Estimating the Temporal Evolution of Synaptic Weights from Dynamic Functional Connectivity

Marco Celotto[1,2,3(✉)], Stefan Lemke[1,4], and Stefano Panzeri[1,3]

[1] Neural Computation Laboratory, Center for Neuroscience and Cognitive Systems, Istituto Italiano di Tecnologia, Rovereto, Italy
marco.celotto@iit.it
[2] Department of Pharmacy and Biotechnology, University of Bologna, Bologna, Italy
[3] Department of Excellence for Neural Information Processing, Center for Molecular Neurobiology (ZMNH), University Medical Center Hamburg-Eppendorf (UKE), Hamburg, Germany
s.panzeri@uke.de
[4] Department of Cell Biology and Physiology, University of North Carolina, Chapel Hill, USA
stefan.lemke@unc.edu

Abstract. How to capture the temporal evolution of synaptic weights from measures of dynamic functional connectivity (DFC) between the activity of different simultaneously recorded neurons is an important and open problem in systems neuroscience. To address this issue, we first simulated models of recurrent neural networks of spiking neurons that had a spike-timing-dependent plasticity mechanism generating time-varying synaptic and functional coupling. We then used these simulations to test analytical approaches that relate dynamic functional connectivity to time-varying synaptic connectivity. We investigated how to use different measures of directed DFC, such as cross-covariance and transfer entropy, to build algorithms that infer how synaptic weights evolve over time. We found that, while both cross-covariance and transfer entropy provide robust estimates of structural connectivity and communication delays, cross-covariance better captures the evolution of synaptic weights over time. We also established how leveraging estimates of connectivity derived from entire simulated recordings could further boost the estimation of time-varying synaptic weights from the DFC. These results provide useful information to estimate accurately time variations of synaptic strength from spiking activity measures.

Keywords: Dynamic functional connectivity · Spiking neural network · Communication delay · Transfer entropy · Cross-covariance

1 Introduction

Neurons in biological networks are connected by directed, plastic synapses. Neurons are sparsely connected and the identity, the strength, and the

M. Mahmud et al. (Eds.): BI 2022, LNAI 13406, pp. 3–14, 2022
https://doi.org/10.1007/978-3-031-15037-1_1

communication delay of the connections between cells-pairs determine the network dynamics [6,15,16]. Importantly, the strength of each synapse can change over different time scales - ranging from tenths of milliseconds to days - due to processes including synaptic potentiation and depression [2]. Such changes in synaptic weights are thought to be neural-activity dependent and driven by Hebbian mechanisms of plasticity such as spike-timing dependent-plasticity (STDP).

Many electrophysiological *in vivo* experiments record simultaneously the spiking activity of several neurons within a network, but without the ability to measure directly synaptic activity. Robust methods to estimate synaptic weights and how they evolve over time from functional measurements of neural activity are thus critical to investigate several neuroscientific questions. For example, sleep is thought to play an essential role in synaptic homeostasis and memory formation. Several theories and experimental findings support the idea that specific features of non-REM sleep might contribute to the up- and down-scaling of synaptic weights [23]. Experimentally, it has been shown that the nesting between spindles and slow oscillations can increase the dynamic functional connectivity (DFC), measured as peaks of cross-correlation between pairs of putatively connected cells, over the temporal range of minutes [14]. However, the corresponding synaptic changes over the same time span are difficult to characterize. In general, it remains unclear how changes in DFC measures relate to the temporal dynamics of synaptic weights in spiking neural networks.

Previous works investigating the relationship between functional connectivity measures and ground truth synaptic connectivity have often utilized the Izhikevich network [11] as a reasonably realistic model of a cortical spiking neural network [9,18]. These studies highlighted that bivariate connectivity measures, such as cross-covariance and transfer entropy, can provide robust estimates of the underlying directed connectivity in simulated networks. However, they did not examine the temporal evolution of functional and structural connectivity within spiking networks incorporating STDP. Here, we examined the performance of several different DFC methods in estimating the temporal dynamics of synaptic weights (termed dynamic structural connectivity or DSC) from up to 90 min of spiking activity in simulated spiking networks with STDP. We first determined the performance of DFC measures in inferring static properties of the simulated networks (such as pairwise synaptic connectivity and the associated communication delays). We then applied these measures with a sliding window approach to compute DFC and quantify its relationship with DSC. We found, that, among all tested DFC measures, the cross-covariance better captured the evolution of synaptic weights over time. Importantly, we also established how to use the information obtained from the static, time-averaged analysis of the network, to enhance the estimate of DSC from DFC.

2 Simulated Spiking Network and Inference Pipeline

To investigate the relationship between DSC and DFC, we simulated a spiking neural network in which the strength of synaptic weights changed over time

according to an STDP rule. We then compared the performance of different functional connectivity measures in estimating both the ground truth structure of the network (i.e. which pairs of neurons were connected, their communication lag, and the type of synapse), and how the strengths of the synaptic weights changed over time (Fig. 1). We simulated a spiking network of N = 100 neurons in which the dynamics of each neuron was described using the Izhikevich neuron model [10]. This model has a good tradeoff between biological plausibility and computational efficiency. The structure of the network was set by Izhikevich [11] to mimic the connectivity of a real population of cortical neurons (Fig. 1A). 80% of neurons in the network were excitatory and 20% were inhibitory. Excitatory neurons were randomly connected to 10 postsynaptic neurons which could be either excitatory or inhibitory (800 excitatory synapses in total). Each excitatory synapse had a random communication delay (δ) whose value was uniformly distributed between 1 and 20 ms and was constant over time. Inhibitory neurons were randomly connected to 10 postsynaptic excitatory neurons (200 inhibitory synapses), therefore no inhibitory-to-inhibitory (I-I) connections were present in the network. The lack of I-I synapses caused the average firing rate of excitatory neurons (5.11 ± 0.03 Hz) to be lower than the one of inhibitory neurons (8.23 ± 0.04 Hz). Inhibitory connections had a communication delay of 1 ms. The simulation ran with 1 ms temporal precision for a duration decided by the user. During the simulation, the strength of excitatory synapses - which were all initialized to the same, positive, value - changed dynamically due to an STDP rule: when a presynaptic neuron i fired before a postsynaptic neuron j the strength of the synapse from i to j (w_{ij}) was strengthened, on the other hand when j fired before i w_{ij} got weaker (Fig. 1B). The decay time of the STDP rule was $\tau = 20$ ms and synaptic weights were updated every 1 s with a memory factor which made the weights change, on average, over the timescale of 1–2 min (obtained measuring the synaptic weights autocorrelation, not shown).

We used different measures to compute the static and dynamic functional connectivity of the network from the spiking activity (Fig. 1C). Such measures were all directed (meaning that, for each pair of neurons, they could take different values in the two directions) and allowed computing the strength of communication for different delays (δ). When computing static functional connectivity, we used data from the whole simulated recording to compute a single connectivity value for each pair of neurons (i, j). We computed all connectivity measures with δ ranging from 1 to 50 ms then, for each pair, we determined the static functional connectivity (w_{ij}) as the maximum connectivity value across delays. We selected the communication delay (δ_{ij}) as the lag that maximized the functional connectivity. Calling $f_{ij}(\delta)$ the generic measure of functional connectivity, then: $w_{ij} = max_{\delta}(f_{ij}(\delta))$ and $\delta_{ij} = argmax_{\delta}(f_{ij}(\delta))$. By taking the top percentile of connectivity values for each measure we obtained sparse static networks (Fig. 1D). If the measure f was signed we could also infer whether a synapse was excitatory or inhibitory. Then, we used a sliding window approach to compute, for each measure, the DFC of all the synapses that were inferred as present (Fig. 1E). We exploited the static measures of communication delay

between pairs to compute delay-consistent DFC and then evaluated the performance of the different measures in recovering the ground-truth dynamics of synaptic weights.

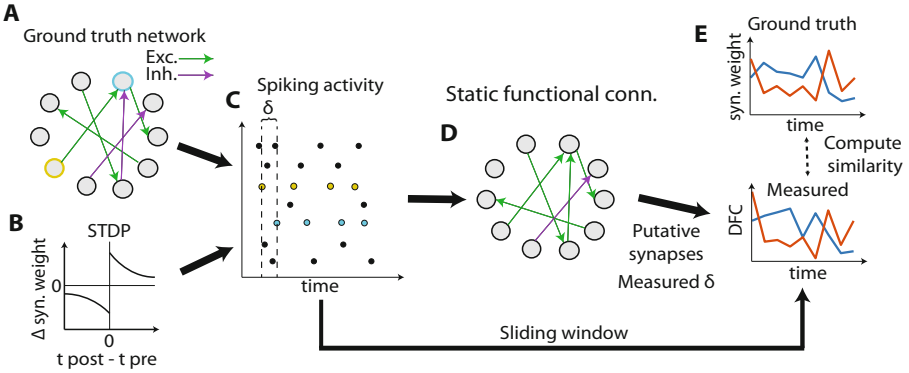

Fig. 1. Graphical depiction of the method. A) Structural connectivity of the simulated network for N = 10 neurons. Synaptic weights could be either excitatory (green) or inhibitory (purple). Excitatory connections had randomly distributed communication delays. B) The strength of the synaptic weights changed over time due to STDP. C) Structural and biophysical properties of the network determined the spiking activity of the neural population. D) Static functional connectivity was measured from spiking activity. E) Dynamic functional connectivity was measured from activity, also leveraging on the inferred static connectivity of the network. (Color figure online)

3 Inferring the Presence of Synapses

We tested the performance of different measures of functional connectivity in estimating the presence of synapses from spiking activity. Two of these measures were based on Pearson correlation, which is commonly used to estimate the connectivity between pairs of neurons [3,9,14]. The first method was normalized cross-correlation ($XCorr$):

$$XCorr_{ij}(\delta) = \frac{E[i_{t-\delta} j_t]}{\sigma_i \sigma_j} \tag{1}$$

where i_t and $j_{t'}$ are the binary values of the spike trains from neurons i and j at times t and t', and the expected value was computed across time. σ_i and σ_j are standard deviations of the spike trains of neurons i and j, respectively.

The second method was the normalized cross-covariance ($XCov$), which is insensitive to correlations in the average firing rate due to subtraction of the average activity value from the spike trains before computing the correlation:

$$XCov_{ij}(\delta) = \frac{E[(i_{t-\delta} - \bar{i})(j_t - \bar{j})]}{\sigma_i \sigma_j} \tag{2}$$

Here $\overline{i}$ and $\overline{j}$ are the average firing rates of neurons i and j, respectively.

Additionally, we computed the functional connectivity using two variants of the information-theoretic measure of information transfer known as transfer entropy [8,21], a measure that has been successfully used to characterize time-dependent changes in recurrent connectivity between mass signals [1]. Transfer entropy has the theoretical advantage - with respect to correlation measures - of being assumption-free in terms of the joint probability distribution of the lagged activity of neuron i and j. This also means that transfer entropy does not assume that the interactions between neurons are linear. Additionally, this measure respects the Wiener-Granger causality principle of causal communication by conditioning the information between the past of the emitter and the present of the receiver neuron on the past activity of the receiver neuron. Our first implementation of transfer entropy uses single time-points statistics to build the probability distribution of lagged neural activity. We refer to this implementation as TE:

$$TE_{ij}(\delta) = I(i_{t-\delta}; j_t | j_{t-1}) = \sum p(i_{t-\delta}, j_t, j_{t-1}) \log_2 \frac{p(j_t | i_{t-\delta}, j_{t-1})}{p(j_t | j_{t-1})} \quad (3)$$

where $p(i_{t-\delta}, j_t, j_{t-1})$ is the joint probability distribution of the present state of the receiver neuron j_t, its past lagged by one time step j_{t-1} and the past state of the emitter neuron lagged by δ time steps $i_{t-\delta}$. The sum occurs over all the $(i_{t-\delta}, j_t, j_{t-1})$ triplets of events in the probability space. The probability distribution is sampled across time. The lag of the receiver past is set to -1 since it has been proven to be theoretically optimal for determining real communication delays [24].

The second implementation of transfer entropy uses multidimensional pasts of the emitter and the receiver neuron to consider the possible relevance of time windows longer than 1 ms when transmitting information. Using the terminology of [9] we refer to this measure as Higher Order Transfer Entropy ($HOTE$):

$$HOTE_{ij}(\delta) = I(i^{(k)}_{t-\delta}; j_t | j^{(l)}_{t-1}) = \sum p(i^{(k)}_{t-\delta}, j_t, j^{(l)}_{t-1}) \log_2 \frac{p(j_t | i^{(k)}_{t-\delta}, j^{(l)}_{t-1})}{p(j_t | j^{(l)}_{t-1})} \quad (4)$$

where k and l are the dimensions of the past activity of the emitter and the receiver neuron i and j, respectively. For the analysis reported in this paper we set $k = l = 5$ ms.

We computed these four functional connectivity measures between all pairs of neurons in the network and estimated the communication strength and delay for each pair as described in the previous section. We then evaluated the performance of the different metrics in determining the presence or absence of synapses between pairs of neurons, varying the threshold probability of connectivity strength incrementally from 0 to 1 in steps of 0.01. Since the two classes of present and absent synapses were unbalanced (only 10% of all the possible synapses were present in the network) we used precision-recall (PR) curves to study the performance in this classification task [4] (Fig. 2A). Calling TP, FP

and FN the number of true positive, false positive and false negative inferred synapses, respectively, we have that $precision = \frac{TP}{TP+FP}$ and $recall = \frac{TP}{TP+FN}$. Therefore, if for a given measure the two distributions of present and absent links were perfectly separable, we would get that for $recall = 1$ also $precision = 1$. On the other hand, a random classifier would always have a precision equal to the ratio of synapses present in the model (10%, dashed line in Fig. 2A) for each recall value.

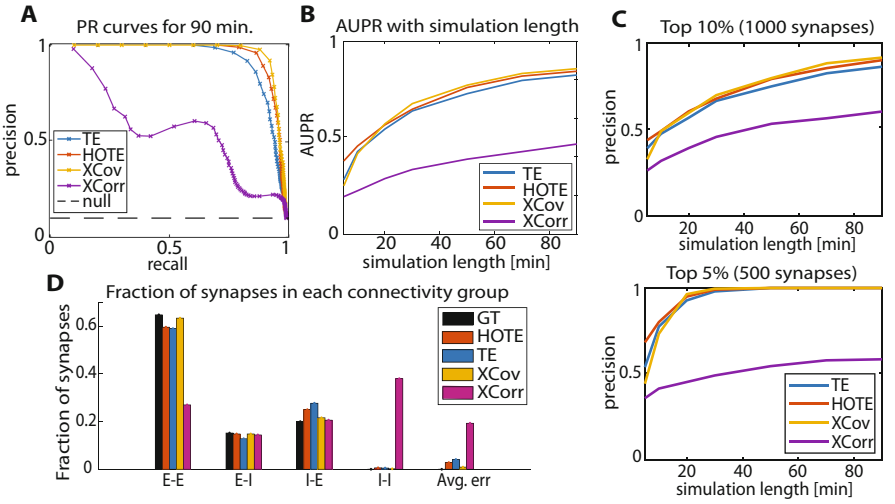

Fig. 2. Performance of functional connectivity measures in estimating structural connectivity. A) Precision-recall (PR) curves computed from 90 min of simulated activity for TE, HOTE, XCov and XCorr. Each point is one percentile of the distribution of functional connectivity values across pairs. B) AUPR trend with simulation length (length ranges from 5 to 90 min). C) Comparison of precision in identifying connected pairs with simulation lengths, for top 10th (1000 pairs) and top 5th (500 pairs) percentiles of each measure's distribution. D) Fraction of pairs belonging to each group of synapses, from 90 min simulation and using the top 10th percentile of connections. GT = ground truth. (Color figure online)

After 90 min of simulation, XCov, TE and HOTE all performed well in the classification task, having a PR curve whose shape approached the optimal one. Among these three measures, XCov showed the best PR curve and TE the worst one. XCorr, on the other hand, performed poorly, with a PR curve far from optimal. The area under the precision-recall curve (AUPR) is a useful metric to summarize the goodness of a PR curve; a perfect classifier has an AUPR equal to one. We computed how AUPR scales with simulation length for different measures. This analysis confirmed that XCov and HOTE were the best metrics in evaluating which links were present for long recordings, while HOTE worked better than XCov and TE for recording shorter than 10 min (Fig. 2B). We measured how the precision of the different measures scaled with the simulation time

for the top 10th and top 5th percentile of inferred synapses. For the top 10th percentile (i.e. 1000 inferred synapses, which equals the ground truth number of connections) we found that the maximum precision in the classification was obtained with XCov, which topped at 92% for 90 min of simulated recording (Fig. 2C top). With a more conservative threshold of the top 5th percentile of connections (i.e. half of the true total number), we captured the top 500 real connections after 30 min of simulation (Fig. 2C bottom) for all measures but XCorr. To investigate why XCorr performance was so poor when compared to the other measures, we computed the fraction of links inferred by each measure as the top 10th percentile of synapses in the four subgroups of excitatory-to-excitatory (E-E), excitatory-to-inhibitory (E-I), inhibitory-to-excitatory (I-E) and inhibitory-to-inhibitory (I-I) synapses (Fig. 2D). XCov performed best in determining the correct fraction of synapses belonging to each group, while XCorr overestimated the number of I-I connections and underestimated the number of E-E connections. This behavior of XCorr is due to the aforementioned differences in average firing rate between inhibitory and excitatory neurons, with a higher firing rate for inhibitory neurons, as XCorr is sensitive to the correlation between average firing rates. Given the poor performance of XCorr in estimating the presence of synapses, we discarded it in the following analyses.

4 Inferring Synapse Type and Communication Delay

We studied how, for each ground truth synapse, different functional connectivity measures performed in inferring whether the synapse was excitatory or inhibitory, and the value of the communication delay of that pair of neurons.

We could not use information-theoretic measures to infer whether synapses were excitatory or inhibitory as these measures are only positively defined. Therefore, we only examined the XCov performance on this excitatory/inhibitory classification task. We classified a connection as excitatory and inhibitory based on XCov value, with positive correlation values assigned as excitatory connections and negative correlation values as inhibitory connections. After 90 min of recording XCov could reliably separate excitatory and inhibitory synapses (Fig. 3A). We found that the performance of the classifier increased with recording time for both the excitatory and the inhibitory class (Fig. 3B).

We also compared how functional connectivity measures performed in inferring ground truth communication delays. After 90 min of simulation, all measures estimated delays with a correlation across synapses that was above 0.85 (see Fig. 3C for the relationship between the ground truth delays and those estimated using XCov - on the top - and using HOTE - on the bottom). The trend of the correlation between ground truth and estimated delays with simulation lengths was approximately linear in the explored range (Fig. 3D). Nonetheless, HOTE estimated the delays more precisely than XCov and TE. After 90 min of simulation, HOTE had an average delay error, measured as the absolute value of the difference between ground truth and inferred delay, below 1 ms. On the other hand, XCov and TE showed a systematic error in the delay estimation of approximately 2 ms (see Fig. 3C and Fig. 3E).

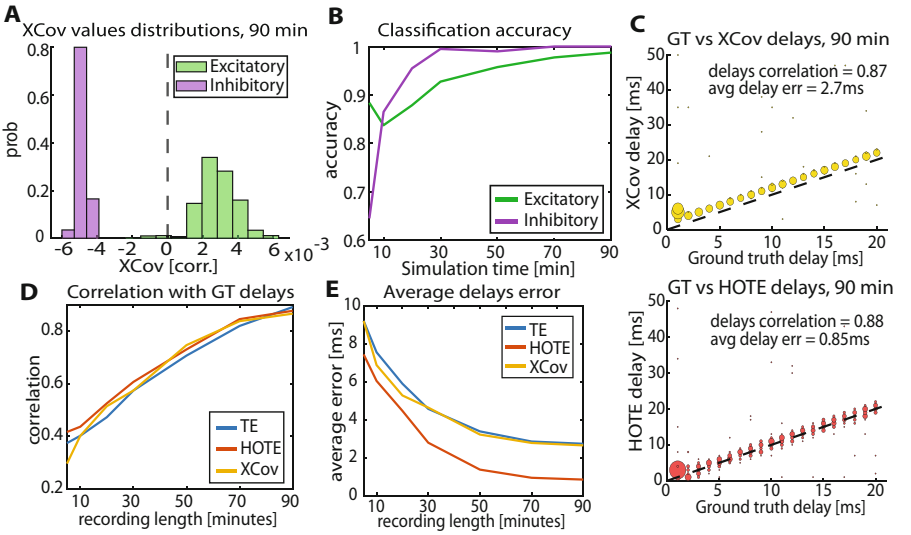

Fig. 3. Performance of the measures in estimating connection type and delays. A) Distributions of functional connectivity values measured using XCov for excitatory (green) and inhibitory (purple) cells. B) Performance of a classifier in identifying excitatory and inhibitory synapses with simulation length. The decision boundary of the classifier was set to $XCov = 0$. C) Scatter plots of real and estimated delays across cell pairs using XCov (top) and HOTE (bottom). The size of the markers is proportional to the number of pairs having that specific combination of ground truth and estimated delay. The dashed line is the identity line $x = y$. Black dots far from the identity line correspond to pairs of measured and real delays that occurred only once. D) Correlation between ground truth and estimated delays with simulation length. E) Average error in delay estimation with simulation length. (Color figure online)

5 Relationship Between Dynamic Functional Connectivity and the Temporal Evolution of Synaptic Weights

Finally, we investigated how the ground truth evolution of the synaptic weights, that is the DSC, related to the measured DFC. We computed DFC using a sliding window approach. We first selected a size for the sliding window T and then shifted it through the simulated recording in steps of length T. We computed DFC only for pairs of neurons that were putatively connected, which we selected as the top 5th percentile of links for each measure after 90 min of simulation (Fig. 1C), and only at the communication delay that we measured for each pair (Fig. 3C). Moreover, we computed DFC only for excitatory synapses since the inhibitory ones had a constant synaptic weight in the simulated network. We calculated the across-time correlation between DFC and DSC for all synapses to quantify the performance of each functional connectivity measure in estimating

the DSC. To do this, we averaged the DSC over windows of width T, so that the number of DSC and DFC samples over time were matched.

In Fig. 4A we show the DSC (top left), the DFC computed using TE (top right), HOTE (bottom left) and XCov (bottom right) for three example synapses and $T = 10\,\text{min}$. It is visible that, while all measures work reasonably well in tracking how the strength of the gray and the green synapses change over time, TE and HOTE fail in quantifying the temporal evolution of the brown synapse. We found that, on average, DFC computed via XCov correlates with DSC better than the DFC computed via TE or HOTE (Fig. 4B). In particular, while DFC computed via TE and HOTE had a high temporal correlation with DSC (above 0.7) for the majority of synapses, their distributions showed a large tail of synapses whose correlation between DSC and DFC was distributed around zero (such as the brown one in Fig. 4A). For XCov, the number of synapses whose DSC was poorly estimated decreased rapidly with the correlation strength, and the average correlation was 0.82 (Fig. 4B, right). Therefore, the DFC computed using XCov outperformed the one obtained from TE and HOTE in inferring the simulated changes of the synaptic weights over time.

We then studied how the across-time correlation between DSC and DFC depends on the width of the sliding window T. The correlation between DFC and DSC increased with the window size, reaching a plateau around $T = 5\,\text{min}$ (Fig. 4C, left). Below $T = 5\,\text{min}$ the correlation dropped due to the limited sample size used to compute DFC, manifesting a tradeoff between the temporal precision of the DFC measures (T) and their performance in estimating DSC. We repeated the same analysis without keeping the delay consistent when computing DFC but simply taking the maximum connectivity value across delays (between 1 and 50 ms) for each window (Fig. 4C, middle). When not keeping the delay consistent with the previously measured one, the correlation between DSC and DFC dropped substantially. For sizes of the sliding window lower than $T = 5\,\text{min}$, the advantage of keeping a consistent delay was particularly evident, with a boost in the correlation between DSC and DFC larger than 0.2 (Fig. 4C, right). This result showed a clear benefit in leveraging estimates of delay derived from entire simulated recordings when inferring DSC from DFC.

6 Discussion

We studied how different measures of functional connectivity can be used to infer the static and dynamic properties of synapses from spiking activity in a simulated neural network. This problem is of relevance because in many *in vivo* experiments only spiking activity is measured, but it is important to also infer the changes in synaptic connectivity to understand the evolution of the neural network under study. We addressed the problem at the level of simulated recordings with single-neuron cellular resolution. As such, our approach differs from and complements other studies of DFC at the level of mass neural activity [7], which lack the ability to resolve interactions between pairs of individual neurons.

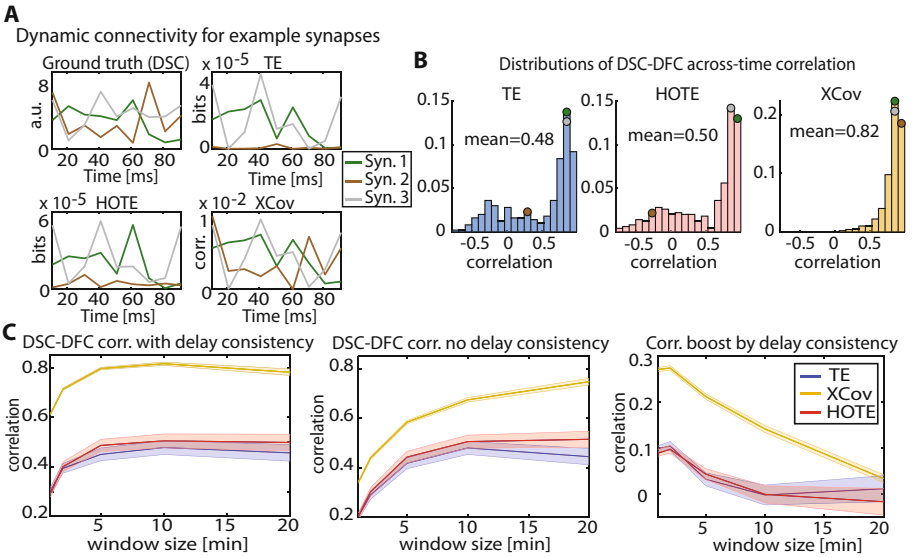

Fig. 4. Relationship between dynamic structural and functional connectivity. A) Dynamic connectivity for 3 example synapses, $T = 10$ min. Top left: ground truth dynamics of synaptic weights (DSC). Top right: Transfer entropy DFC. Bottom left: HOTE DFC. Bottom right: Cross-covariance DFC. B) Distribution of the across-time correlation coefficients between DSC and DFC, $T = 10$ min. Left: Transfer entropy. Middle: HOTE. Right: Cross-covariance. Colored dots show where the synapses in panel A are in the correlation distributions. C) Average correlation between DSC and DFC over time for different sizes of the moving window. Shaded areas are SEM across synapses. Left: DFC keeping delay consistency (i.e. measures computed only at previously estimated delay); Middle: DFC without delay consistency; Right: Boost in correlation between DFC and DSC when keeping delay consistency (difference between left and middle panels). (Color figure online)

Consistent with previous studies, we found that among the considered functional connectivity measures, XCov and HOTE performed best in identifying which pairs of neurons were connected [9]. Cross-covariance could also reliably classify excitatory and inhibitory synapses, while HOTE was the best measure in recovering the ground-truth communication delay between neurons. Cross-covariance performed best in inferring DSC, with an across-time correlation above 0.8 between DFC and DSC for sliding window sizes larger than 5 min. We also found that, when computing DFC, keeping the communication delay consistent with the one obtained from the static network analysis boosted the relationship between DFC and DSC, especially for moving windows shorter than 5 min. It is possible that this correlation boost by keeping the delay constant is because considering delays that differ from the ground truth one enhances the detection of spurious correlations. This specifically holds under the assumption that communication delays are constant in the recording period as is the case

of our spiking network. Such spurious correlations might possibly be induced by other neurons in the population connected with a lag to both neurons in a putatively connected pair.

The present study has limitations that we plan to address in future works. First of all, it will be important to validate the DFC measures on more biologically realistic simulated neural networks presenting global oscillations, correlated inputs to neurons (mimicking sensory perception and motion), and more heterogeneity in the firing rates and in the average synaptic weights over time. Indeed, such effects could act as confounders of the relationship between DFC and DSC [19] or could require refined null hypotheses based on permutation tests to assess the presence of links. In the model we also assumed that (i) the communication delays are constant and (ii) no synapses are formed or eliminated over time. (i) assumes that the main parameters determining the conductance velocity of action potentials (e.g. axons diameters and myelin levels) are approximately constant over time scales of tens of minutes. Experimental finding suggest that this assumption is reasonable, especially in adult mice where the formation of new myelin occurs in the range of weeks [17]. Assumption (ii) is more delicate since in mice it has been shown that, especially during sleep, dendritic spines can be formed and eliminated within hours [25]. It will be important to investigate how much we can relax these hypotheses while still exploiting the knowledge obtained from the static network inference. Moreover, we plan to test the performance of other bivariate (e.g. Granger causality) and, especially, other multivariate measures for estimating DSC. These measures might include using Granger Causality estimates based on Generalized Linear Models (GLMs) [5,13,22] and maximum entropy models [12,20]. Such multivariate measures could be useful e.g. to alleviate the effect of confounders such as common inputs.

To conclude, here we lay down foundations for relating dynamic functional connectivity to the temporal evolution of synaptic weights in spiking neural networks. The results obtained here provide a benchmark for further improving methodologies that infer DSC from DFC.

References

1. Besserve, M., Lowe, S.C., Logothetis, N.K., Schölkopf, B., Panzeri, S.: Shifts of gamma phase across primary visual cortical sites reflect dynamic stimulus-modulated information transfer. PLoS Biol. **13**(9), e1002257 (2015)
2. Citri, A., Malenka, R.C.: Synaptic plasticity: multiple forms, functions, and mechanisms. Neuropsychopharmacology **33**(1), 18–41 (2007)
3. Cutts, C.S., Eglen, S.J.: Detecting pairwise correlations in spike trains: an objective comparison of methods and application to the study of retinal waves. J. Neurosci. **34**(43), 14288–14303 (2014)
4. Davis, J., Goadrich, M.: The relationship between Precision-Recall and ROC curves. In: Proceedings of the 23rd International Conference on Machine Learning - ICML 2006, pp. 233–240 (2006)
5. Francis, N.A., Mukherjee, S., Koçillari, L., Panzeri, S., Babadi, B., Kanold, P.O.: Sequential transmission of task-relevant information in cortical neuronal networks. Cell Rep. **39**(9), 110878 (2022)

6. Ganguli, S., Sompolinsky, H.: Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. Ann. Rev. Neurosci. **35**(1), 485–508 (2012)
7. Hindriks, R., et al.: Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? Neuroimage **127**, 242–256 (2016)
8. Hlavackovaschindler, K., Palus, M., Vejmelka, M., Bhattacharya, J.: Causality detection based on information-theoretic approaches in time series analysis. Phys. Rep. **441**(1), 1–46 (2007)
9. Ito, S., Hansen, M.E., Heiland, R., Lumsdaine, A., Litke, A.M., Beggs, J.M.: Extending transfer entropy improves identification of effective connectivity in a spiking cortical network model. PLoS ONE **6**(11), e27431 (2011)
10. Izhikevich, E.: Simple model of spiking neurons. IEEE Trans. Neural Netw. **14**(6), 1569–1572 (2003)
11. Izhikevich, E.: Polychronization: computation with spikes. Neural Comput. **18**(2), 245–282 (2006)
12. Jaynes, E.T.: Information theory and statistical mechanics. Phys. Rev. **106**(4), 620–630 (1957)
13. Kobayashi, R., et al.: Reconstructing neuronal circuitry from parallel spike trains. Nat. Commun. **10**(1), 4468 (2019)
14. Lemke, S.M., Ramanathan, D.S., Darevksy, D., Egert, D., Berke, J.D., Ganguly, K.: Coupling between motor cortex and striatum increases during sleep over long-term skill learning. eLife **10**, e64303 (2021)
15. Mastrogiuseppe, F., Ostojic, S.: Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. Neuron **99**(3), 609–623 (2018)
16. Ostojic, S., Brunel, N., Hakim, V.: How connectivity, background activity, and synaptic properties shape the cross-correlation between spike trains. J. Neurosci. **29**(33), 10234–10253 (2009)
17. Pan, S., Mayoral, S.R., Choi, H.S., Chan, J.R., Kheirbek, M.A.: Preservation of a remote fear memory requires new myelin formation. Nat. Neurosci. **23**(4), 487–499 (2020)
18. Pastore, V.P., Massobrio, P., Godjoski, A., Martinoia, S.: Identification of excitatory-inhibitory links and network topology in large-scale neuronal assemblies from multi-electrode recordings. PLoS Comput. Biol. **14**(8), e1006381 (2018)
19. Satterthwaite, T.D., et al.: An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. Neuroimage **64**, 240–256 (2013)
20. Schneidman, E., Berry, M.J., Segev, R., Bialek, W.: Weak pairwise correlations imply strongly correlated network states in a neural population. Nature **440**(7087), 1007–1012 (2006)
21. Schreiber, T.: Measuring information transfer. Phys. Rev. Lett. **85**(2), 461–464 (2000)
22. Sheikhattar, A., et al.: Extracting neuronal functional network dynamics via adaptive Granger causality analysis. Proc. Natl. Acad. Sci. U.S.A. **115**(17), E3869–E3878 (2018)
23. Tononi, G., Cirelli, C.: Sleep and the price of plasticity: from synaptic and cellular homeostasis to memory consolidation and integration. Neuron **81**(1), 12–34 (2014)
24. Wibral, M., et al.: Measuring information-transfer delays. PLoS ONE **8**(2), e55809 (2013)
25. Yang, G., Lai, C.S.W., Cichon, J., Ma, L., Li, W., Gan, W.B.: Sleep promotes branch-specific formation of dendritic spines after learning. Science **344**(6188), 1173–1178 (2014)