A new approach for estimating effective connectivity from activity in neural networks

Niklas Laasch[1*], Wilhelm Braun[1*], Lisa Knoff[1], Jan Bielecki[3] , Claus C. Hilgetag[1,2]

[1] Institute of Computational Neuroscience, Center for Experimental Medicine, University Medical Center Hamburg-Eppendorf, Hamburg, Martinistraße 52, 20246 Hamburg, Germany

[2] Department of Health Sciences, Boston University, 635 Commonwealth Avenue, Boston, Massachusetts 02215, USA

[3] Institute of Physiology, Kiel University, 24118 Kiel, Germany

*** Correspondence:**
Corresponding Authors
niklas.laasch@posteo.de, w.braun@uke.de

Keywords: Computational Connectomics, Dynamical Systems, Hopf model, Effective Connectivity, Structural Connectivity, C. Elegans

Abstract

Inferring and understanding the underlying connectivity structure of a system solely from the observed activity of its constituent components is a challenge in many areas of science. In neuroscience, such link inference techniques for estimating connectivity are paramount when attempting to understand the network structure of neural systems from their recorded activity patterns. To date, no universally accepted method exists for the inference of effective connectivity, which describes how the activity of a neural node mechanistically affects the activity of other nodes. One practical challenge is that, without ground truth structural connectivity data, the inferred underlying structural connections cannot be validated. In this case, information on the nodal dynamics is needed to obtain a more complete causal understanding of the system in the form of its effective connectivity. Here, we describe a systematic comparison of different approaches for estimating effective connectivity. Starting with the Hopf neuron model with known ground truth structural connectivity, we reconstruct the system's structural connectivity matrix using a variety of algorithms. We show that, in sparse networks, combining a lagged-cross-correlation (LCC) approach with a recently published derivative-based correlation analysis method provides the most reliable estimation of the known ground truth connectivity matrix. We then use the estimated structural connectivity matrix as the basis for a forward simulation of the system dynamics, in order to recreate the observed node activity patterns. We show that, under certain conditions, our best method, LCC, works better than derivative-based methods for sparse noise-driven systems. Finally, we apply our LCC method to empirical biological data. Specifically, we reconstruct the structural connectivity of a subset of the nervous system of the nematode *C. Elegans*. We show that our computationally simple method performs better than another recently published, computationally more expensive reservoir computing-based method. Our results show that a comparatively simple method can be

used to reliably estimate directed effective connectivity in sparse neural systems in the presence of noise. We provide a perspective for future work by advocating the combined use of structural network estimation and activity recreation techniques for system identification, thus bridging the gap between structural and effective connectivity.

1 Introduction

What can be learnt about the network structure of a neural system from observations of the dynamics of its constituent components? In network neuroscience, three principal modes of connectivity can be distinguished (Farahani, Karwowski, and Lighthall 2019; Bullmore and Sporns 2009), specifically structural, functional, and effective connectivity (SC, FC, and EC, respectively). In this paper, we aim to estimate the SC, and finally, with knowledge of the nodal dynamics, the EC of a network-based neural system.

To start, SC (Yeh et al. 2021) refers to the actual anatomical connectivity among neural elements, for example by white matter tracts. In humans, diffusion tensor imaging is used to approximate SC (Maier-Hein et al. 2017). Obtaining reliable and detailed information on SC requires invasive experiments, which are prohibitive in humans (Sporns, Tononi, and Kötter 2005). In fact, the only organism with presently completely known SC is the nematode *C.elegans* (White et al. 1986; Yan et al. 2017). Recently, also the SC of the larval stage of the fruit fly *Drosophila* has become accessible (Winding et al. 2023). Second, FC (K. J. Friston et al. 1993) is the most commonly considered connectivity mode, for example in the context of resting-state functional magnetic resonance imaging (fMRI). By using correlation measures, this mode comprises statistical dependencies between brain regions that are not necessarily spatially close or structurally connected to each other. Finally, EC (Karl J. Friston 1994; 2011) estimates causal connectivity, via which the activity in one neural node mechanistically affects the activity in other nodes. Standard methods to obtain EC are Granger causality, dynamic causal modeling, Bayesian networks and transfer entropy (Farahani, Karwowski, and Lighthall 2019). In addition, there are also perturbation-based methods to infer EC, see (Valdes-Sosa et al. 2011; Klink et al. 2021; Fakhar and Hilgetag 2022).

A straightforward way to compute FC is to assign high connectivity to nodes whose time series co-fluctuate, hence showing a high (positive or negative) correlation, and assigning low connectivity to nodes that only have small correlation, possibly followed by thresholding. Following this approach, one can compute a FC matrix (Fornito, Zalesky, and Bullmore 2016; Rubinov and Sporns 2010) and compare it to reconstructed anatomical connections (i.e., SC) forming the structural connectome (Sporns, Tononi, and Kötter 2005). This approach has revealed that FC does not necessarily reflect static SC, e.g. (Fakhar, Hadaeghi, and Hilgetag 2022; Honey et al. 2007). The FC approach, although it is traditionally used in network neuroscience due to its mathematical simplicity, has a number of shortcomings. Because correlation is intrinsically symmetric, estimates of FC based on pairwise correlations result in symmetric connectivity matrices. Moreover, owing to false positives due to confounder motifs, activity correlations may not necessarily reflect structural connectivity (Chen, Rosen, and Sejnowski 2022). In this scenario, two structurally uncoupled nodes show high correlation because of common input from a third shared neighbor node, or from outside of the network. Because high correlation is equated with high connectivity, this setting incorrectly infers a connection between two highly correlated, yet structurally unconnected nodes. Improved techniques, such as partial correlation, have been devised to alleviate this problem, but they also result in symmetric connectivity matrices.

Recent work on the nematode *C.elegans* has also shown that structural connections alone cannot account for signal propagation and, instead, extrasynaptic mechanisms must be taken into account (Randi et al. 2023). Furthermore, neuromodulation is thought to significantly shape the activity of neural circuits, and therefore, SC again might only be the 'minimal structure' required for a causal understanding of any neural system (Marder 2012). Therefore, a causal understanding of neural network function is not provided solely by computing FC or SC alone.

In general, estimating SC, and EC, in neural systems is hard, both experimentally (Hagmann et al. 2008) and theoretically (Lainscsek et al. 2019). Recently, an approach called dynamic differential covariance (DDC) has been put forward to infer EC (Chen, Rosen, and Sejnowski 2022). Methods aiming to infer EC typically make assumptions about the structure of the network or the individual node dynamics before inferring connectivity. DDC assumes that any given time series data is either generated by a linear or a non-linear autonomous dynamical system and uses this structure to estimate EC, essentially by a series of matrix inversions. The method was shown to reliably estimate the directionality of connections even in the presence of confounder motifs, and to exhibit high noise tolerance, in the sense that individual trajectories could be perturbed by noise without a massive decrease in the accuracy of the algorithm (Chen, Rosen, and Sejnowski 2022). It also showed good performance when applied to non-stationary data, spiking neural networks and resting-state fMRI data. The validation of the method was carried out in cases when a ground truth (GT) connectivity matrix was known, either by design or by biological knowledge (Chen, Rosen, and Sejnowski 2022).

Besides shortcomings of any particular connectivity estimation algorithm, another frequent scenario poses a significant challenge: the GT connectivity of a system might not be known. In the absence of GT connectivity data and without a detailed mechanistic understanding of each node's internal dynamics, it is an open theoretical challenge to reliably estimate EC structure solely from neural time series recordings. Similarly, if after application of many different SC estimation algorithms, a common structure, for example a network core, becomes apparent, one can also be confident that the estimation procedure has reliably estimated at least some elements of the real network structure. Merely knowing a system's SC does not suffice to fully comprehend the system dynamics. However, if the system dynamics can be recreated from an estimated structural connectivity, assuming valid node dynamics, it would be a strong indication that the connectivity reflects physical connections in the system. For most systems, this is a formidable task. Indeed, the task of jointly inferring node dynamics and connectivity structure is in general a hard problem in system identification (Goodwin and Payne 1977; Shandilya and Timme 2011). In neural systems, this problem is often exacerbated by their high dimensionality and non-linearity of the underlying generating system, resulting in many parameters that can give rise to a particular observed activity data set (Prinz, Bucher, and Marder 2004; Deistler, Macke, and Gonçalves 2022). In a biological context, scenarios where a GT is lacking are especially relevant when it is hard or even impossible to trace actual physical connections, either because of the lack of transgenic animals or because the network under study is not self-contained so that some of the observed dynamics come from outside the system. In general, this scenario cannot be ruled out in real biological systems where rarely all network nodes can be imaged with high temporal and spatial resolution. Apart from lacking information on GT connectivity, in biological systems, nodal dynamics might also not be known. The reason is that again, the nodal dynamics might not be experimentally accessible in detail. This is especially true

for non-mammalian model systems, such as cnidarians, among them jellyfish and Hydra (Weissbourd et al. 2021; Pallasdies et al. 2019; Dupre and Yuste 2017). For example, in the case of jellyfish neurons, it might not be feasible to experimentally determine the electrical properties of spike or transmembrane signal transduction. In this case, one has to resort to either qualitative descriptions of the dynamics or first determine the form of nodal dynamics by specialized estimation methods (Ladenbauer et al. 2019; Gonçalves et al. 2020).

Thus, in the absence of GT connectivity data, any single estimated connectivity (EStC) matrix is not really valuable per se, because it cannot be fundamentally validated. Hence, if only the number of nodes and their corresponding activities are known, it might not be sufficient to just estimate connectivity. Therefore, in the absence of GT connectivity data, the task of connectivity estimation must be augmented by a recreation step to obtain knowledge on EC. If the nodal dynamics from this regeneration agrees well with that of the real system, the agreement indicates that our choice of dynamical system and SC faithfully reflects the processes that generated the original time series data. Then, graph theory methods can be used to study the connectivity structure of the biological network of interest (Bassett and Sporns 2017).

In this paper, we are interested in unraveling the network interactions that produce observed node activity time series, which amounts to understanding how the system functions by observing node activity. For a complete understanding, we need to infer the EC of the system. EC can be classically decomposed into SC among the nodes plus a model of the local node dynamics. Hence, the task can be broken down into two steps: (1) infer the SC and (2) define a node model and use it together with the SC to reconstruct the system dynamics. For (1), the different inference approaches need to be validated against GT connectivity data. For (2), they need to be further validated against the observed node activity patterns. In general, SC is represented by sparse directed weighted graphs (Rubinov and Sporns 2010). When equipped with a model for the node dynamics, it is possible to simulate the entire underlying system, which opens access to EC. Hence, obtaining EC, on top of SC, will result in a more complete understanding of the neural system under investigation.

Here, we introduce a new method, lagged-cross-correlation (LCC) analysis, to estimate SCin a wide range of systems. Our method does not make assumptions on the network topology or dynamics. We first focus on mathematical models of neural networks as exemplified by the Hopf model (Deco et al. 2017; Cakan and Obermayer 2020) as a system with known GT connectivity. We show that in sparse networks, LCC-based methods as well as methods combining LCC with derivative-based DDC methods outperform other methods of connectivity estimation, including purely derivative-based methods by themselves. We then systematically study the performance of different algorithms in terms of the quality of network reconstruction as the connection density varies. For sparse connectivity, our LCC methods perform best, whereas for denser connectivities they perform on a par with other methods, such as standard correlation based methods. Next, we ask whether our inferred SC matrices can be used to recreate the observed nodal dynamics, and hence gain insight into EC, to which the answer is affirmative under strict conditions on the random seed of simulated and recreated data. We compare the quality of the data reconstruction when connectivity matrices estimated with different algorithms are used to recreate the system dynamics. Again, we find that our method outperforms derivative-based methods, because it results in higher trace-to-trace correlations. Using our approach, these correlations were

perfect for some of the recreated traces. Finally, we apply our method to biological data from the nematode *C.Elegans*, and show that it performs at least as well as a recently devised computationally more expensive reservoir-based method (Banerjee, Chandra, and Ott 2023).

In addition to providing a systematic validation of different methods in a paradigmatic scenario, our work constitutes a first important step in tackling the problem of validating connectivity in the absence of GT connectivity data. We propose that combining SCestimation with a recreation simulation can serve as a phenomenological system identification step (Brunton, Proctor, and Kutz 2016). Furthermore, this reconstruction step can help to bridge the gap between SC and EC connectivity measures. Indeed, having estimated an SC matrix, it can be used as an input for a forward simulation of a system. If the system dynamics is recreated by this step, then the estimated SC matrix also bears relevance for the EC structure of the system.

2 Materials and Methods

2.1 Methods for connectivity estimation

2.1.1 Lagged Correlation - Method Explanation

Our principal new method, LCC, is based on Fourier transforms. In short, the method estimates the lag at which cross-correlation between two signals is maximized and uses this maximal value of the correlation as an approximation for the strength and directionality of connectivity between two nodes. As such, our method is related to the work by (Cecchi et al. 2007; Schiefer et al. 2018; Gilson et al. 2016), who used lagged covariance or inverse correlation matrices to study the directionality of connections in fMRI experiments.

Cross-correlation between two signals $X(t)$ and $Y(t)$ at lag τ is defined as a measure of similarity between $X(t)$ and a time-shifted version of $Y(t)$, $Y(t+\tau)$. The time shifted version $Y(t+\tau)$ is then used in the integral

$$R_{XY}(\tau) := (X \star Y)(\tau) = \int_{-\infty}^{\infty} \overline{X(t)} Y(t+\tau) dt$$

over time to compute cross-correlation between X and Y. $\overline{X(t)}$ denotes complex conjugation of $X(t)$. This integral calculates how much the signal $X(t)$ correlates with $Y(t+\tau)$ as a function of the time lag τ .

The integral can be computed directly in the time domain. However, calculating the cross-correlation can be computationally demanding, especially when the signals get longer. Therefore, to efficiently calculate the cross-correlation, the method is using Fast Fourier transformations (FFTs) in multiple steps. The first step is to transform the signals into frequency domain using an FFT:

$$F_X(f) = \mathrm{FFT}\{X(t)\} \qquad (1)$$

$$F_Y(f) = \mathrm{FFT}\{Y(t)\} \qquad (2)$$

To now compute the cross-spectral density, the method uses the power spectral density by multiplying the Fourier transform of the first time-series with the complex conjugate of the Fourier transform of the second time-series:

$$S_{XY}(f) = \overline{F_X}(f)F_Y(f) \qquad (3)$$

Afterwards, to obtain the cross-correlation in the time domain, the inverse of the Fourier transform is calculated by using the Inverse Fast Fourier Transformation (IFFT)

$$R_{XY}(\tau) = \mathrm{IFFT}\{S_{XY}(f)\} \qquad (4)$$

By virtue of the convolution theorem, this inverse is the cross-correlation in the time domain.This calculation outputs a sequence consisting of the cross-correlation of the two time series at different lags. The lag with the highest cross-correlation is determined by finding the index in the sequence at which the correlation is highest, yielding a lag:

$$\tau_{\max} = \mathrm{argmax}\{|R_{XY}(\tau)|\} \qquad (5)$$

By examining $R_{XY}(\tau)$ and $R_{YX}(\tau)$, the method assesses the directionality of the interaction. If $\tau_{\max}$ for $R_{XY}(\tau)$ is smaller than for $R_{YX}(\tau)$ it implies a directional influence from X to Y. Conversely, a smaller $\tau_{\max}$ for $R_{YX}(\tau)$ implies a directional influence from Y to X. Equally large values for $\tau_{\max}$ in both directions with a high correlation suggest a bidirectionality of the influence. Then, the Pearson correlation coefficients are calculated for each neuron pair. With the use of the directionality calculated in the step before, the EStC matrix is filled with the corresponding Pearson correlation coefficients. In summary, this method first used the cross-correlation to determine the directionality of connections, and then estimates connectivity with the Pearson correlation coefficient.

2.1.2 Lagged Correlation in conjunction with DDC

Another method that was explored was to use the LCC method in conjunction with the DDC algorithm (Chen, Rosen, and Sejnowski 2022). For this approach, two matrices obtained with the LCC approach and the DDC approach were compared against each other and only values over a certain threshold, typically set at 0.1, that were present in both estimated connectivities were used. The DDC approach was used in its linear form, using a linear estimation non-linearity.

2.1.3 Cut dataset

For the comparison, the methods described above were also applied to a dataset which was sampled at every 15th time point (i.e. every 15th timepoint was kept in the time series), resulting in a naive smoothing of the data. This was done to reduce signal-to-noise ratio, which can sometimes be beneficial when working with neuronal data.

2.1.4 Thresholded methods

For comparison, some methods were also thresholded after estimation. These are marked as “thresholded” in the comparison matrices. For this, a threshold of 0.1 was chosen, which clamps any value lower than the set threshold to 0. Values above the threshold are not changed. This can be helpful when dealing with sparse connectivity matrices, as these naturally contain a lot of zeros.

2.1.5 Transfer Entropy

For the transfer entropy estimated connectivity, the implementation *(https://github.com/netsiphd/netrd*) from *netrd* (McCabe et al. 2021) was used. This allowed for calculation of the naive transfer entropy (NTE).

2.1.6 Normalized methods

To further compare the methods, some data sets were normalized prior to applying the method. The data was therefore z-scored across all the traces prior to the connectivity estimation. This potentially helps the method in estimating connectivity if for example the data has varying magnitudes of activity.

2.2 Neuron model

To validate our methods, a Hopf neuron model was implemented using the neurolib library (Cakan, Jajcay, and Obermayer 2023) Python package. The Hopf model is a non-linear dynamical system that simulates, through a bifurcation mechanism, the activity of a neuronal network. This model is particularly good at capturing the point of transition between stable and oscillatory neuron states, leading to a good representation of the physiological transition of firing states (Deco et al. 2017). In this implementation, each neuron’s dynamics are governed by the Stuart-Landau equations, which form the dynamics underlying the Hopf model. Specifically, the two-dimensional dynamics (one variable each for real and imaginary parts of the complex variable whose time evolution is given by the Stuart-Landau equation) of each neuron are described by:

$$\dot{x}_i = (a - x_i^2 - y_i^2)x_i - wy_i + K_{gl} \sum_{l=1}^{n} C_{il}(x_l(t - \tau_{il}) - x_i) + \text{external input} \qquad (6)$$

$$\dot{y}_i = (a - x_i^2 - y_i^2)y_i + wx_i + K_{gl} \sum_{l=1}^{n} C_{il}(y_l(t - \tau_{il}) - y_i) + \text{external input} \qquad (7)$$

where x and y represent the real and imaginary parts of the neuron’s state and $\dot{x}$ and $\dot{y}$ denote a time derivative. The Hopf bifurcation parameter a and the natural frequency parameter w describe the oscillator. K_{gl} is the global coupling strength of the model that is multiplied with the connectivity of the current neuron to adjacent neurons from the connectivity matrix C. The difference term $x_l(t - \tau_{il}) - x_i$ (and similarly for y) describe the effect of a diffusive coupling scheme. There additionally is a distance-dependent time delay between neurons. The external inputs are inputs independent of the network dynamics, such as noise in the form of Ornstein-Uhlenbeck processes. These are used for representing sensory input or other forms of stimuli that the neurons might experience in a biological context.

In the context of validating the connectivity estimation method, employing this Hopf model presents a robust approach. First, the dynamics of the model is rich and non-linear. Secondly, by generating a GT before running the model, we can later validate the estimation method by comparing the GT matrix to the estimated matrix. Different measures, such as accuracy and correlation between estimated and GT connectivity matrices can be used to quantitatively validate the estimation. This ensures that its efficacy is assessed in a robust way by application of multiple statistical measures.

2.2.1 Data Generation

The neural network model was structured as a network consisting of 10 neurons using the Hopf neuronal model described in Sec.2.2. The connections between these neurons were chosen based on a pairwise probability p. For method validation, this parameter was initially set to 0.1 and then increased to 0.9. Our application of this method resulted in a directed, non-symmetric graph (directed random graph): if there is a connection from node i to j, it does not mean that there also is a connection from j to i. p=10% is the probability that two given neurons in the network will be connected, but only in one direction, so that the probability of bidirectional connections in our adjacency matrix is low, but increases with increasing p. The total number of connections is given by $pN(N-1)$ where N is the number of neurons. The diagonal is also filled with zeros, such that there exists no connection from a neuron to itself. Such sparse connectivity was a made choice, aimed to emulate sparse connectivity patterns observed in biological neural networks. For each neuron, the dynamics were governed by the Hopf model differential equations, for which we used the default parameters. The generated dataset at the end consisted of a time series of 20000 time steps per neuron. All statistics were computed on the entire length of these time series.

2.2.2 Forward simulation

To validate the accuracy of the estimated connectivities, we used the same method as described in Sec.2.1. First, we estimated the connectivity matrix as described in Sec.2.1. Then, the dynamical system in Sec.2.2 was simulated forward in time with the EStC matrix. Applied with the same initial parameters as the original data, this approach was used to recreate the time-series data. Furthemore, the seed for the noisy inputs in Eq.6 and 7 had to be set exactly the same, as otherwise the background noise and corresponding oscillation would not match up in the end. Furthermore, this recreation of the node dynamics allows for easy validation of the used method.

2.2.3 Distance similarity measure

To further validate the accuracy of the EStC, a distance similarity measure is used. To compare two connectivity matrices X and Y a measure is computed through a series of mathematical operations. Firstly, the absolute distance between the corresponding elements of X and Y is calculated as:

$$\mathrm{dist}_{ij} = |X_{ij} - Y_{ij}| \qquad (8).$$

This difference is then normalized by the absolute values of non-zero elements in X yielding the normalized distance:

$$\mathrm{dist}_{ij} = \frac{\mathrm{dist}_{ij}}{|X_{ij}|},\ \mathrm{for} X_{ij} \neq 0 \qquad (9)$$

This is then transformed into their absolute complements:

$$\mathrm{dist_absolute}_{ij} = 1 - |\mathrm{dist}_{ij}| \qquad (10)$$

quantifying the similarity between X and Y. Negative values in $\mathrm{dist_absolute}$ are set to zero:

$$\mathrm{dist_absolute}_{ij} = \max(0, \mathrm{dist_absolute}_{ij}) \qquad (11)$$

ensuring that only meaningful values are used. We then define a performance measure as follows:

$$\mathrm{performance} = \frac{1}{n^2} \sum_{i=1}^{n} \sum_{j=1}^{n} \mathrm{dist_absolute}_{ij} \qquad (12)$$

This performance is then used as the distance similarity measure in Fig. 3. A higher value of this measure indicates a better agreement between the EC and GT connectivity matrices, whereas a lower value indicates a larger discrepancy. The distance similarity measure is complementary to the Pearson correlation coefficient, as it captures the absolute difference between the matrices. Therefore, using both measures can provide a more comprehensive evaluation of the performance of our algorithms.

3 Results

3.1 Comparison of methods for connectivity estimation in the Hopf model

We started to use our LCC method (Sec.2.1.1) on small networks of the Hopf model (Sec.2.2) to estimate a known GT connectivity. The GT connectivity was chosen to be a sparse matrix with p=10% connectivity. Our method demonstrated a remarkably high accuracy in estimating the underlying connectivity of the Hopf neuron model using the proposed LCC approach (Fig. 1). The Pearson correlation coefficient between the EC and GT connectivities was calculated to be 0.97 (Fig. 1B), which is exceptionally high. For the same data, the DDC algorithm only showed a corresponding correlation coefficient of 0.52 (Fig. 1C).

Using the LCC method, we could sometimes even estimate connectivities that near-perfectly recovered the GT connectivity (meaning that both the presence or absence and in case of a connection also its strength was correctly estimated), as shown in Fig. 1 A and B.

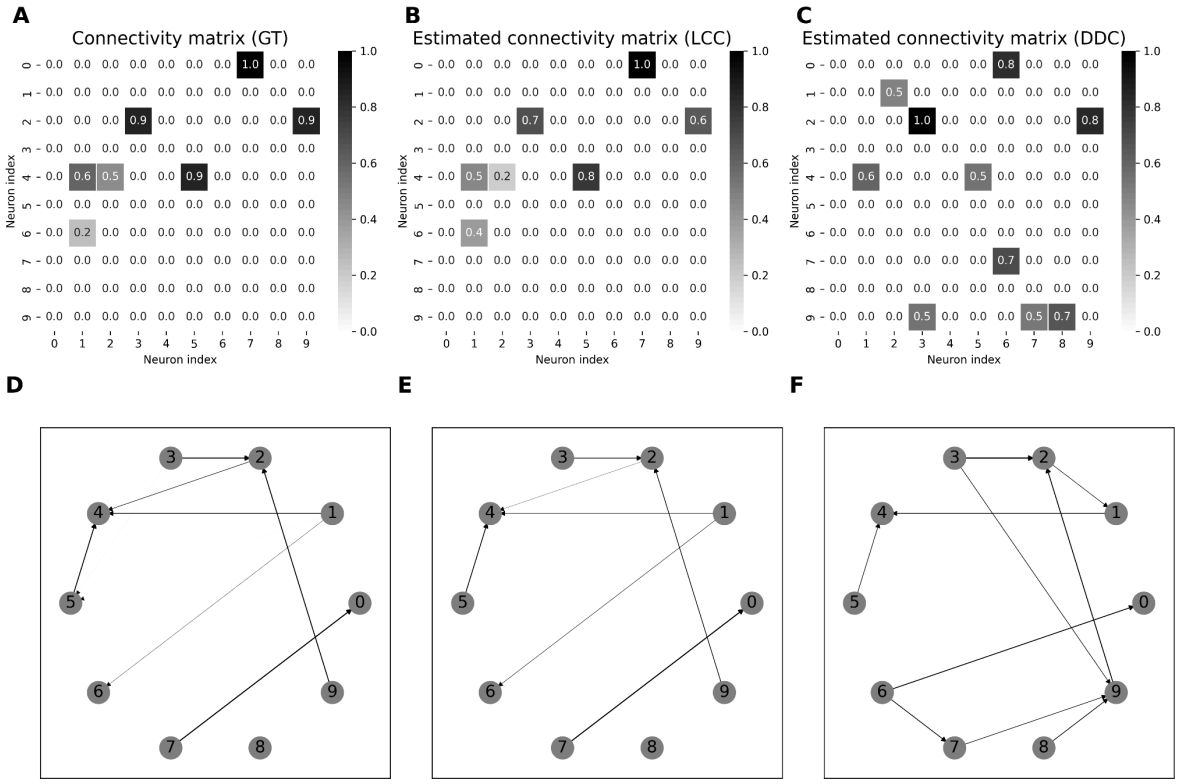

Figure 1: Comparison of GT and EStC connectivity matrices in a Hopf Model using the thresholded LCC method, connection probability p = 0.1. Panel A shows the original GT connectivity matrix used to generate time-series data with a Hopf model. Panel B shows the EStC matrix using the LCC method (Sec. 2.1.1). Nearly perfect agreement with the GT matrix is observed. The Pearson correlation coefficient is 0.97 between the matrices. Panel C shows the EStC matrix using the DDC method. This method finds some of the right connections, but the connection strength is off-target, and there are many false positives. This leads to a lower correlation (0.52) between the EStC matrix computed with DDC and the GT connectivity matrix (panel A). Due to rounding to one decimal place, 2 values are not shown in the GT connectivity matrix, this leads to 7 (instead of 9) visible connections in the GT connectivity matrix. Panels D-F show the corresponding network graphs for the GT connectivity (Panel D) and the two estimated graphs for LCC (Panel E) and DDC (Panel F). Two connections in Panel D are hardly visible because of their small numerical value.

Figure 2: Pearson correlation coefficients between EStC matrices estimated using different methods and the GT matrix of the given model parameters for a Hopf model. The number of neurons is set to $N = 10$, while the connection probability p increases from 0.1 to 0.9. For small p ($p < 0.4$), LCC in conjunction with the DDC algorithm yielded the highest correlation out of all the methods used (row 2). At higher values of p normalized LCC (row 1) as well as correlation-based methods (rows 5 and 7) and cut LCC (row 6) perform best. LCC and thresholded LCC (rows 2 and 9) perform worse, but still better than DDC-based methods (rows 2,3 and 9). The value per cell is calculated by taking the mean of the correlation in 100 simulations rounded to 2 decimal places.

Which method performs best for this synthetic Hopf model data set? To answer this question, we systematically studied the performance of different algorithms in networks of fixed size and increasing connectivity (Fig. 2). Performance here was quantified with the Pearson correlation coefficient between GT connectivity matrices and their reconstructed EStC counterparts. Our first result is that the newly proposed LCC method in conjunction with the DDC algorithm (Sec.2.1.2) showed the best performance for small values of p (Fig. 2, row 2). By just using the LCC method (Fig. 2, row 4), the correlation coefficient was not decreasing as rapidly, compared to DDC, with increasing p. For all values of p, DDC alone (Fig. 2, row 3) did not perform well compared to the other methods. Moreover, the LCC method performed nearly equally as well on its own (Fig. 2, row 4) as in conjunction with DDC (Fig. 2, row 2), while either normalized or cut LCC (Fig. 2, rows 1 and 6) and methods using standard correlation (Fig. 2, rows 5 and 7) were the best methods for higher values of p. The agreement with the GT connectivity, however, was much lower than for small values of p for all methods. The thresholded DDC method (Fig. 2, row 8) generally yielded a higher correlation for larger values of p compared to DDC alone (Fig. 2, row 3), as this method produces true zeros in the EStC matrix leading to a higher correlation

especially in the simulations with sparse connectivity. LCC thresholded (Fig. 2, row 9) performed nearly as well as LCC (Fig. 2, row 4) for all values of p. Again, the methods that did not use DDC worked better than those use DDC for larger values of p. As a comparison with more traditional methods for estimating directed connectivity, we added NTE (Schreiber 2000; McCabe et al. 2021). In Fig. 2, row 10 it can be seen that while NTE performs worse than LCC for small p, but increasingly better at large p, where it even slightly outperforms LCC (Fig. 2, row 4). However, TE still performs slightly worse than standard correlation-based measures (Fig.2, rows 5 and 7) and also than normalized LCC (Fig. 2, row 1) and cut LCC (Fig. 2, row 6). For all values of p, TE performs better than methods involving DDC (Fig. 2, rows 3 and 8).

Summarizing, we have shown that our novel and comparatively simple method and its variants can successfully reconstruct model connectivity from dynamics. In general, all methods perform worse when p is increased. Thus, our methods work best for sparse networks, such as those found in the cnidarian *Hydra* (Dupre and Yuste 2017), the nematode *C. Elegans* or jellyfish (Bielecki et al. 2023). In higher animals like zebrafish, drosophila, mice and rats, networks may be less sparse (Svara et al. 2022; Winding et al. 2023).

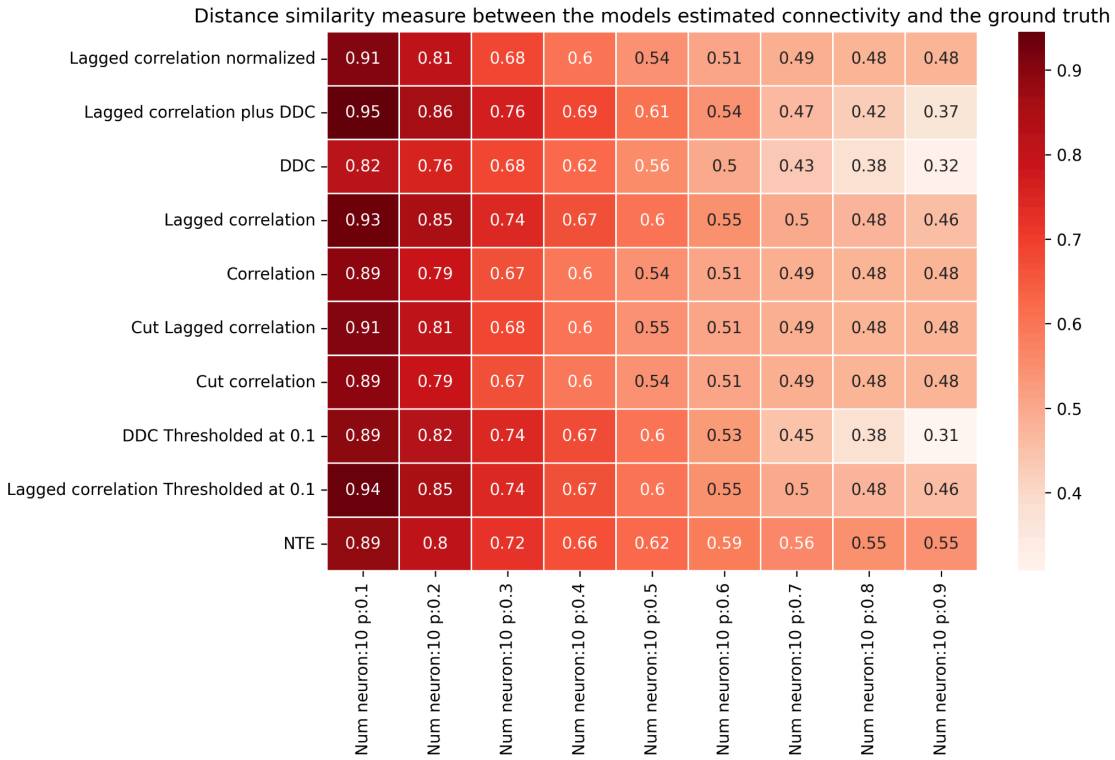

Figure 3: Distance Similarity Measure (Sec.2.2.3) between GT and EStC Matrices in Neural Network Modeling. This heat map provides a visual representation of the distance similarity measure between the GT and EStC matrices, crucial in evaluating the accuracy of connectivity estimation in neural network models. The similarity measure is quantified on a scale from 0 to 1, where a value of 1 indicates perfect correspondence (identical matrices), and a value of 0 denotes complete dissimilarity (opposite values in the matrices). The value per cell is calculated by taking the mean of the correlation in 100 simulations rounded to 2 decimal places.

To further evaluate the performance of our methods in estimating connectivity, a distance similarity measure (Sec.2.2.3) is used to show how much the EStC matrices deviate from the GT matrices in terms of the differences of individual matrix entries. This helps with comparing methods, as the correlation compares the correctness of the pattern (Fig. 2), while the distance similarity shows the deviation from the original values, in which slight deviations do not affect the metric as much. This means false positive zero values are not punished as much as in the correlation, as the distance is not much affected by this. This leads to higher accuracy values in Fig. 3. The distance similarity measure (Fig. 3), which again is highest for LCC in conjunction with the DDC algorithm for small p ($p < 0.5$), but also for LCC thresholded (Fig. 3, row 9). The LCC algorithm alone (Fig. 3, row 4) performs nearly equally as well without relying on the computationally more expensive DDC algorithm. This, in addition to the higher correlation coefficients for higher connection probability values- compared to the DDC method- leads us to suggest that LCC presents a method that is performing well in estimating Hopf network connectivities. Of note, the standard measures Correlation and Cut Correlation also perform well for higher values of p (Fig. 3, rows 5 and 7), but perform slightly worse than methods using LCC for smaller p. TE (Fig. 3, row 10) outperforms all other measures for higher values of p ($p>0.4$), albeit not by a large margin. As in Fig. 2, the thresholded LCC method (Fig. 3, row 9) performs nearly equally as well as its non-thresholded counterpart (Fig. 3, row 4). In contrast, the thresholded DDC algorithm (Fig. 3 row 8) performs better than its non-thresholded counterpart (Fig. 3, row 3) at lower values of p; the performance at higher p is comparable. The thresholded DDC method performs better than its non-thresholded counterpart, because the thresholding allows for zeros. In contrast, the DDC algorithm can produce negative values, which do not make sense for our model's connectivity that can only range from 0 to 1.

3.2 Regeneration of node dynamics with different EStC matrices

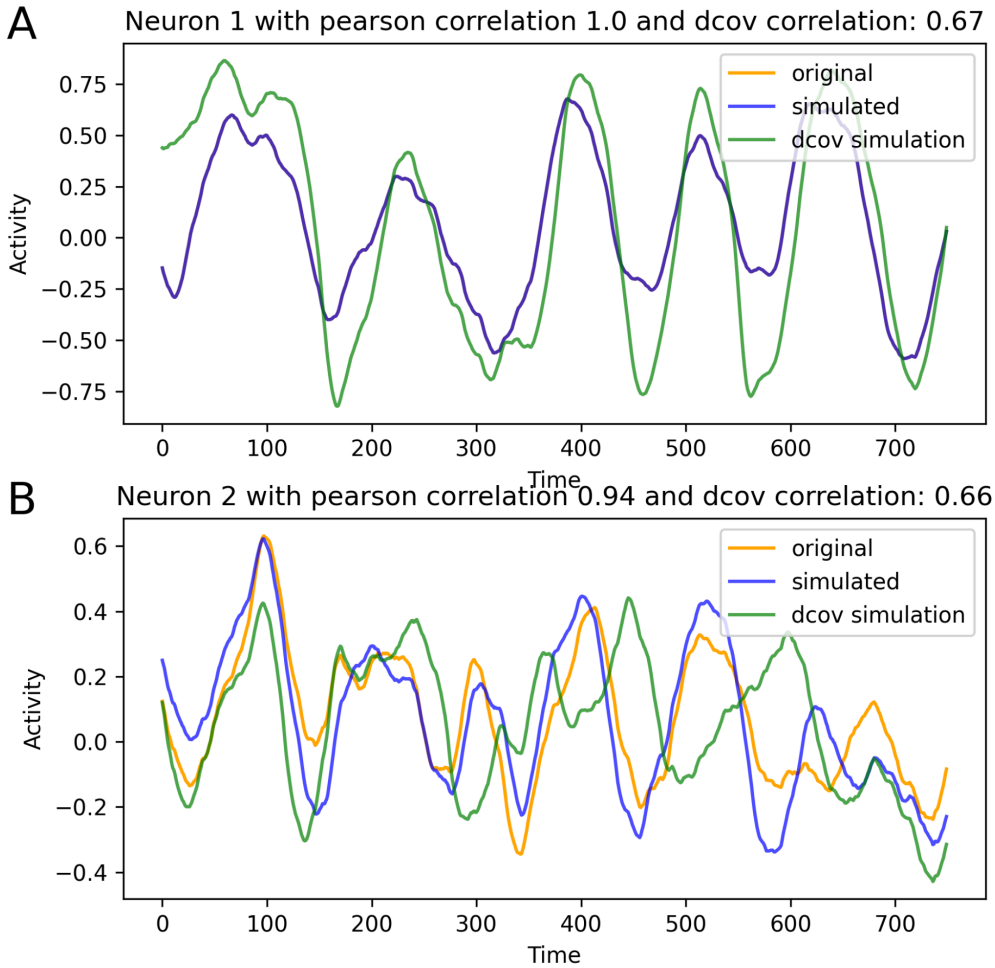

Figure 4: Comparison of the simulated data using the GT, LCC- estimated connectivity and DDC-estimated connectivity for a network of 10 neurons with p=0.1. Orange: Original simulation. Blue: Recreation of the dynamics using LCC to obtain EC. Green: Recreation of the dynamics using DDC to obtain EC. Panel A: Results for neuron 1. Panel B: Results for neuron 2. The real part of the corresponding Stuart-Landau equation (Eq.6 and 7) is here plotted as activity. The same random seed for the stochastic input to each node's internal dynamics was used in original and both recreated trajectories.

We then asked the question whether connectivity matrices estimated with our methods can be used to recreate the time series of each node (see Sec.2.2.1 for details). Fig. 4 presents a comprehensive comparison between two distinct methods for estimating neuronal connectivity: LCC (traces labeled 'simulated') and DDC (traces labeled 'dcov simulation)'. The analysis is conducted on simulated data from two neurons. Each plot

illustrates three trajectories: the original data (‘original’, GT simulation), the simulated data reconstructed using the LCC method (‘simulated’), and the simulated data reconstructed via the DDC approach (‘dcov simulation’). The GT matrix is shown in Fig. 1A. Fig. 4A shows neuron activity traces where the LCC method correctly identified all the connections as 0, therefore yielding almost perfectly similar traces. This is a special case for our sparse connectivities, as some neurons are not connected to any others (cf. row for neuron 1 in Fig. 1A) and if this is estimated successfully, the correlation is 1 as all the activity is derived from the underlying oscillation (Eq.6 and 7). In other words, the coupling matrix does not influence the dynamics of this particular neuron 1. Panel B shows neuron activity traces for a neuron in the model where the LCC method estimated a very similar connectivity compared to the GT connectivity This results in a trace-trace correlation equal to 0.94. Using DDC to determine the EStC, the trace-trace correlation drops to 0.66. The correlation of the simulated data using LCC with the original GT data is 0.98 over all the 10 neurons, while the correlation using DDC only reaches 0.51. The main reason for the lower correlation when using DDC is that the EStC matrices are less sparse: incorrectly identified connections of a neuron can affect other neurons indirectly connected with it. A typical scenario is that of secondary connected neurons. If the EStC matrix wrongly results in a connection between neuron 1 and neuron 2, and neuron 2 is correctly connected to neuron 3, then neuron 1 wrongly influences the activity of neuron 3. This lowers the correlation of neuron 3’s reconstructed activity with the GT activity. Thus, false positives in the connectivity matrix result in low correlations between reconstructed and GT dynamics. In summary, we have shown that if the EStC is close to the GT matrix, then the node dynamics can be accurately recreated in a forward simulation of the system. For good quantitative agreement of the regenerated traces with the original traces, evidenced by high trace-trace correlations, the EStC matrix must nearly perfectly resemble the GT matrix. Small differences in the two matrices result in larger differences between the traces.

3.3 Application to *C.elegans* data

Finally, to test our methods using empirical biological data, we applied the LCC method to a dataset of the 8 most active motor neurons of the nematode *C. Elegans* (Kato et al. 2015). This is the same data as used by the recent study (Banerjee, Chandra, and Ott 2023), in which a reservoir computing approach was used to infer connectivity. Applying our LCC method to the same data results in a higher count of true positive values (21 compared to 14) for the same amount of false positives (4 FP), as shown in Fig.5. Note that the GT data in this case is binary, so a suitable binarization threshold is needed after application of the LCC method.

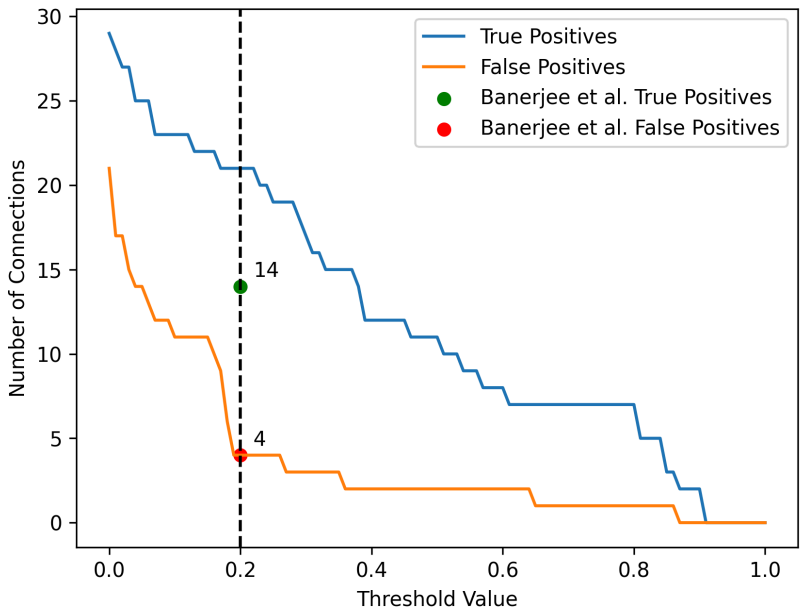

Figure 5: Performance analysis of LCC method binarization thresholds on true and false positive connections in C. Elegans time series data. The figure depicts the relationship between the cut threshold, seen on the x-axis and the corresponding counts of TPs and FPs. A threshold value of 0.2 (black dashed line) was chosen to have comparability against the methods proposed by Banerjee et al. (Banerjee, Chandra, and Ott 2023).

The number of true positives (TPs) and false positives (FPs) as a function of the binarization threshold value is shown in Fig. 5. The threshold clamps all the values above it to 1 and everything below to 0 to obtain a binary connectivity matrix. Considering the slope of the curve for FPs in Fig. 5, another way to choose the threshold value is given by placing it just before the number of FPs increases sharply, which is just below 0.2. Concretely, in decreasing the threshold from 0.2 to 0, FP increases from 4 to 21, whereas TP increases much less steeply from 21 to 29. The green and red dots in the plot show the TP (14) and FP (4) values from the reservoir computing shown in the recent study (Banerjee, Chandra, and Ott 2023).

This shows that the suggested method using LCC does not only perform well when applied to simulated Hopf neuron model data, but also actual biologically recorded *C. Elegans* neuron activity data. Comparing our method with that obtained by a reservoir computing method, we have achieved a higher count of TP (21 compared to 14) while keeping the same number of FPs (4). In summary, our results so far propose a newly more generalized method for connectivity estimation for neuron activity data, which is crucial for research in neuronal connectomics. For comparison, our best connectivity matrix with a threshold of 0.2 is shown in Fig. 6A, together with the GT connectivity in Fig. 6B.

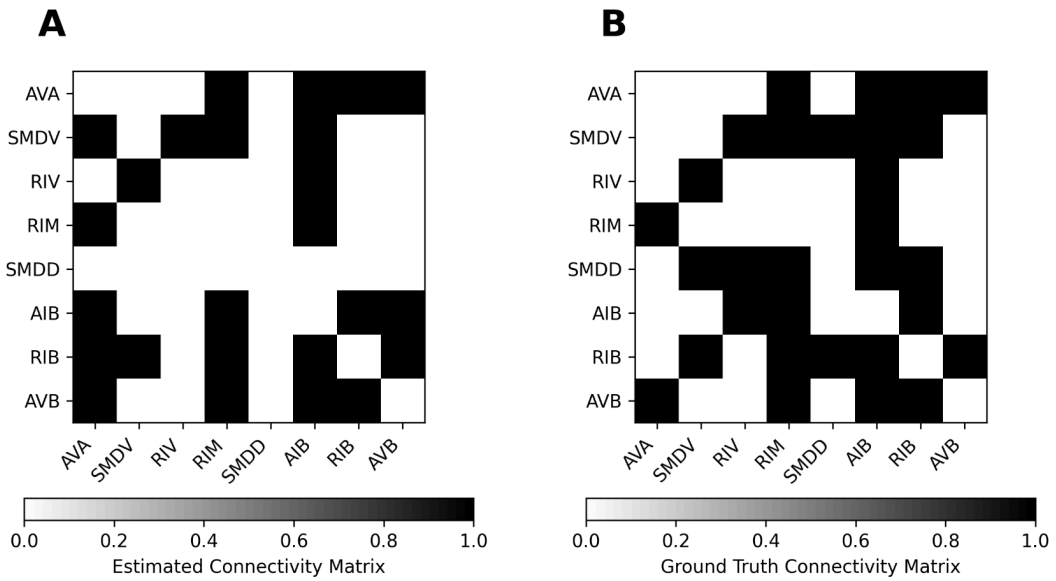

Figure 6: Comparison of the GT and EStC connectivity for the 8 most active motor neurons of *C. Elegans.* Panel A shows the EStC matrix using the LCC method with a binarization: values above a threshold of 0.2 are clipped to 1 and values below it to 0. Panel B shows the GT connectivity matrix for the 8 most active motor neurons of C.elegans. Comparing the matrices shows a TP count of 21 and an FP count of 4 as indicated in Fig.5 by the black dashed line.

4 Discussion

The goal of this paper was to develop and validate a new method to estimate the effective network connectivity from observed neural activity data, without relying on prior knowledge of the network topology or dynamics. The goal was pursued in two steps, by first inferring the underlying structural connectivity of the system and comparing it to ground truth structural connectivity data, and second using the structural network in conjunction with a model of the local node dynamics to recreate the observed activity patterns.

Our approach for estimating the SC of the system, LCC, is based on lagged correlation analysis, which first estimates the directionality and then the strength of connections between neurons by finding the optimal time lag that maximizes the cross-correlation between the neurons. Furthermore, this method was used in conjunction with the recently published method of DDC, which uses a derivative-based approach to infer EC. We then applied the different methods to synthetic data generated by a Hopf neuron model to validate the methods prior to testing them on real data. The estimation algorithm was finally applied to biological data derived from the eight most active motor neurons of the nematode *C. Elegans*, a well-known model organism for connectomics with a known GT connectivity matrix.

Our results showed that our new method performed well in estimating the connectivity of the Hopf models, especially when the networks were sparsely connected. For higher connectivities, our results suggest that the well-established measure of transfer entropy slightly outperforms our LCC methods. Given that TE methods are known to be computationally demanding (Shao et al. 2014), and given that TE does not perform much better than LCC (Fig.3), we are confident that our LCC method provides a good tradeoff between accuracy and computational cost, at least for the examples considered here. We further compared our method with other established methods such as pairwise correlation, and DDC. For sparse connectivity, our method outperformed others in terms of Pearson correlation and distance similarity between the EStC and GT connectivity matrices. For denser connectivity, when more bidirectional connections are present, standard correlation-based methods show comparable performance to our LCC-based methods. This is expected, because, in denser networks, estimating the directionality of pairwise connections is more difficult due to the influence of third nodes. Therefore, our LCC-based methods do not outperform standard correlation-based measures in this case. In all scenarios, our LCC-based methods performed better than DDC-based methods. As such, our LCC method, like established correlation-based methods for inferring connectivity, is parameter free. The thresholded LCC method needs one input parameter, namely the threshold, above which potential links are considered connections.

We then demonstrated that our method could be used to recreate the system dynamics by using the estimated SC matrix as the new connectivity for a Hopf model with the same input parameters as used for the activity data generated using the GT connectivity. We found that this produced high correlations between the original and recreated data, indicating that our approach captured the essential dynamical features encoded in the connectivity matrix. Thus, our approach could be used to bridge the gap between SC and EC in this scenario.

Finally, we applied our method to *C. Elegans* activity data and found that it could infer the connectivity better than currently used methods such as the recently proposed reservoir computing approach by (Banerjee, Chandra, and Ott 2023). Our method achieved a higher count of true positives at the same amount of false positives compared to the reservoir computing approach. This finding suggests that our method is robust and reliable in estimating connectivity from biological data, which often struggles with noise and low temporal resolution.

Our work contributes to the field of computational connectomics by proposing a new method to estimate SC, which can then be used in conjunction with a model of local node dynamics to derive the effective connectivity of the system and recreate the observed activity patterns. The two-step approach is surprisingly simple, which endows it with generality and flexibility, allowing for an application to different types of neural activity data. Another advantage of our LCC method is that we need no prior assumptions about network topology or node dynamics, in contrast to methods involving DDC, which assume linear or non-linear node dynamics of a specific type and no external input.

How can our methods be improved? First, it remains to be shown how the methods perform in large networks that are not sparse. Surprisingly, initial simulations show that

increasing the network size does not impact performance negatively, whereas increasing sparsity usually does.

Second, dealing with the problem of partially observed dynamics remains a challenge. Even if our LCC method does not need the assumption of full observability, it certainly aids performance when network dynamics are generated intrinsically, and not driven from outside. When strong external driving is present, our approach can still result in errors arising from confounder motifs (see Introduction). To alleviate this problem, either the network needs to be observed more completely or external input must be taken into account in the formulation of the estimation algorithm.

Another area for future work would be to combine our approach with methods from simulation-based inference (Cranmer, Brehmer, and Louppe 2020; Boelts et al. 2023; Gonçalves et al. 2020). Concretely, in the absence of GT connectivity, our LCC-based methods could be used to estimate structural connectivity. This connectivity could be fed into a mechanistic model of neural activity serving as a system generating the empirical activity data. Its parameters, but not the connectivity itself, could then be estimated using methods from simulation-based inference. When these parameters are found, the system identification loop is closed: the data can be recreated by forward-simulating the estimated model with the estimated SC matrix. Alternatively, for small networks, dynamical causal modeling (DCM) (K. J. Friston, Harrison, and Penny 2003) could be used. In case of good agreement of the simulated, recreated traces with the original activity data, this step closes the loop and bridges the gap between SC and EC. In other words, a system that generates the data has been successfully identified. These exciting questions are left for future work.

5 Conflict of Interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

6 Author Contributions

Conceptualization	NL, WB, LK, JB, CCH
Data curation	NL
Formal analysis	NL, WB
Funding acquisition	CCH
Investigation	NL, WB, CCH, LK

Methodology	NL, WB, LK
Project administration	WB, CCH
Resources	NL, CCH
Software	NL
Supervision	WB, CCH
Validation	NL, WB
Visualization	NL
Writing – original draft	NL, WB
Writing – review & editing	NL, WB, LK, JB, CCH

7 Funding

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 434434223 – SFB 1461.

8 Acknowledgments

The authors would like to thank Kayson Fakhar, Alexander Schaum, Fatemeh Hadaeghi and Heike Siebert for useful comments.

9 Data Availability Statement

Code used in this study can be found under www.github.com/UKE-ICNS/LCC .

Bibliography

Andalman, Aaron S., Vanessa M. Burns, Matthew Lovett-Barron, Michael Broxton, Ben Poole, Samuel J. Yang, Logan Grosenick, et al. 2019. "Neuronal Dynamics Regulating Brain and Behavioral State Transitions." *Cell* 177 (4): 970-985.e20. https://doi.org/10.1016/j.cell.2019.02.037.

Banerjee, Amitava, Sarthak Chandra, and Edward Ott. 2023. "Network Inference from Short, Noisy, Low Time-Resolution, Partial Measurements: Application to C. Elegans Neuronal Calcium Dynamics." *Proceedings of the National Academy of Sciences* 120 (12): e2216030120. https://doi.org/10.1073/pnas.2216030120.

Bassett, Danielle S., and Olaf Sporns. 2017. "Network Neuroscience." *Nature Neuroscience* 20 (3): 353–64. https://doi.org/10.1038/nn.4502.

Bielecki, Jan, Sofie Katrine Dam Nielsen, Gösta Nachman, and Anders Garm. 2023. "Associative Learning in the Box Jellyfish Tripedalia Cystophora." *Current Biology* 33 (19): 4150-4159.e5. https://doi.org/10.1016/j.cub.2023.08.056.

Boelts, Jan, Philipp Harth, Richard Gao, Daniel Udvary, Felipe Yáñez, Daniel Baum, Hans-Christian Hege, Marcel Oberlaender, and Jakob H. Macke. 2023. "Simulation-Based Inference for Efficient Identification of Generative Models in Computational Connectomics." *PLOS Computational Biology* 19 (9): e1011406. https://doi.org/10.1371/journal.pcbi.1011406.

Brunton, Steven L., Joshua L. Proctor, and J. Nathan Kutz. 2016. "Discovering Governing Equations from Data by Sparse Identification of Nonlinear Dynamical Systems." *Proceedings of the National Academy of Sciences* 113 (15): 3932–37. https://doi.org/10.1073/pnas.1517384113.

Bullmore, Ed, and Olaf Sporns. 2009. "Complex Brain Networks: Graph Theoretical Analysis of Structural and Functional Systems." *Nature Reviews Neuroscience* 10 (3): 186–98. https://doi.org/10.1038/nrn2575.

Cakan, Caglar, Nikola Jajcay, and Klaus Obermayer. 2023. "Neurolib: A Simulation Framework for Whole-Brain Neural Mass Modeling." *Cognitive Computation* 15 (4): 1132–52. https://doi.org/10.1007/s12559-021-09931-9.

Cakan, Caglar, and Klaus Obermayer. 2020. "Biophysically Grounded Mean-Field Models of Neural Populations under Electrical Stimulation." *PLOS Computational Biology* 16 (4): e1007822. https://doi.org/10.1371/journal.pcbi.1007822.

Cecchi, Guillermo A., A. Ravishankar Rao, Maria V. Centeno, Marwan Baliki, A. Vania Apkarian, and Dante R. Chialvo. 2007. "Identifying Directed Links in Large Scale Functional Networks: Application to Brain fMRI." *BMC Cell Biology* 8 (1): S5. https://doi.org/10.1186/1471-2121-8-S1-S5.

Chen, Yusi, Burke Q. Rosen, and Terrence J. Sejnowski. 2022. "Dynamical Differential Covariance Recovers Directional Network Structure in Multiscale Neural Systems." *Proceedings of the National Academy of Sciences* 119 (24): e2117234119. https://doi.org/10.1073/pnas.2117234119.

Cranmer, Kyle, Johann Brehmer, and Gilles Louppe. 2020. "The Frontier of Simulation-Based Inference." *Proceedings of the National Academy of Sciences* 117 (48): 30055–62. https://doi.org/10.1073/pnas.1912789117.

Deco, Gustavo, Morten L. Kringelbach, Viktor K. Jirsa, and Petra Ritter. 2017. "The Dynamics of Resting Fluctuations in the Brain: Metastability and Its Dynamical Cortical Core." *Scientific Reports* 7 (1): 3095. https://doi.org/10.1038/s41598-017-03073-5.

Deistler, Michael, Jakob H. Macke, and Pedro J. Gonçalves. 2022. "Energy-Efficient Network Activity from Disparate Circuit Parameters." *Proceedings of the National Academy of*

Sciences 119 (44): e2207632119. https://doi.org/10.1073/pnas.2207632119.
Dupre, Christophe, and Rafael Yuste. 2017. "Non-Overlapping Neural Networks in Hydra Vulgaris." *Current Biology* 27 (8): 1085–97. https://doi.org/10.1016/j.cub.2017.02.049.
Durstewitz, Daniel, Nicole M. Vittoz, Stan B. Floresco, and Jeremy K. Seamans. 2010. "Abrupt Transitions between Prefrontal Neural Ensemble States Accompany Behavioral Transitions during Rule Learning." *Neuron* 66 (3): 438–48. https://doi.org/10.1016/j.neuron.2010.03.029.
Faisal, A. Aldo, Luc P. J. Selen, and Daniel M. Wolpert. 2008. "Noise in the Nervous System." *Nature Reviews Neuroscience* 9 (4): 292–303. https://doi.org/10.1038/nrn2258.
Fakhar, Kayson, Fatemeh Hadaeghi, and Claus C. Hilgetag. 2022. "Causal Influences Decouple From Their Underlying Network Structure In Echo State Networks." In *2022 International Joint Conference on Neural Networks (IJCNN)*, 1–8. https://doi.org/10.1109/IJCNN55064.2022.9892782.
Fakhar, Kayson, and Claus C. Hilgetag. 2022. "Systematic Perturbation of an Artificial Neural Network: A Step towards Quantifying Causal Contributions in the Brain." *PLOS Computational Biology* 18 (6): e1010250. https://doi.org/10.1371/journal.pcbi.1010250.
Farahani, Farzad V., Waldemar Karwowski, and Nichole R. Lighthall. 2019. "Application of Graph Theory for Identifying Connectivity Patterns in Human Brain Networks: A Systematic Review." *Frontiers in Neuroscience* 13. https://www.frontiersin.org/articles/10.3389/fnins.2019.00585.
Fornito, Alex, Andrew Zalesky, and Edward T. Bullmore, eds. 2016. "Front Matter." In *Fundamentals of Brain Network Analysis*, i–ii. San Diego: Academic Press. https://doi.org/10.1016/B978-0-12-407908-3.09996-9.
Friston, K. J., C. D. Frith, P. F. Liddle, and R. S. J. Frackowiak. 1993. "Functional Connectivity: The Principal-Component Analysis of Large (PET) Data Sets." *Journal of Cerebral Blood Flow & Metabolism* 13 (1): 5–14. https://doi.org/10.1038/jcbfm.1993.4.
Friston, K. J., L. Harrison, and W. Penny. 2003. "Dynamic Causal Modelling." *NeuroImage* 19 (4): 1273–1302. https://doi.org/10.1016/S1053-8119(03)00202-7.
Friston, Karl J. 1994. "Functional and Effective Connectivity in Neuroimaging: A Synthesis." *Human Brain Mapping* 2 (1–2): 56–78. https://doi.org/10.1002/hbm.460020107.
———. 2011. "Functional and Effective Connectivity: A Review." *Brain Connectivity* 1 (1): 13–36. https://doi.org/10.1089/brain.2011.0008.
Gilson, Matthieu, Ruben Moreno-Bote, Adrián Ponce-Alvarez, Petra Ritter, and Gustavo Deco. 2016. "Estimation of Directed Effective Connectivity from fMRI Functional Connectivity Hints at Asymmetries of Cortical Connectome." *PLOS Computational Biology* 12 (3): e1004762. https://doi.org/10.1371/journal.pcbi.1004762.
Gonçalves, Pedro J, Jan-Matthis Lueckmann, Michael Deistler, Marcel Nonnenmacher, Kaan Öcal, Giacomo Bassetto, Chaitanya Chintaluri, et al. 2020. "Training Deep Neural Density Estimators to Identify Mechanistic Models of Neural Dynamics." Edited by John R Huguenard, Timothy O'Leary, and Mark S Goldman. *eLife* 9 (September): e56261. https://doi.org/10.7554/eLife.56261.
Goodwin, Graham Clifford, and Robert L. Payne. 1977. *Dynamic System Identification: Experiment Design and Data Analysis*. Academic Press.
Hagmann, Patric, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J. Honey, Van J. Wedeen, and Olaf Sporns. 2008. "Mapping the Structural Core of Human Cerebral Cortex." *PLOS Biology* 6 (7): e159. https://doi.org/10.1371/journal.pbio.0060159.
Honey, Christopher J., Rolf Kötter, Michael Breakspear, and Olaf Sporns. 2007. "Network Structure of Cerebral Cortex Shapes Functional Connectivity on Multiple Time Scales." *Proceedings of the National Academy of Sciences* 104 (24): 10240–45. https://doi.org/10.1073/pnas.0701519104.

Kato, Saul, Harris S. Kaplan, Tina Schrödel, Susanne Skora, Theodore H. Lindsay, Eviatar Yemini, Shawn Lockery, and Manuel Zimmer. 2015. "Global Brain Dynamics Embed the Motor Command Sequence of Caenorhabditis Elegans." *Cell* 163 (3): 656–69. https://doi.org/10.1016/j.cell.2015.09.034.

Klink, P. Christiaan, Jean-François Aubry, Vincent P. Ferrera, Andrew S. Fox, Sean Froudist-Walsh, Béchir Jarraya, Elisa E. Konofagou, et al. 2021. "Combining Brain Perturbation and Neuroimaging in Non-Human Primates." *NeuroImage* 235 (July): 118017. https://doi.org/10.1016/j.neuroimage.2021.118017.

Ladenbauer, Josef, Sam McKenzie, Daniel Fine English, Olivier Hagens, and Srdjan Ostojic. 2019. "Inferring and Validating Mechanistic Models of Neural Microcircuits Based on Spike-Train Data." *Nature Communications* 10 (1): 4933. https://doi.org/10.1038/s41467-019-12572-0.

Lainscsek, Claudia, Christopher E. Gonzalez, Aaron L. Sampson, Sydney S. Cash, and Terrence J. Sejnowski. 2019. "Causality Detection in Cortical Seizure Dynamics Using Cross-Dynamical Delay Differential Analysis." *Chaos: An Interdisciplinary Journal of Nonlinear Science* 29 (10): 101103. https://doi.org/10.1063/1.5126125.

Maier-Hein, Klaus H., Peter F. Neher, Jean-Christophe Houde, Marc-Alexandre Côté, Eleftherios Garyfallidis, Jidan Zhong, Maxime Chamberland, et al. 2017. "The Challenge of Mapping the Human Connectome Based on Diffusion Tractography." *Nature Communications* 8 (1): 1349. https://doi.org/10.1038/s41467-017-01285-x.

Marder, Eve. 2012. "Neuromodulation of Neuronal Circuits: Back to the Future." *Neuron* 76 (1): 1–11. https://doi.org/10.1016/j.neuron.2012.09.010.

McCabe, Stefan, Leo Torres, Timothy LaRock, Syed Arefinul Haque, Chia-Hung Yang, Harrison Hartle, and Brennan Klein. 2021. "Netrd: A Library for Network Reconstruction and Graph Distances." *Journal of Open Source Software* 6 (62): 2990. https://doi.org/10.21105/joss.02990.

McCormick, David A., Dennis B. Nestvogel, and Biyu J. He. 2020. "Neuromodulation of Brain State and Behavior." *Annual Review of Neuroscience* 43 (1): 391–415. https://doi.org/10.1146/annurev-neuro-100219-105424.

Pallasdies, Fabian, Sven Goedeke, Wilhelm Braun, and Raoul-Martin Memmesheimer. 2019. "From Single Neurons to Behavior in the Jellyfish Aurelia Aurita." Edited by Ronald L Calabrese, Richard Satterlie, and Eva Kanso. *eLife* 8 (December): e50084. https://doi.org/10.7554/eLife.50084.

Prinz, Astrid A., Dirk Bucher, and Eve Marder. 2004. "Similar Network Activity from Disparate Circuit Parameters." *Nature Neuroscience* 7 (12): 1345–52. https://doi.org/10.1038/nn1352.

Randi, Francesco, Anuj K. Sharma, Sophie Dvali, and Andrew M. Leifer. 2023. "Neural Signal Propagation Atlas of Caenorhabditis Elegans." *Nature* 623 (7986): 406–14. https://doi.org/10.1038/s41586-023-06683-4.

Rubinov, Mikail, and Olaf Sporns. 2010. "Complex Network Measures of Brain Connectivity: Uses and Interpretations." *NeuroImage*, Computational Models of the Brain, 52 (3): 1059–69. https://doi.org/10.1016/j.neuroimage.2009.10.003.

Schiefer, Jonathan, Alexander Niederbühl, Volker Pernice, Carolin Lennartz, Jürgen Hennig, Pierre LeVan, and Stefan Rotter. 2018. "From Correlation to Causation: Estimating Effective Connectivity from Zero-Lag Covariances of Brain Signals." *PLOS Computational Biology* 14 (3): e1006056. https://doi.org/10.1371/journal.pcbi.1006056.

Schreiber, Thomas. 2000. "Measuring Information Transfer." *Physical Review Letters* 85 (2): 461–64. https://doi.org/10.1103/PhysRevLett.85.461.

Shandilya, Srinivas Gorur, and Marc Timme. 2011. "Inferring Network Topology from Complex Dynamics." *New Journal of Physics* 13 (1): 013004. https://doi.org/10.1088/1367-2630/13/1/013004.

Shao, Shengjia, Ce Guo, Wayne Luk, and Stephen Weston. 2014. “Accelerating Transfer Entropy Computation.” In *2014 International Conference on Field-Programmable Technology (FPT)*, 60–67. https://doi.org/10.1109/FPT.2014.7082754.

Sporns, Olaf, Giulio Tononi, and Rolf Kötter. 2005. “The Human Connectome: A Structural Description of the Human Brain.” *PLOS Computational Biology* 1 (4): e42. https://doi.org/10.1371/journal.pcbi.0010042.

Svara, Fabian, Dominique Förster, Fumi Kubo, Michał Januszewski, Marco dal Maschio, Philipp J. Schubert, Jörgen Kornfeld, et al. 2022. “Automated Synapse-Level Reconstruction of Neural Circuits in the Larval Zebrafish Brain.” *Nature Methods* 19 (11): 1357–66. https://doi.org/10.1038/s41592-022-01621-0.

Valdes-Sosa, Pedro A., Alard Roebroeck, Jean Daunizeau, and Karl Friston. 2011. “Effective Connectivity: Influence, Causality and Biophysical Modeling.” *NeuroImage* 58 (2): 339–61. https://doi.org/10.1016/j.neuroimage.2011.03.058.

Weissbourd, Brandon, Tsuyoshi Momose, Aditya Nair, Ann Kennedy, Bridgett Hunt, and David J. Anderson. 2021. “A Genetically Tractable Jellyfish Model for Systems and Evolutionary Neuroscience.” *Cell* 184 (24): 5854-5868.e20. https://doi.org/10.1016/j.cell.2021.10.021.

White, J. G., E. Southgate, J. N. Thomson, and S. Brenner. 1986. “The Structure of the Nervous System of the Nematode Caenorhabditis Elegans.” *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 314 (1165): 1–340. https://doi.org/10.1098/rstb.1986.0056.

Winding, Michael, Benjamin D. Pedigo, Christopher L. Barnes, Heather G. Patsolic, Youngser Park, Tom Kazimiers, Akira Fushiki, et al. 2023. “The Connectome of an Insect Brain.” *Science (New York, N.Y.)* 379 (6636): eadd9330. https://doi.org/10.1126/science.add9330.

Yan, Gang, Petra E. Vértes, Emma K. Towlson, Yee Lian Chew, Denise S. Walker, William R. Schafer, and Albert-László Barabási. 2017. “Network Control Principles Predict Neuron Function in the Caenorhabditis Elegans Connectome.” *Nature* 550 (7677): 519–23. https://doi.org/10.1038/nature24056.

Yeh, Chun-Hung, Derek K. Jones, Xiaoyun Liang, Maxime Descoteaux, and Alan Connelly. 2021. “Mapping Structural Connectivity Using Diffusion MRI: Challenges and Opportunities.” *Journal of Magnetic Resonance Imaging* 53 (6): 1666–82. https://doi.org/10.1002/jmri.27188.