfer and modification [7,50–52]. And finally, as these measures can be

estimated non-parametrically [53], they are not dependent on model assumptions and can capture any form of non-linear relationship.

Not only does this work present the first validation of TE for multivariate effective network inference on spike-train data, but it presents the first validation study of using the recently-developed continuous-time estimator of TE on event-based data such as spike trains in the context of network inference [27]. This estimator has been demonstrated to have many substantial advantages over the traditional discrete-time approach. These include consistency, lower bias and faster convergence. Of particular relevance to network inference, by representing the temporal pattern or history embeddings of processes using inter-event intervals, it is able to represent histories of substantial length using few dimensions and without any loss of time precision. This efficient use of dimensions facilitates building conditioning sets of significant size.

These various benefits culminate in a technique that is highly effective at the inference of networks from event times. We have demonstrated its strong performance, with high precision, in dynamical regimes ranging from low to high network synchrony (Fig 1) as well as with varying levels of unobserved noise sources in the system (Fig 2). Furthermore, high quality inferences were made with relatively low numbers of target spikes. In some instances (eg: Fig 1A and 1B), near perfect reconstruction was achieved with only 3000 events per target and a reasonable reconstruction was achieved with an order of magnitude fewer events. We note that all of these advantages do come at a cost of increased run-time in comparison to the pairwise TE as a measure of directed function connectivity although this increase is by a constant factor (see Sect 4.1), and moreover the run-time performance remained faster than the GLM (which required many iterations for parameter selection).

These results all point to the strong potential for deploying this methodology in the inference of networks from recordings of the spike times of biological neurons. Tantalising hints of the results that might be expected were provided in Sect 2.4. Further such applications remain a focus of future work. It is of particular interest to note that, in these effective networks, there was a strong preference towards inferring edges between nodes that are spatially close together, especially when compared with the functional network approach. This is likely due to the fact that effective networks are known to conform closer to the underlying structural networks than those inferred using pairwise, functional, methods [45]. This fits with the known fact that, in neural cell cultures, closer neurons are more likely to be connected [54]. However, despite this change in the topology of the networks, it is worth noting that the lock-in of information flows early in development, which was previously observed in functional networks [31], remained in these effective networks.

## 4. Methods

### 4.1. Greedy algorithm

The greedy network inference algorithm used here was proposed in a range of papers [8,28–30], as summarised and studied in depth by Novelli et al. [7], for traditional time series (a continuous-valued signal sampled at regular time intervals). We describe it here for completeness, and also to highlight some small changes that we made to adapt to the context of event-based data. The most significant change made is that only one inter-spike interval per source is considered as a candidate in each selection round, and sequentially, from the most recent. This is in contrast to the original algorithm, where several lags from each source could be considered in the same selection round, and no ordering was imposed on the addition of these lagged samples.

The greedy algorithm is specified in [Algorithm 1](#). We walk through its operation here, with reference to the line numbers in [Algorithm 1](#).

We iterate over each process  $R_i$  in the set of processes  $\mathcal{R}$  (line 1). These processes are the raw timestamps of events (spikes). Each process is being considered as a target, for which the sources need to be inferred. It is worth noting that the computations performed for each target are considered completely independent of one another. As such, it is easy to parallelise this algorithm across the different targets. Indeed, such parallelisation was performed for the experiments presented in this paper.

We initialise a data structure to keep track of the last interval added to the conditioning set for each source and the target itself (line 2). This algorithm makes the assumption that more recent inter-event intervals from a given source (or the target itself) always have more influence over the target than inter-event intervals further in the past. Based on this assumption, inter-event intervals for a given source are only considered as candidates once more recent intervals for that source have been added to the conditioning set. As above, this is distinct from the operation of the algorithm for traditional time-series.

Before considering candidate sources, we determine the number of target history intervals to condition on. We always include at least one (the most recent) such interval. We continue incrementing the total number of intervals that we are conditioning on until the next interval does not provide a statistically significant reduction in the uncertainty of the target (line 4). This reduction in uncertainty is measured by the conditional Active Information Storage (AIS) [50], which is the mutual information between the oldest target history interval being considered and the current state of the target, conditioned on the more recent target history intervals. Note that this quantity can be easily estimated using the continuous-time TE estimator by simply considering this oldest target interval as a source interval itself. The conditional active information storage is estimated on the original spike train and on  $N_{\text{surrogates}}$  surrogate processes (lines 5 and 6), constructed using the local permutation method described in [Sect 4.5](#). The  $p$  value for the significance test is then the proportion of AIS estimates on the surrogates which are larger than the AIS estimate on the original process. If  $p < p_{\text{cutoff}}$ , then the null hypothesis of zero additional contribution to the AIS of the oldest interval is rejected, and the number of target intervals being added is incremented.

Returning to the candidate sources then, we continuously iterate over rounds of parent selection for the target until the candidate source interval with the highest TE in a given round is no longer statistically significant (line 13). For each of these selection rounds, we iterate over every process other than the target under consideration (line 16). For each such candidate source, we estimate the TE between the most recent inter-event interval of that source that has not yet been added to the conditioning set and the target, conditioned on all intervals of all sources already added to the conditioning set (line 17). We then estimate the TE for  $N_{\text{surrogates}}$  surrogate processes between the same source and target and conditioned on the same conditioning set (line 18) (this is performed for every TE calculation in order to support the maximum statistic test in the next step; see [Sect 4.2](#)). We also bias-correct the original and surrogate TE estimates by subtracting the mean value of the surrogates from each estimate (lines 19 and 20).

Once the conditional TEs and surrogates have been computed for every candidate source interval, we select the interval which had the highest bias-corrected conditional TE (line 22). We then estimate the  $p$  value associated with the null hypothesis of the conditional TE from this source interval being zero (line 23). This is done using the maximum statistic test (see [Sect 4.2](#)).

Algorithm 1. Greedy TE algorithm for network inference from event-based data.

```

input  : /* Set of the event times of each process */
         $\mathcal{R} = \{R_i\}_{i=1}^{N_{proc}}$ 
        /* Cutoff significance value for adding sources */
         $p_{cutoff}$ 
        /* Number of surrogate estimates to perform per TE estimation */
         $N_{surrogates}$ 

output : /* Set of the indices of the sources of each process */
         $\mathcal{S} = \{S_i\}_{i=1}^{N_{proc}}$ 

/* Iterate over all target processes. */
1 for  $i \leftarrow 1$  to  $N_{proc}$  do
  /* Variable to keep track of added sources, as well as the added target inter-event
  intervals. If  $c_k$  is zero, this implies the source is not added. Otherwise, the value of  $c_k$ 
  indicates the number of added inter-event intervals. */
2   $C \leftarrow \{0\}_{k=1}^{N_{proc}}$ 
3   $p \leftarrow 0$ 
  /* Keep trying the next target interval, until the added information is no longer
  significant. */
4  while  $p < p_{cutoff}$  do
    /* Estimate the TE from the next interval of the target under consideration,
    conditional on all target intervals already added. */
5     $a \leftarrow \text{estimateConditionalAIS}(i, c_i + 1, R_i)$ 
    /* Estimate the associated surrogate values. */
6     $\{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}} \leftarrow \text{estimateConditionalAISonSurrogates}(i, c_i + 1, R_i, N_{surrogates})$ 
7     $p \leftarrow \text{calculatePVal}(t, \{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}})$ 
    /* If the null hypothesis is rejected, we increment the number of target intervals
    added */
8    if  $p < p_{cutoff}$  then
9       $c_i \leftarrow c_i + 1$ 
10   end
11 end
12  $p_{max} \leftarrow 0$ 
  /* Keep looking for sources until the candidate with maximum TE is not significant. */
13 while  $p_{max} < p_{cutoff}$  do
  /* Initialise TE and surrogate TE values to 0. */
14   $\{t_j\}_{j=1}^{N_{proc}} \leftarrow \{0\}_{j=1}^{N_{proc}}$ 
15   $\{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}} \leftarrow \{0\}_{n=1}^{N_{surrogates}}$ 
  /* Iterate over all processes except the target under consideration. */
16  for  $j = 1 \dots N_{proc}$  where  $j \neq i$  do
    /* Estimate the TE from the next interval of the source under consideration,
    conditional on all source intervals already added. */
17     $t_j \leftarrow \text{estimateConditionalTE}(i, j, c_j + 1, R_i, C)$ 
    /* Estimate the associated surrogate values. */
18     $\{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}} \leftarrow \text{estimateConditionalTEonSurrogates}(i, j, c_j + 1, R_i, C, N_{surrogates})$ 
    /* Bias-correct both the original and surrogate TE estimates by subtracting the
    mean of the surrogates. */
19     $t_j \leftarrow t_j - \text{mean}(\{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}})$ 
20     $\{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}} \leftarrow \{t_{surrogate,j,n} - \text{mean}(\{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}})\}_{n=1}^{N_{surrogates}}$ 
21  end
  /* Find which candidate source interval had the highest bias-corrected TE estimate. */
22   $j_{max} \leftarrow \text{findIndexOFMax}(\{t_j\}_{j=1}^{N_{proc}})$ 
  /* Estimate the p value corresponding to the null hypothesis that this source had zero
  TE. */
23   $p_{max} \leftarrow \text{calculateMaxStatPVal}(t_{j_{max}}, \{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}})$ 
  /* If the null hypothesis is rejected, add the candidate source interval to the set of
  sources. */
24  if  $p_{max} < p_{cutoff}$  then
25     $c_{j_{max}} \leftarrow c_{j_{max}} + 1$ 
26  end
27 end
  /* We now check that all added source intervals are still significant, when conditioning
  on all other added sources. */
28  $p_{max} \leftarrow 0$ 
  /* Iterate until no checked source interval is not significant. */
29 while  $p_{max} > p_{cutoff}$  do
30   $\{p_j\}_{j=1}^{N_{proc}} \leftarrow \{1\}_{j=1}^{N_{proc}}$ 
  /* Iterate over all processes that had any intervals added. */
31  for  $j = 1$  to  $N_{proc}$  where  $j \neq i$  and  $c_j > 0$  do
    /* Estimate the TE from the last interval of the source under consideration,
    conditional on all other added source intervals. */
32     $t \leftarrow \text{estimateConditionalTE}(i, j, c_j, R_i, C)$ 
    /* Estimate the associated surrogate values. */
33     $\{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}} \leftarrow \text{estimateConditionalTEonSurrogates}(i, j, c_j, R_i, C, N_{surrogates})$ 
    /* Estimate the p value corresponding to the null hypothesis that this source had
    zero TE. */
34     $p_j \leftarrow \text{calculatePVal}(t, \{t_{surrogate,j,n}\}_{n=1}^{N_{surrogates}})$ 
35  end
  /* Find the index of the source whose final interval had the highest p value. */
36   $j_{max} \leftarrow \text{findIndexOFMax}(\{p_j\}_{j=1}^{N_{proc}})$ 
37   $p_{max} \leftarrow p_{j_{max}}$ 
  /* If the null hypothesis is not rejected for this source interval, remove the
  candidate source interval from the set of sources. */
38  if  $p_{max} > p_{cutoff}$  then
39     $c_{j_{max}} \leftarrow c_{j_{max}} - 1$ 
40  end
41 end
  /* Add the indices of the sources which had any intervals remaining in the conditional set
  to the final set of inferred sources for the given target. */
42   $S_i \leftarrow \{j : j \in \{1, 2, \dots, N_{proc}\} \text{ and } c_j > 0\}$ 
43 end

```

The above process continues until the selected candidate source interval (with maximum bias-corrected conditional TE) fails the significance test (line 13). The algorithm then moves onto the final pruning step, where it is checked that every source retains a statistically significant conditional TE, once conditioning on every other process added to the conditioning set. This step is necessary as a source added early in the process might be providing information about the target that is fully redundant with that held by sources added later in the greedy building of the conditioning set. Such redundant sources need to be removed.

To perform the pruning, we continually try removing source intervals from the conditioning set one-by-one, until every final source interval in the set is found to be statistically significant. In a mirror image to how the candidate intervals, for a given source, are added iteratively from the most recent and then further back in time, they are removed in order from the last interval to the most recent. In each round of pruning, we iterate over all sources which had an interval added to the conditioning set (line 31). We then estimate the conditional TE and associated surrogates for the last remaining added interval for that source (lines 32 and 33). We then calculate the  $p$  value corresponding to the null hypothesis (line 34) of zero TE in the normal manner (that is, not using the maximum statistic test). After iterating over all sources in the conditioning set, we then find the source index with the maximum  $p$  value (line 36). If this  $p$  value is greater than the specified  $\alpha$  cutoff, then we remove the last added interval for that source from the conditioning set (line 39).

As per Novelli et al. [7, Supporting Information], the number of TE calculations required, for each target, scales as  $\mathcal{O}(N_{\text{proc}}dS)$ , where in this context  $N_{\text{proc}}$  is the number of source processes,  $d$  is the average number of inferred source intervals for that target, and  $S$  is the number of surrogates per TE calculation. The number of TE calculations across all  $N_{\text{proc}}$  targets is  $\mathcal{O}(N_{\text{proc}}^2dS)$ , though of course this can be parallelised over targets. The time complexity of each TE calculation is described in Sect 4.4.

## 4.2. Maximum statistic test

When considering adding sources to the conditioning set, we test the candidate source with the highest TE using the maximum statistic test (line 23 of Algorithm 1).

It is worth briefly describing the usual method for testing for non-zero TE using surrogates. We generate  $N_{\text{surrogates}}$  surrogates, which conform to the null hypothesis of no temporal relationship (zero TE), using a given surrogate generation algorithm (see Sect 4.5 for a description of the surrogate generation method used here). We then estimate the TE on each of these generated surrogate series. The proportion of these estimates which are greater than or equal to the estimate on the original data is then an estimate of the probability that we would observe a value greater than or equal to what we estimated on the original data, under the null hypothesis of zero TE (and therefore it is our  $p$  value).

Novelli et al. [7] highlighted the fact that using this test as is, when adding sources to the conditioning set, would lead to high false-positive rates. This is effectively a multiple comparisons problem, in that the test is being performed on the *maximum* estimated TE value from the set of candidate sources.

In order to compensate for this, we replicate the selection of the maximum candidate source in the significance testing step. Specifically, for each  $i \in \{1, 2, \dots, N_{\text{surrogates}}\}$ , we compare the TE estimates on the  $i^{\text{th}}$  surrogate for each of the candidate sources. We select the maximum such estimate for each  $i$ . The population of surrogate values for the maximum statistic test is then made up of the resulting  $N_{\text{surrogates}}$  maximum values. The test then proceeds as normal.

### 4.3. Transfer entropy estimation—time-binning method

The transfer entropy [12,13] is defined as the information provided by the past state or history embedding  $\mathbf{y}_{<t}$  of a source variable (capturing its sequential temporal patterns) about the next value  $x_t$  of a target variable, given the past state of that target  $\mathbf{x}_{<t}$ . Specifically, this is a conditional mutual information:

$$\mathbf{T}_{Y \rightarrow X} = I(\mathbf{y}_{<t}; x_t | \mathbf{x}_{<t}). \quad (3)$$

For standard time series processes  $x_t$  and  $y_t$ , the relevant past states are usually represented by Takens embedding vectors up to some embedding dimensions  $l$  and  $k$ , e.g.  $\mathbf{y}_{<t} = \{y_{t-1}, y_{t-2}, \dots, y_{t-l}\}$  and  $\mathbf{x}_{<t} = \{x_{t-1}, x_{t-2}, \dots, x_{t-k}\}$  (and potentially with embedding delays between the samples utilised as well); see [13]. The traditional method applying TE to spike trains operates by first discretising the spiking process into time bins (using some time bin width  $\Delta t$ ), providing a time series of binary values (spiking or not) on which the TE is then estimated. As above, the estimation of TE requires the use of embedding vectors of these binary time series to represent histories for the target, source and conditioning processes in the relevant conditional probability distributions for the target process. In order for these embeddings to both extend over a reasonable period of time and also capture fine subtleties in event timings, each embedding vector needs to consist of multiple time bins. Capturing effects occurring on both fine and large time scales is necessary as it is known that correlations in spike trains exist over distances of (at least) hundreds of milliseconds [55,56]. Moreover, it is established that correlations at the millisecond and sub-millisecond scale play a role in neural function [57–60]. The use of multi-bin embedding vectors causes an explosion in the dimensionality of the state space over which probability distributions need to be estimated as conditioning processes are added, rendering the estimation of TE with substantial sets of conditioning processes infeasible. In the next section we discuss a recently developed TE estimator for spike trains which more efficiently handles the state space of interactions in such data.

### 4.4. Transfer entropy estimation for spike trains in continuous time

It has, relatively recently, been shown that, for event-based data such as spike-trains, in the limit of small bin size  $\Delta t$ , that the expected TE rate  $\dot{\mathbf{T}}_{Y \rightarrow X} = \lim_{\Delta t \rightarrow 0} \frac{\mathbf{T}_{Y \rightarrow X}}{\Delta t}$  is given by the following expression [61]:

$$\dot{\mathbf{T}}_{Y \rightarrow X} = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{i=1}^{N_X} \ln \frac{\lambda_{x|\mathbf{x}_{<t}, \mathbf{y}_{<t}}[\mathbf{x}_{<x_i}, \mathbf{y}_{<x_i}]}{\lambda_{x|\mathbf{x}_{<t}}[\mathbf{x}_{<x_i}]}. \quad (4)$$

Here,  $\lambda_{x|\mathbf{x}_{<t}, \mathbf{y}_{<t}}[\mathbf{x}_{<x_i}, \mathbf{y}_{<x_i}]$  is the instantaneous firing rate of the target conditioned on the histories of the target  $\mathbf{x}_{<x_i}$  and source  $\mathbf{y}_{<x_i}$  at the time points  $x_i$  of the spike events in the target process.  $\lambda_{x|\mathbf{x}_{<t}}[\mathbf{x}_{<x_i}]$  is the instantaneous firing rate of the target conditioned on its history alone, ignoring the history of the source. It is important to note that the sum is being taken over the  $N_X$  spikes of the target during the sampling period  $\tau$ : thereby evaluating log ratios of the expected spike rates of the target given source and target histories versus target histories alone, *when* the target does spike. As this expression allows us to ignore the “empty space” between events, it presented clear potential for allowing for more efficient estimation of TE on spike trains.

This potential was recently realised in a new continuous-time estimator of TE presented in [27], and all TE estimates in this paper were performed using this new estimator. In [27] it is

demonstrated that this continuous-time estimator is far superior to the traditional discrete-time approach to TE estimation on spike trains. For a start, unlike the discrete-time estimator, it is consistent. That is, in the limit of infinite data, it will converge to the true value of the TE. It was also shown to have much preferable bias and convergence properties. Most significantly, perhaps, this new estimator utilises the inter-spike intervals to efficiently represent the history embeddings  $\mathbf{x}_{<x_i}$  and  $\mathbf{y}_{<x_i}$  in estimating the relevant conditional spike rates in (4).

This is in contrast with the traditional discrete-time estimator which uses the presence or absence of spikes in an array of time bins as its history embeddings (it sometimes also uses the number of spikes occurring in a bin). In order to avoid the dimensionality of the estimation problem becoming sufficiently large so as to render estimation infeasible, only a small number of bins can be used in these embeddings. To focus in on cell-culture data, previous applications of TE to this type of data have used a variety of bin sizes: 40  $\mu\text{s}$  [19], 0.3 ms [15], and 1 ms [20,23]. Some studies chose to examine the TE values produced by multiple different bin widths, specifically: 0.6 ms and 100 ms [21], 1.6 ms and 3.5 ms [24] and 10 different widths ranging from 1 ms to 750 ms [22]. And specifically, those studies demonstrated the unfortunate high sensitivity of the discrete-time TE estimator to the bin width parameter. Moreover, all of these studies have only used a single bin in the history embeddings. In the instances where narrow ( $< 5$  ms) bins were used, only a very narrow slice of history is being considered in the estimation of the history-conditional spike rate. This is problematic, as it is known that correlations in spike trains exist over distances of (at least) hundreds of milliseconds [55,56]. Conversely, in the instances where broad ( $> 5$  ms) bins were used, relationships occurring on fine time scales will be completely missed. This is significant given that it is established that correlations at the millisecond and sub-millisecond scale play a role in neural function [57–60]. In other words, previous applications of transfer entropy to electrophysiological data from cell cultures either captured some correlations occurring with fine temporal precision or they captured relationships occurring over larger intervals, but never both simultaneously. This can be contrasted with the inter-spike interval history representation used by the continuous-time estimator. To take a concrete example, in the *in vitro* data we used, for the recording on day 24 of culture 1-3, the average interspike interval was 0.71 seconds. This implies that the history embeddings used are at least on average 0.71 seconds long, being longer than this in cases where multiple intervals are being used. This is despite the fact that our history representations retain the precision of the raw data (40  $\mu\text{s}$ ) and the ability to measure relationships on this scale where they are relevant (via the underlying nearest-neighbour estimators). Furthermore, the innovative representation of history embeddings as an array of inter-spike intervals allows for the application of the highly effective nearest-neighbour family of information-theoretic estimators [53,62], which bring estimation efficiency and bias correction.

The challenges of using the discrete-time estimator only become more severe when one attempts to infer networks using conditional TEs. As there are now more processes being considered by the estimator (those in the conditioning set) the dimensionality of the estimation problem increases faster as we increase the embedding length. This places further pressure on keeping the number of bins in each embedding low, thus increasing the harshness in the tradeoff between history length and temporal precision. This is the likely reason behind the fact that almost all previous studies which evaluated the use of TE for the inference of spiking networks only made use of pairwise TE estimates [15–18]. This is as opposed to the multivariate conditional TE estimation used here, which takes into account the relationship of the target to other processes when considering its relationship to the given source.

The parameters used with this estimator for the simulated data are shown in Table 1 and those used for *in vitro* spike recordings are shown in Table 2. The effect of varying these

**Table 1. The parameter values used in the continuous-time TE estimator when used for the inference of the simulated spiking networks.** A complete description of these parameters, along with analysis and discussion of their effects can be found in [27]. Note that the history embedding length is not a fixed parameter. Rather, it is determined by Algorithm 1.

Parameter	Description	Value
$N_X$	Number of spikes in the target spike train	varied (see text)
$k_{\text{global}}$	Number of nearest neighbours to find in the initial search	4
$k_{\text{perm}}$	Number of nearest neighbours to consider during surrogate generation	20
$N_U$	Number of random samples of histories at non-spiking points in time	$20N_X$
$N_{U,\text{surrogates}}$	Number of random samples of histories at non-spiking points in time used for surrogate generation	$20N_X$
$N_{\text{surrogates}}$	Number of surrogates to generate for each node pair	100

<https://doi.org/10.1371/journal.pcbi.1013500.t001>

**Table 2. The parameter values used in the continuous-time TE estimator when used for network inference on the *in vitro* spike recordings.** A complete description of these parameters, along with analysis and discussion of their effects can be found in [27].

Parameter	Description	Value
$N_X$	Number of spikes in the target spike train	varied (see text)
$k_{\text{global}}$	Number of nearest neighbours to find in the initial search	10
$k_{\text{perm}}$	Number of nearest neighbours to consider during surrogate generation	20
$N_U$	Number of random samples of histories at non-spiking points in time	$20N_X$
$N_{U,\text{surrogates}}$	Number of random samples of histories at non-spiking points in time used for surrogate generation	$20N_X$
$N_{\text{surrogates}}$	Number of surrogates to generate for each node pair	100

<https://doi.org/10.1371/journal.pcbi.1013500.t002>

parameters is explored in the original paper regarding the continuous-time estimator [27]. The chief difference in the parameter values used in these situations is that, for the *in vitro* recordings a larger value of  $k_{\text{global}}$  (the number of nearest neighbours to consider in the initial searches) was employed (10 compared to 4). This was due to the observation that the estimates on the *in vitro* recordings exhibited much higher variance than those on the simulated data. It is a known property of nearest-neighbour information theoretic estimators that considering larger numbers of neighbours reduces their variance [53].

The time complexity of each TE calculation using this nearest neighbour algorithm is approximately  $\mathcal{O}(k_{\text{global}}N_X \log(N_XN_U))$  (as  $N_X$  instances of fast  $k_{\text{global}}$  nearest neighbour searches in spaces of  $N_X$  and  $N_U$  samples respectively).

As in the authors' previous work applying the continuous-time TE estimator to *in vitro* spike recordings [31], a small change was made to the estimation procedure described in [27]. This was made in how random sample points were placed along the process both for the estimation of the TE and the generation of surrogates. Instead of laying out the  $N_U$  and  $N_{U,\text{surrogates}}$  sample points randomly uniformly, we placed each one at an existing target spike, with the addition of uniform noise on the interval  $[-80 \text{ ms}, 80 \text{ ms}]$ . This was due to the fact that these recordings contain incredibly dense bursts. Such a sampling strategy is required in order to adequately sample these regions of intense activity.

An implementation of the estimator contained in the Java Information Dynamics Toolkit (JIDT) [63] software package was used in this study.

#### 4.5. Surrogate generation

Surrogate processes were generated by applying an adaptation of the permutation method of Runge [64] to the spiking TE estimator, as detailed in [27]. In brief, this method permutes

the history embedding vectors to destroy the relationship between the source intervals and the existence or absence of spiking in the target. However, it retains the relationship between the source history embedding intervals and the embedding intervals from the target and conditioning processes.

#### 4.6. Spiking network simulation

All network simulations were conducted using Leaky-Integrate-and-Fire (LIF) model neurons [37]. In this model, the membrane potential of the  $i^{\text{th}}$  evolves according to:

$$\tau \frac{dV_i}{dt} = (V_0 - V_i) + R_m I_{\text{syn},i}. \quad (5)$$

When  $V_i$  crosses the threshold  $V_{\text{threshold}}$ , the timestamp of crossing is recorded as a spike.  $V_i$  is then set to  $V_{\text{reset}}$  and the evolution of the membrane potential is subsequently paused for the duration of the hard refractory period.  $I_{\text{syn},i}$  is the synaptic input current into neuron  $i$ . Neurons were connected using alpha synapses [65]. Each synapse connecting neuron  $j$  to neuron  $i$  evolves according to:

$$I_{i,j}(t) = a_{i,j} \bar{\alpha}_j \sum_{t_s \in S_{i,j}} \frac{t - t_s}{\tau_{\text{syn}}} \exp\left(-\frac{t - t_s}{\tau_{\text{syn}}}\right). \quad (6)$$

$A$  is the connectivity matrix, with  $a_{i,j} = 1$  indicating that neuron  $j$  is a pre-synaptic input to neuron  $i$  and  $a_{i,j} = 0$  indicating otherwise.  $\bar{\alpha}_j$  is the connection strength of the afferent connections from neuron  $j$ . All excitatory neurons share the same afferent connection strength  $\bar{\alpha}_{\text{exc}}$ . Inhibitory neurons, by contrast, have connection strength  $g\bar{\alpha}_{\text{exc}}$ . The sum is taken over the set of spike times in neuron  $j$  occurring before time  $t$ ,  $S_{i,j}$ . The synaptic current for neuron  $i$  is then the sum of the currents from all other neurons in the network, that is,  $I_{\text{syn},i} = \sum I_{i,j}$ .

Each neuron was connected to exactly  $n_{\text{exc}}$  excitatory sources, chosen randomly from the set of excitatory neurons. Similarly, each neuron was connected to  $n_{\text{inh}}$  inhibitory sources, chosen randomly from the inhibitory sources. The specific parameter values used in the experiments described in Sect 2 are shown in Tables 3 and 4.

**Table 3. The parameter values used in the LIF network simulations at various levels of synchrony, presented in Sect 2.1.**

Parameter	Description	Value
$N_{\text{exc}}$	Number of excitatory neurons	30
$N_{\text{inh}}$	Number of inhibitory neurons	20
$\tau$	Membrane time constant	20 ms
$R_m$	Membrane resistance	1 $\Omega$
$V_{\text{reset}}$	Membrane reset potential	0 mV
$V_0$	Membrane resting potential	0 mV
$V_{\text{threshold}}$	Spike threshold potential	40 mV
$\tau_{\text{syn}}$	Synaptic time constant	20 ms
$g$	Ratio of inhibitory to excitatory connection strength	{1, 1.5, 3}
$\bar{\alpha}_{\text{exc}}$	Excitatory connection strength	20 mA
$n_{\text{exc}}$	Number of excitatory sources per neuron	3
$n_{\text{inh}}$	Number of inhibitory sources per neurons	2
$\lambda_{\text{stim}}$	Rate of the Poisson stimuli	200 Hz
$\bar{\alpha}_{\text{stim}}$	Stimulus connection strength	6 mA

<https://doi.org/10.1371/journal.pcbi.1013500.t003>

**Table 4. The parameter values used in the LIF network simulations at various levels of synchrony, presented in Sect 2.2.**

Parameter	Description	Value
$N_{exc}$	Number of excitatory neurons	30
$N_{inh}$	Number of inhibitory neurons	20
$\tau$	Membrane time constant	20 ms
$R_m$	Membrane resistance	1 $\Omega$
$V_{reset}$	Membrane reset potential	0 mV
$V_0$	Membrane resting potential	0 mV
$V_{threshold}$	Spike threshold potential	40 mV
$\tau_{syn}$	Synaptic time constant	20 ms
$g$	Ratio of inhibitory to excitatory connection strength	3
$\bar{\alpha}_{exc}$	Excitatory connection strength	17.5 mA
$n_{exc}$	Number of excitatory sources per neuron	3
$n_{inh}$	Number of inhibitory sources per neurons	2
$\bar{\alpha}_{stim}$	Stimulus connection strength	4 mA

<https://doi.org/10.1371/journal.pcbi.1013500.t004>

Each neuron also received an independent stimulus. In the experiments presented in Sect 2.1, this source was an homogeneous Poisson point process. In the experiments presented in Sect 2.2, it contained varying amounts of regularity. Specifically, in Sect 2.2, each neuron received a stimulus with a spike rate  $\lambda$  drawn from a normal distribution with mean 500 hertz and standard deviation of 25 hertz. In the fully random case, the stimulus was generated as an homogeneous Poisson point process, with rate  $\lambda$ . In the fully regular case, spikes were placed at a fixed interval of  $1/\lambda$ . In the semi-regular case, the spikes were placed with this same fixed interval, but gaussian noise with mean 0 and standard deviation 0.5 millisecond was added to each spike time. The connectivity strength between each stimulus and its target was specified by the parameter  $\bar{\alpha}_{stim}$ .

#### 4.7. Analysis of *in vitro* data

We made use of the same dataset as in the authors' previous study [31] and analysed it in a very similar fashion. As such, the following section closely follows the discussion of this dataset in that previous work.

The spike train recordings used in this study were collected by Wagenaar et al. [34] and are freely available online [46]. The details of the methodology used in these recordings can be found in the original publication [34]. A short summary of their methodology follows:

Dissociated cultures of rat cortical neurons had their activity recorded. This was achieved by plating 8x8 Multi-Electrode Arrays (MEAs), operating at a sampling frequency of 25 kHz with neurons obtained from the cortices of rat embryos. The spacing between the electrodes was 200  $\mu\text{m}$  center-to-center. The MEAs did not have electrodes on their corners and one electrode was used as ground, resulting in recordings from 59 electrodes. In all recordings, electrodes with less than 100 spikes were removed from the analysis. This resulted in electrodes 37 and 43 being removed from every recording as no spikes were recorded on them. The spatial layout of the electrodes is available from the website associated with the dataset [46], allowing us to overlay the inferred networks onto this spatial layout as is done in Fig 4.

Inspection of Fig 2 in the paper describing the original dataset [34] strongly suggests that these electrode recordings are sub-sampling the population of neurons. That is, spikes are only being detected from a fraction of the actual population. This is a possible cause of the low degree of the nodes in Fig 4.

Recordings were conducted on most days, starting from 3-4 Days *In Vitro* (DIV). The end point of recording varied between 25 and 39 DIV. Longer overnight recordings were also conducted on some cultures at sparser intervals. In this work we make use of these longer overnight recordings, since longer data sets are required to support statistical power of our analysis methods. In alignment with our previous work [31], we select cultures which have at least three overnight recordings, so as to track change in connectivity over time. These recordings were split into multiple files. The specific files used, along with the names of the cultures and days of the recordings are listed in Table 5. 30 Minute windows of spiking activity were extracted and used for network inference. Specifically, the number of target spikes  $N_X$  was set as the number of spikes that fell within this 30 minute window for the given target neuron.

The original study plated the electrodes with varying densities of cortical cells. However, overnight recordings were only performed on the ‘dense’ cultures, plated with a density of 2500 cells/ $\mu$ L.

The original study performed threshold-based spike detection by determining that a spike was present in the case of an upward or downward excursion beyond 4.5 times the estimated RMS noise of the recorded potential on a given electrode. The analysis presented in this paper makes use of these detected spike times. No spike sorting was performed and, as such, we are studying multi-unit activity (MUA) [66].

As the data was sampled at 25 kHz, uniform noise distributed between  $-20 \mu$ s and  $20 \mu$ s was added to each spike time. This is to prevent the TE estimator from exploiting the fact that, in the raw data, inter-spike intervals are always an integer multiple of  $40 \mu$ s.

#### 4.8. GLM model

The implementation of Generalised Linear Models (GLMs) of spiking activity followed that of Song et al. [43] very closely. We briefly list the few minor differences.

For the B-spline basis functions, we excluded all knot locations beyond 100 ms. This was done due to the membrane potential decay time constant ( $\tau$ ) in the simulated models being set to 20 ms (see Table 4), implying that statistical relationships beyond 100 ms would be very unlikely.

Song et al. [43] propose finding the penalty weight parameter  $\lambda$  using the Bayesian information criterion (BIC), by iteratively trialling various penalty weight values. Performing this for each target spike train would have been computationally prohibitive given the large networks and long simulation times used in this work. Instead, this step was performed on a few trial runs and a single value of  $\lambda = 1 \times 10^{-3}$  was chosen as it closely approximated that chosen by the BIC in all such trial runs.

We chose to designate the existence of a connection between a source and target when the GLM for the given target contained one or more non-zero weights assigned to the basis-splines associated with a given source.

Fitting of the GLM models was performed using the statsmodels [67] Python library.

**Table 5. File numbers used for each culture on each day.** These correspond to the file numbering used in the freely available dataset used in this study, provided by Wagenaar et al.[34,46]

<b>Culture 1-1</b>	day 4	day 14	day 20	
	2	2	2	
<b>Culture 1-3</b>	day 5	day 10	day 16	day 24
	2	2	2	2
<b>Culture 2-2</b>	day 9	day 15	day 21	day 33
	2	2	2	2

<https://doi.org/10.1371/journal.pcbi.1013500.t005>

## Supporting information

**S1 Fig. CoNNECT inference performance.** Plots showing the resulting precision and recall from running the CoNNECT algorithm [33] on networks of LIF neurons composed of 30 excitatory neurons and 20 inhibitory neurons. The ratio of the inhibitory to excitatory connection strength was varied in order to change the degree of synchrony in the network. Plots are shown for three different synchrony levels. Each plot contains points for the precision and recall for the inhibitory and excitatory neurons as well as the combined precision and recall. The lines pass through the means of these points.

(EPS)

**S2 Fig. GLM inference performance.** Plots showing the resulting precision and recall from using a GLM model of spiking activity [43] to infer the connectivity of networks of LIF neurons composed of 30 excitatory neurons and 20 inhibitory neurons. The ratio of the inhibitory to excitatory connection strength was varied in order to change the degree of synchrony in the network. Plots are shown for three different synchrony levels. Each plot contains points for the precision and recall for the inhibitory and excitatory neurons as well as the combined precision and recall. The lines pass through the means of these points.

(EPS)

**S3 Fig. GLM inference with large bin width.** Identical plots to S2 Fig, but showing the use of a larger value of the time bin width  $\Delta t$  (40 ms as oppose to 20 ms). There is a substantial drop in precision at this larger bin width. Note that, without access to the ground truth, there is no principled way to choose a best bin size. Although the results shown in S2 Fig are quite promising, it is unlikely that they could ever be achieved in practice, as finding the optimal parameters (such as bin size) requires access to the ground truth.

(EPS)

**S4 Fig. GLM inference with original knot locations.** Identical plots to S2 Fig, but including B-spline knot locations beyond 100 ms. All 16 knot locations used Song et al. [43] were incorporated, as opposed to only using the first 6, which was done elsewhere in this paper. Refer to Section IV H for more details on the use of B-splines in the GLM model. There is a substantial drop in precision when using these knot locations. Note that, without access to the ground truth, there is no principled way to choose the optimal knot locations. Although the results shown in S2 Fig are quite promising, it is unlikely that they could ever be achieved in practice, as finding the optimal parameters (such as knot locations) requires access to the ground truth.

(EPS)

## References

1. Stevenson IH, Kording KP. How advances in neural recording affect data analysis. *Nat Neurosci.* 2011;14(2):139–42. <https://doi.org/10.1038/nn.2731> PMID: 21270781
2. Stevenson IH. Tracking advances in neural recordings. [stevenson.lab.uconn.edu/scaling/#](http://stevenson.lab.uconn.edu/scaling/#).
3. Yuan X, Schröter M, Obien MEJ, Fiscella M, Gong W, Kikuchi T, et al. Versatile live-cell activity analysis platform for characterization of neuronal dynamics at single-cell and network level. *Nat Commun.* 2020;11(1):4854. <https://doi.org/10.1038/s41467-020-18620-4> PMID: 32978383
4. Sejnowski TJ, Churchland PS, Movshon JA. Putting big data to good use in neuroscience. *Nat Neurosci.* 2014;17(11):1440–1. <https://doi.org/10.1038/nn.3839> PMID: 25349909
5. Bassett DS, Sporns O. Network neuroscience. *Nat Neurosci.* 2017;20(3):353–64. <https://doi.org/10.1038/nn.4502> PMID: 28230844
6. Bzdok D, Yeo BTT. Inference in the age of big data: future perspectives on neuroscience. *Neuroimage.* 2017;155:549–64. <https://doi.org/10.1016/j.neuroimage.2017.04.061> PMID: 28456584
7. Novelli L, Wollstadt P, Mediano P, Wibral M, Lizier JT. Large-scale directed network inference with multivariate transfer entropy and hierarchical statistical testing. *Netw Neurosci.* 2019;3(3):827–47.

- [https://doi.org/10.1162/netn\\_a\\_00092](https://doi.org/10.1162/netn_a_00092) PMID: 31410382
8. Lizier J, Rubinov M. Multivariate construction of effective computational networks from observational data. 25. Max-Planck-Institut für Mathematik in den Naturwissenschaften; 2012.
  9. Stramaglia S, Faes L, Cortes JM, Marinazzo D. Disentangling high-order effects in the transfer entropy. *Phys Rev Research*. 2024;6(3):1L032007. <https://doi.org/10.1103/physrevresearch.6.1032007>
  10. Lizier JT, Bertschinger N, Jost J, Wibral M. Information decomposition of target effects from multi-source interactions: perspectives on previous, current and future work. 2018.
  11. Faes L, Porta A, Nollo G, Javorka M. Information decomposition in multivariate systems: definitions, implementation and application to cardiovascular networks. *Entropy*. 2016;19(1):5. <https://doi.org/10.3390/e19010005>
  12. Schreiber T. Measuring information transfer. *Phys Rev Lett*. 2000;85(2):461–4. <https://doi.org/10.1103/PhysRevLett.85.461> PMID: 10991308
  13. Bossomaier T, Barnett L, Harré M, Lizier JT. *Transfer entropy*. Springer; 2016.
  14. Vicente R, Wibral M, Lindner M, Pipa G. Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *J Comput Neurosci*. 2011;30(1):45–67. <https://doi.org/10.1007/s10827-010-0262-3> PMID: 20706781
  15. Garofalo M, Nieus T, Massobrio P, Martinoia S. Evaluation of the performance of information theory-based methods and cross-correlation to estimate the functional connectivity in cortical networks. *PLoS One*. 2009;4(8):e6482. <https://doi.org/10.1371/journal.pone.0006482> PMID: 19652720
  16. Ito S, Hansen ME, Heiland R, Lumsdaine A, Litke AM, Beggs JM. 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
  17. Stetter O, Battaglia D, Soriano J, Geisel T. Model-free reconstruction of excitatory neuronal connectivity from calcium imaging signals. *PLoS Comput Biol*. 2012;8(8):e1002653. <https://doi.org/10.1371/journal.pcbi.1002653> PMID: 22927808
  18. 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
  19. Nigam S, Shimono M, Ito S, Yeh F-C, Timme N, Myroshnychenko M, et al. Rich-club organization in effective connectivity among cortical neurons. *J Neurosci*. 2016;36(3):670–84. <https://doi.org/10.1523/JNEUROSCI.2177-15.2016> PMID: 26791200
  20. Shimono M, Beggs JM. Functional clusters, hubs, and communities in the cortical microconnectome. *Cereb Cortex*. 2015;25(10):3743–57. <https://doi.org/10.1093/cercor/bhu252> PMID: 25336598
  21. Matsuda E, Mita T, Hubert J, Oka M, Bakkum D, Frey U, et al. Multiple time scales observed in spontaneously evolved neurons on high-density CMOS electrode array. In: *Artificial Life Conference Proceedings 13*. MIT Press; 2013. p. 1075–82.
  22. Timme N, Ito S, Myroshnychenko M, Yeh F-C, Hiolski E, Hottowy P, et al. Multiplex networks of cortical and hippocampal neurons revealed at different timescales. *PLoS One*. 2014;9(12):e115764. <https://doi.org/10.1371/journal.pone.0115764> PMID: 25536059
  23. Kajiwara M, Nomura R, Goetze F, Kawabata M, Isomura Y, Akutsu T, et al. Inhibitory neurons exhibit high controlling ability in the cortical microconnectome. *PLoS Comput Biol*. 2021;17(4):e1008846. <https://doi.org/10.1371/journal.pcbi.1008846> PMID: 33831009
  24. Timme NM, Ito S, Myroshnychenko M, Nigam S, Shimono M, Yeh F-C, et al. High-degree neurons feed cortical computations. *PLoS Comput Biol*. 2016;12(5):e1004858. <https://doi.org/10.1371/journal.pcbi.1004858> PMID: 27159884
  25. Mijatovic G, Antonacci Y, Loncar-Turukalo T, Minati L, Faes L. An information-theoretic framework to measure the dynamic interaction between neural spike trains. *IEEE Trans Biomed Eng*. 2021;68(12):3471–81. <https://doi.org/10.1109/TBME.2021.3073833> PMID: 33872139
  26. Gourévitch B, Eggermont JJ. Evaluating information transfer between auditory cortical neurons. *J Neurophysiol*. 2007;97(3):2533–43. <https://doi.org/10.1152/jn.01106.2006> PMID: 17202243
  27. Shorten DP, Spinney RE, Lizier JT. Estimating transfer entropy in continuous time between neural spike trains or other event-based data. *PLoS Comput Biol*. 2021;17(4):e1008054. <https://doi.org/10.1371/journal.pcbi.1008054> PMID: 33872296
  28. Faes L, Nollo G, Porta A. Information-based detection of nonlinear granger causality in multivariate processes via a nonuniform embedding technique. *Phys Rev E Stat Nonlin Soft Matter Phys*. 2011;83(5 Pt 1):051112. <https://doi.org/10.1103/PhysRevE.83.051112> PMID: 21728495

29. Sun J, Taylor D, Bollt EM. Causal network inference by optimal causation entropy. *SIAM J Appl Dyn Syst*. 2015;14(1):73–106. <https://doi.org/10.1137/140956166>
30. Vlachos I, Kugiumtzis D. Nonuniform state-space reconstruction and coupling detection. *Phys Rev E Stat Nonlin Soft Matter Phys*. 2010;82(1 Pt 2):016207. <https://doi.org/10.1103/PhysRevE.82.016207> PMID: 20866707
31. Shorten DP, Priesemann V, Wibral M, Lizier JT. Early lock-in of structured and specialised information flows during neural development. *Elife*. 2022;11:e74651. <https://doi.org/10.7554/eLife.74651> PMID: 35286256
32. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, et al. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*. 2008;454(7207):995–9. <https://doi.org/10.1038/nature07140> PMID: 18650810
33. Endo D, Kobayashi R, Bartolo R, Averbeck BB, Sugase-Miyamoto Y, Hayashi K, et al. A convolutional neural network for estimating synaptic connectivity from spike trains. *Sci Rep*. 2021;11(1):12087. <https://doi.org/10.1038/s41598-021-91244-w> PMID: 34103546
34. Wagenaar DA, Pine J, Potter SM. An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neurosci*. 2006;7:11. <https://doi.org/10.1186/1471-2202-7-11> PMID: 16464257
35. Runge J. Causal network reconstruction from time series: From theoretical assumptions to practical estimation. *Chaos*. 2018;28(7):075310. <https://doi.org/10.1063/1.5025050> PMID: 30070533
36. Spirtes P, Glymour CN, Scheines R, Heckerman D. Causation, prediction, and search. MIT Press; 2000.
37. Burkitt AN. A review of the integrate-and-fire neuron model: I. Homogeneous synaptic input. *Biol Cybern*. 2006;95(1):1–19. <https://doi.org/10.1007/s00422-006-0068-6> PMID: 16622699
38. Gerstner W, Kistler WM, Naud R, Paninski L. Neuronal dynamics: From single neurons to networks and models of cognition. Cambridge University Press; 2014.
39. Kriener B, Enger H, Tetzlaff T, Plesser HE, Gewaltig M-O, Einevoll GT. Dynamics of self-sustained asynchronous-irregular activity in random networks of spiking neurons with strong synapses. *Front Comput Neurosci*. 2014;8:136. <https://doi.org/10.3389/fncom.2014.00136> PMID: 25400575
40. Reconstructing neuronal circuitry from spike trains. <https://s-shinomoto.com/CONNECT/>
41. Linderman S, Stock CH, Adams RP. A framework for studying synaptic plasticity with neural spike train data. In: *Advances in Neural Information Processing Systems*. 2014. <https://proceedings.neurips.cc/paper/2014/file/4122cb13c7a474c1976c9706ae36521d-Paper.pdf>
42. Calabrese A, Schumacher JW, Schneider DM, Paninski L, Woolley SMN. A generalized linear model for estimating spectrotemporal receptive fields from responses to natural sounds. *PLoS One*. 2011;6(1):e16104. <https://doi.org/10.1371/journal.pone.0016104> PMID: 21264310
43. Song D, Wang H, Tu CY, Marmarelis VZ, Hampson RE, Deadwyler SA, et al. Identification of sparse neural functional connectivity using penalized likelihood estimation and basis functions. *J Comput Neurosci*. 2013;35(3):335–57. <https://doi.org/10.1007/s10827-013-0455-7> PMID: 23674048
44. Gerwinn S, Macke JH, Bethge M. Bayesian inference for generalized linear models for spiking neurons. *Front Comput Neurosci*. 2010;4:12. <https://doi.org/10.3389/fncom.2010.00012> PMID: 20577627
45. Novelli L, Lizier JT. Inferring network properties from time series using transfer entropy and mutual information: Validation of multivariate versus bivariate approaches. *Netw Neurosci*. 2021;5(2):373–404. [https://doi.org/10.1162/netn\\_a\\_00178](https://doi.org/10.1162/netn_a_00178) PMID: 34189370
46. Network activity of developing cortical cultures in vitro. <http://neurodatasharing.bme.gatech.edu/development-data/html/index.html>
47. Fisher RA. Statistical methods for research workers. 1936.
48. Varley T, Sporns O, Scherberger H, Dann B. Information dynamics in neuronal networks of macaque cerebral cortex reflect cognitive state and behavior. *bioRxiv*. 2021. 2021.09.05.458983.
49. Zalesky A, Fornito A, Cocchi L, Gollo LL, van den Heuvel MP, Breakspear M. Connectome sensitivity or specificity: which is more important?. *Neuroimage*. 2016;142:407–20. <https://doi.org/10.1016/j.neuroimage.2016.06.035> PMID: 27364472
50. Lizier JT, Prokopenko M, Zomaya AY. Local measures of information storage in complex distributed computation. *Information Sciences*. 2012;208:39–54. <https://doi.org/10.1016/j.ins.2012.04.016>
51. Spinney RE, Lizier JT. Characterizing information-theoretic storage and transfer in continuous time processes. *Phys Rev E*. 2018;98(1–1):012314. <https://doi.org/10.1103/PhysRevE.98.012314> PMID: 30110808
52. Mijatovic G, Stramaglia S, Faes L. A Model-free method to quantify memory utilization in neural point processes. *IEEE Trans Biomed Eng*. 2025;72(8):2532–43.

- <https://doi.org/10.1109/TBME.2025.3546842> PMID: 40031573
53. Kraskov A, Stögbauer H, Grassberger P. Estimating mutual information. *Phys Rev E Stat Nonlin Soft Matter Phys.* 2004;69(6 Pt 2):066138. <https://doi.org/10.1103/PhysRevE.69.066138> PMID: 15244698
  54. Lv Z-S, Zhu C-P, Nie P, Zhao J, Yang H-J, Wang Y-J, et al. Exponential distance distribution of connected neurons in simulations of two-dimensional in vitro neural network development. *Front Phys.* 2017;12(3). <https://doi.org/10.1007/s11467-017-0602-0>
  55. Aldridge JW, Gilman S. The temporal structure of spike trains in the primate basal ganglia: afferent regulation of bursting demonstrated with precentral cerebral cortical ablation. *Brain Res.* 1991;543(1):123–38. [https://doi.org/10.1016/0006-8993\(91\)91055-6](https://doi.org/10.1016/0006-8993(91)91055-6) PMID: 2054667
  56. Rudelt L, González Marx D, Wibrál M, Priesemann V. Embedding optimization reveals long-lasting history dependence in neural spiking activity. *PLoS Comput Biol.* 2021;17(6):e1008927. <https://doi.org/10.1371/journal.pcbi.1008927> PMID: 34061837
  57. Nemenman I, Lewen GD, Bialek W, de Ruyter van Steveninck RR. Neural coding of natural stimuli: information at sub-millisecond resolution. *PLoS Comput Biol.* 2008;4(3):e1000025. <https://doi.org/10.1371/journal.pcbi.1000025> PMID: 18369423
  58. Kayser C, Logothetis NK, Panzeri S. Millisecond encoding precision of auditory cortex neurons. *Proc Natl Acad Sci U S A.* 2010;107(39):16976–81. <https://doi.org/10.1073/pnas.1012656107> PMID: 20837521
  59. Sober SJ, Sponberg S, Nemenman I, Ting LH. Millisecond spike timing codes for motor control. *Trends Neurosci.* 2018;41(10):644–8. <https://doi.org/10.1016/j.tins.2018.08.010> PMID: 30274598
  60. Garcia-Lazaro JA, Belliveau LAC, Lesica NA. Independent population coding of speech with sub-millisecond precision. *J Neurosci.* 2013;33(49):19362–72. <https://doi.org/10.1523/JNEUROSCI.3711-13.2013> PMID: 24305831
  61. Spinney RE, Prokopenko M, Lizier JT. Transfer entropy in continuous time, with applications to jump and neural spiking processes. *Phys Rev E.* 2017;95(3–1):032319. <https://doi.org/10.1103/PhysRevE.95.032319> PMID: 28415203
  62. Kozachenko L, Leonenko NN. Sample estimate of the entropy of a random vector. *Problemy Peredachi Informatsii.* 1987;23(2):9–16.
  63. Lizier JT. JIDT: an information-theoretic toolkit for studying the dynamics of complex systems. *Front Robot AI.* 2014;1. <https://doi.org/10.3389/frobt.2014.00011>
  64. Runge J. Conditional independence testing based on a nearest-neighbor estimator of conditional mutual information. In: *International Conference on Artificial Intelligence and Statistics.* PMLR; 2018. p. 938–47.
  65. Roth A, van Rossum MCW. Modeling synapses. *Computational modeling methods for neuroscientists.* The MIT Press; 2009. p. 139–60. <https://doi.org/10.7551/mitpress/9780262013277.003.0007>
  66. Schroeter MS, Charlesworth P, Kitzbichler MG, Paulsen O, Bullmore ET. Emergence of rich-club topology and coordinated dynamics in development of hippocampal functional networks in vitro. *J Neurosci.* 2015;35(14):5459–70. <https://doi.org/10.1523/JNEUROSCI.4259-14.2015> PMID: 25855164
  67. Seabold S, Perktold J. Statsmodels: econometric and statistical modeling with python. In: *Proceedings of the 9th Python in Science Conference, Austin, TX, 2010.* p. 61.