

# What does functional connectivity tell us about the behaviorally-functional connectivity of a multifunctional neural circuit?

Eduardo J. Izquierdo, Madhavun Candadai

Cognitive Science Program, Indiana University  
Luddy School of Informatics, Computing, and Engineering, Indiana University  
edizque@iu.edu

## Abstract

What insights can statistical analysis of the time series recordings of neurons and brain regions during behavior give about the neural basis of behavior? With the increasing amount of whole-brain imaging data becoming available, the importance of addressing this unanswered theoretical challenge has become increasingly urgent. We propose a computational neuroethology approach to begin to address this challenge. We evolve dynamical recurrent neural networks to be capable of performing multiple tasks. We then analyze the neural activity using popular network neuroscience tools, specifically functional connectivity using Pearson's correlation, mutual information, and transfer entropy. We compare the results from these tools against a series of informational lesions, as a way to reveal their degree of approximation to the ground-truth. Our initial analysis reveals an overwhelming large gap between the insights gained from statistical inference of the functionality of the circuits based on neural activity and the actual functionality of the circuits as revealed by mechanistic interventions.

## Introduction

A central goal in neuroscience is to understand how the brain, body and environment come together to produce behavior. Specifically, we would like to understand in some detail the functional role of the nervous system in behavior. To this end, researchers are imaging with increasing time and spatial resolution the neural activity of living organisms at various scales, ranging from *C. elegans* to humans (Nguyen et al., 2016; Aimon et al., 2019; Randlett et al., 2015). Furthermore, technological advancements are starting to make recording of neural activity from freely moving animals possible (Lin et al., 2022). This increase in neural activity data has led to a similar increase in statistical measures and methods for inferring function from the time series of the neural activity (Paninski and Cunningham, 2018; Ramaswamy, 2019). Despite the incredible experimental progress and the overwhelming explosion in data availability, a fundamental theoretical challenge remains open (see Fig. 1): What can statistical measures of neural activity during behavior reveal about the function of the components of the nervous system?

Of the wide range of statistical methods that are available, the application of network theoretic tools to interpret

animal brain activity as it pertains to behavior and disease has seen an explosion of interest in the last decade (Sporns, 2010; Fornito et al., 2016). Specifically, there has been a myriad of methods for constructing functional connectivity networks from neural activity to understand the interaction between brain regions at various scales with the ultimate goal of understanding the underlying causal relationships (Van Den Heuvel and Pol, 2010; Smith et al., 2011; Yeo et al., 2011). Of these, the most popular methods include Pearson's correlation, mutual information and transfer entropy (Friston, 1994). While these statistical methods can provide very useful insights about the interactions between the different components of the neural system, they provide no guarantees as to their ability to converge to the ground-truth causal relationships.

Computational models of neural networks have proven to be an excellent test bed for generating and evaluating such statistical methods (Dayan et al., 2003). For instance, using a computational model of a fully-connected spiking neural network, Ito et al. (2011) showed that while transfer entropy can get close, it still cannot estimate the structural connectivity of a neural network from its activity alone. Similarly, using a recurrent dynamical neural network model optimized to perform a task, Candadai and Izquierdo (2020) showed that mutual information cannot disambiguate between predictive information from different sources. Maheswaranathan et al. (2019) showed that analysis of some features of the representation geometry led to conclusions that were not related to the function of the network, while others did. Similar approaches have been taken to show that polyadic interactions and the presence of underlying common inputs present challenges to these methods (Stevenson et al., 2008; James and Crutchfield, 2017). In such studies, two aspects that is often overlooked are: (a) animals don't function as "brains in a vat" but are embodied and embedded in the environments that they are continuously interacting with; and (b) natural systems are multifunctional, however most computational models that are studied are typically built to perform only a single behavior. This is where computational neuroethological approaches to understand-

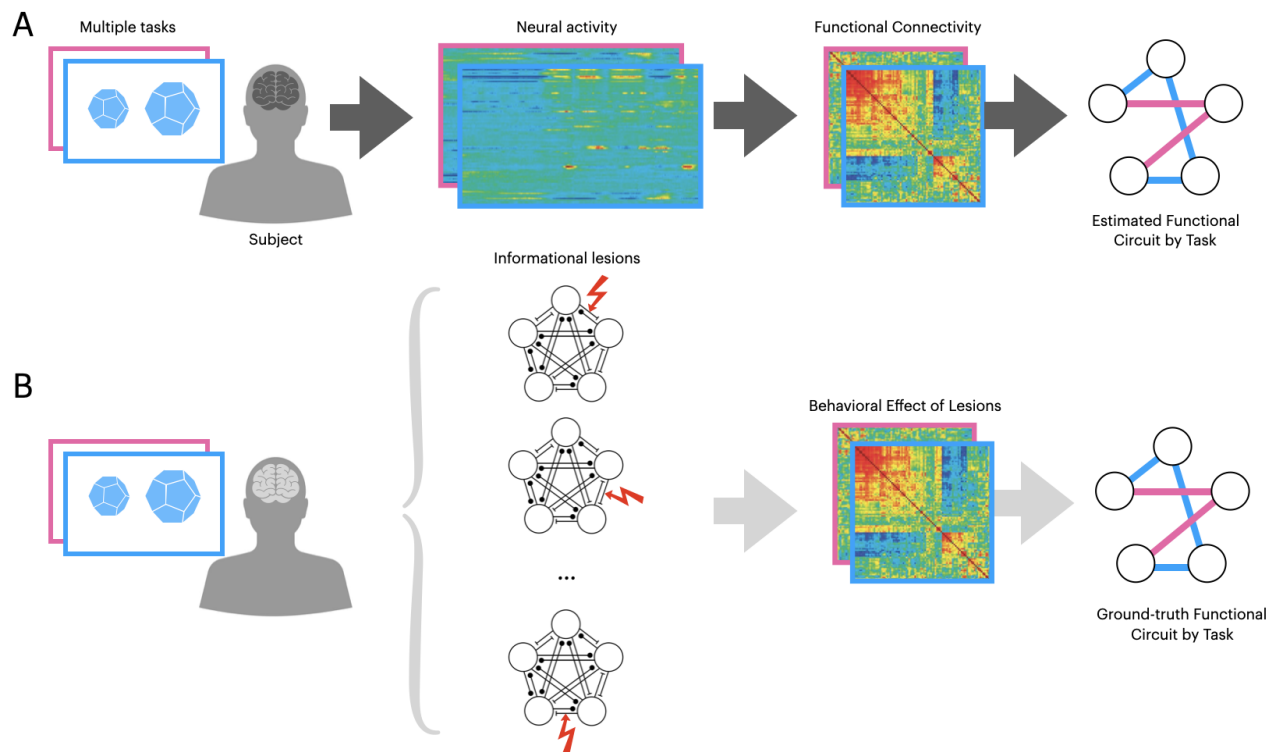


Figure 1: Computational neuroethology approach to uncover what a statistical measure of functional connectivity tell us about the actual functional connectivity of a nervous system. In order to address this theoretical challenge, we propose the following paradigm. In scenario (A), the subject is presented with two different tasks (blue and magenta). For each task, there are multiple trials (different sizes, different starting conditions). For each trial, neural activity is recorded for the subject. From the combined neural recordings of each task, a node functional connectivity (nFC) is created, using one of three techniques: Pearson’s correlation, mutual information, and transfer entropy. Finally, from the nFC the subcircuit for each task is estimated. In scenario (B), the same subject is now tested on the same two tasks, but now the drop in performance is recorded during information lesions to each pair of connections or each individual connection between the components of the subject’s brain. This effect of lesions per pairwise component is considered the actual functional connectivity (aFC). From it, the ground truth functional circuit is obtained for each task, which is used to assess the usefulness of the statistical nFC approach. We cannot do part B of this approach with humans, or with any other living organism, given current experimental limitations and ethical considerations. However, we can use artificial life techniques to: first generate agents capable of multiple tasks, and then analyze them in the way proposed above.

ing the neural basis of behaviors come in (Beer, 1996; Datta et al., 2019; Candadai, 2021): build computational models of brain-body-environment (BBE) systems, optimize them to perform multiple tasks, and use them as ground-truth for neural network operation with *behavior* defining what function means.

We would like to revisit the question that is at the heart of all the work that uses Network Neuroscience tools from a computational neuroethology perspective: What can statistical analysis of the time-series of the neural activity of brain-body-environment systems tell us about the behavioral functionality of the system? Our overarching goals are to: First, build intuition about what the statistical methods tell

us about the ground-truth. Second, use this to guide and inform predictions and generate hypotheses in experiments. Finally, improve the tools of analysis for complex BBE systems. In this paper, we would like to tackle the first of those goals in the simplest set of conditions possible: (1) Evolve a BBE model to perform a pair of visually-guided behaviors; (2) Infer the functional connectivity for each task from neural activity time series using Pearson’s correlation, mutual information, and transfer entropy; (3) Compare the insights gained against the ground-truth obtained from an informational lesion characterization of the circuit.

This paper is organized as follows: the next section describes the design of visually-guided behaviors, and the

agent; the following sections describe our results with regards to how the agent performed said behaviors and the comparison between statistically-inferred functional connectivity and the ground-truth; finally we discuss our results and present ideas for future work.

## Methods

We replicated the visually-guided agent described in (Beer, 1996), including two tasks previously used in (Beer, 1996; Slocum et al., 2000). The model agent is illustrated in Figure 2. The agent has a circular body with a diameter of 30 (in an environment of size  $400 \times 275$ ). The agent possesses an “eye” consisting of a foveated array of distance sensors. The eye consists of 15 proximity sensors of maximum length 220, uniformly distributed over a visual angle of  $\pi/4$ . An intersection between a ray and an object causes an input to be injected into the corresponding sensory neuron. The magnitude of the injected input is inversely proportional to the distance to the object, with values ranging from 0 (no intersection) to 10 (no separation). The agent has two “motors” that produce 1D movement of the entire body. The agent moves according to first-order dynamics, with motor neurons directly specifying the velocity of movement. The agent’s horizontal velocity is proportional to the sum of opposing forces produced by a bilateral pair of effectors (with a constant of proportionality of 8).

The agent’s behavior is controlled by a continuous-time recurrent neural network (CTRNN) with the following state equation:

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^N w_{ji} \sigma(y_j + \theta_j) + I_i \quad (1)$$

where  $y$  is the state of each neuron,  $\tau$  is its time constant,  $w_{ji}$  is the strength of the connection from the  $j^{th}$  to the  $i^{th}$  neuron,  $\theta$  is a bias term,  $\sigma(x) = 1/(1 + e^{-x})$  is the standard logistic activation function, and  $I$  represents an external input (e.g., from a sensor). States were initialized to 0 and circuits were integrated using the forward Euler method with an integration step size of 0.1.

A real-valued genetic algorithm was used to evolve CTRNN parameters. A population of individuals was maintained, with each individual encoded as a length  $M$  vector of real numbers. Initially, a random population of vectors was generated by initializing each component of every individual to random values uniformly distributed over the range  $\pm 1$  (they could not move outside this range during evolution). Individuals were selected for reproduction using a linear rank-based method. Children were generated by either mutation or crossover with an adjustable crossover probability. A selected parent was mutated by adding to it a random displacement vector whose direction was uniformly distributed on the  $M$ -dimensional hypersphere and whose magnitude was a Gaussian random variable with 0

mean and variance  $\sigma^2$  (for details, see Slocum et al. (2000)). Search parameters in the range  $\pm 1$  were mapped linearly into CTRNN parameters with the following ranges: connection weights  $\in [-5, 5]$ , biases  $\in [-5, 5]$ , and time-constants  $\in [1, 2]$ . The parameters of the neural circuit were bilaterally symmetric.

For the first task, the embodied agent must be capable of visually discriminating between objects of different sizes, catching smaller circular objects while avoiding the larger circular objects. Objects fell straight down with an initial horizontal offset in the range  $\pm 25$  and a vertical velocity of 3. The circular objects had a diameter in the range  $[20, 40]$ . Accordingly, the performance measure to be maximized was:

$$f_A = \frac{1}{T} \sum_{i=1}^T p_i \quad (2)$$

where  $p_i = 1 - d_i$  for smaller circular objects and  $p_i = d_i$  for larger circular objects,  $d_i$  is the horizontal distance between the centers of the object and the agent when their vertical separation goes to zero on the  $i$ th trial (clipped to MaxDistance and normalized to run between 0 and 1),  $T$  is the total number of trials, and  $D$  is the maximum distance. The reason that  $d_i$  was clipped to  $D$  was to prevent the avoidance of, for example, larger circles by large distances from dominating the fitness at the expense of accuracy in catching smaller circles. A total of 24 evaluation trials were used during evolution, uniformly distributed over the range of horizontal offsets.

For the second task, the embodied agent must become sensitive to the relationship of its own body to its surroundings and it must be able to perceive the actions that this environment affords. We evolved agents that could accurately distinguish between passageways and obstacles in a falling wall, passing through openings wide enough to accommodate their bodies while avoiding openings that were too narrow. Walls consisting of two squares of width 20 separated by an aperture whose width was in the range  $[20, 40]$  dropped from above with a vertical velocity of 3 and a horizontal offset of  $\pm 25$  relative to the agent. Accordingly, the performance measure to be maximized was:

$$f_B = \frac{1}{T} \sum_{i=1}^T p_i \quad (3)$$

where  $p_i = 1 - d_i$  for an aperture wide enough for the agent to pass through and  $p_i = d_i$  for an opening too narrow for the agent to pass through,  $d_i$  is the horizontal distance between the centers of the object and the agent when their vertical separation goes to zero on the  $i$ th trial (again clipped to MaxDistance and normalized to run between 0 and 1),  $T$  is the total number of trials, and  $D$  is the maximum distance. A total of 24 evaluation trials were used during evolution, uniformly distributed over the range of horizontal offsets.




Figure 2: Agent and tasks set up. **(A)** Agent. Sensory neurons (blue) are interconnected to interneurons (black), then which are in turn connected to the two motor neurons (red). The interneurons are recurrently interconnected. The nervous system is bilaterally symmetric. **(B)** Object-size discrimination task. The agent moves horizontally while a circle of different size falls from above. The rays of the agent's proximity sensors are shown in dashed blue. **(C)** Perceiving affordances task. The agent moves horizontally while a wall with an adjustable aperture falls from above.

Statistical analyses of the neural activity, namely mutual information and transfer entropy, were estimated using the *infotheory* Candadai and Izquierdo (2019) package. Both metrics were estimated using an average shifted histogram based binning of the normalized activity with 100 bins along each dimension.

## Part I: Generating a Multi-Functional Agent

In order to study the relationship between the statistical functional connectivity inferred from neural activity and the actual functional connectivity from lesion analysis, we first need an agent that is capable of performing multiple tasks. Thus, our first step was to generate an ensemble of successful multi-behavioral embodied dynamic recurrent neural systems. The agents were tasked with solving two minimally-cognitive tasks Beer (1996); Slocum et al. (2000): an object-size discrimination task and a perceiving affordances task. We performed 100 evolutionary runs with different random seeds (Fig. 3A). Agents are evolved to solve both tasks. During each fitness evaluation, an agent is tested first on the 24 evaluation trials of the object-size discrimination task and then on 24 evaluation trials of the perceiving affordances task. The performance is calculated for each task according to the defined fitness function (see Methods) and multiplied together to produce a score. After 1000 generations, many of the runs found successful configurations of the neural circuit that could solve the multiple tasks. A histogram of the final performance of each of the best agents in each run is shown in Fig. 3B. At least three of the runs achieved near-perfect performance on the fitness evaluation. As far as we are aware, this is the first report of agents successfully evolved for multiple minimally-cognitive tasks.

Before setting out to analyze one of these agents in some detail, it is important that we examine the generality and

robustness of these solutions across a wider range of behavioral conditions. Our goal is to use this further examination to select which circuit to analyze in detail. We based our selection on which of the evolved agents solved both tasks equally well and also generalized well across a wider range of conditions for both tasks. There were three key changes with respect to the original fitness evaluation: (a) The step size of integration is made smaller (from 0.1 to 0.01). (b) The range of object sizes and aperture sizes was drawn from [20,40] in steps of 0.05 instead of 1. (c) The starting position was drawn from [1,5] in steps of 0.01. Altogether, this corresponds to 200000 trials for the generalization performance analysis (up from 100 trials per evaluation for the fitness function). In Figure 3C, we show the performance of each of the final circuits from each of the 100 evolutionary runs in each of the two tasks. As expected, the same three circuits that performed best in the fitness evaluation also generalized best across the wider range of conditions and are thus most appropriate for further examination. The best performing agent obtained a near-perfect performance 0.977 on the object-size discrimination task and 0.976 on the perceiving affordances task. We focus on this agent for the remainder of this paper. What is the behavior and neural activity of this agent? In Figure 4, we show the behavior and neural activity for this best agent across the two different tasks. Traces are colored according to whether the agent has to catch or avoid the object or pass through the aperture or avoid it. We use these neural traces for the functional connectivity analysis ahead.

## Part II: Comparing Functional Connectivities

Our second step is to analyze how well the insights gained from the statistically-inferred functional connectivity help us understand the actual functional connectivity of the circuit, as determined through informational lesion analysis.

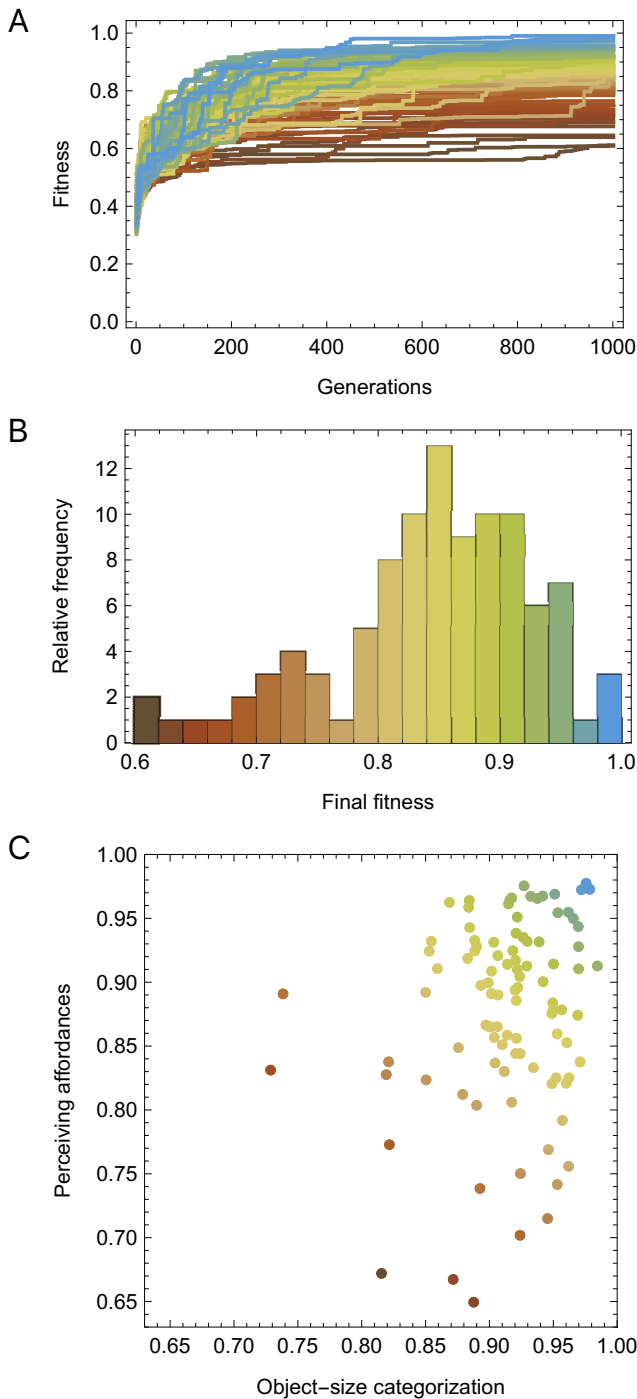


Figure 3: Generating a multifunctional agent using an evolutionary algorithm. (A) Best fitness over generations for hundred evolutionary runs, color coded by final fitness. (B) Relative frequency of the final fitness for each of the evolutionary runs. (C) Generalization performance for each of the final solutions on the two tasks: object-size discrimination and perceiving affordances.

In order to do this, we first characterize the node functional connectivity using traditional methods from network neuroscience (first column, Fig. 5). Specifically, we focus here on estimating the functional connectivity between pairs of neurons using Pearson's correlation, mutual information, and transfer entropy (Candadai and Izquierdo, 2019). We perform the analysis using the neural traces generated from each of the two tasks (Fig. 4B). From the magnitude of each of these measures, we establish the degree of involvement of each pair of neurons to the task. That is, pairs of neurons with little or no correlation in their activity during a task are deemed unlikely to be involved in that task; whereas pairs of neurons with strong correlation (either positive or negative) are deemed likely to be involved in the task.

Second, we characterize the two-way and one-way causal pairwise function of the edges using informational lesions (second column, Fig. 5). For the two-way informational lesions, we clamp the interchange of activity between two neurons in both directions, have the agent perform the task at hand, and measure the deficit in performance. We do this for a large range of potential basal outputs for each neuron in the pair (from an output of 0 to 1 in steps of 0.01 for each neuron), and we select the smallest deficit generated. The two-way lesions is used as the ground-truth for the nFC generated using Pearson's correlation and mutual information, because both provide only symmetric information between two neurons. For the nFC calculated using transfer entropy, which is directional, we perform a one-way information lesion analysis. The method here is the same as the previous one but it is performed for each individual connection in the circuit.

Finally, we compare the difference between the estimated involvement of each of the pairs calculated using nFCs to the ground-truth functional involvement characterized during the lesion studies for the two tasks (third column, Fig. 5). In these plots, each point represents a pair of neurons in the circuit and the different colors represent the different tasks, with the gray line connecting the same pair of neurons across different tasks. Points in the upper left corner represent unimportant connections for the task that are well captured by the statistical measurements. That is, connections between neurons that do not have much correlation and whose physical disruption does not cause a noticeable effect on task performance. Points in the bottom right corner represent important connections for the task that are also well captured by the statistical measurements. That is, connections between neurons that have a strong correlation and whose physical disruption causes a noticeable effect on task performance. Points in the upper right triangle represent points that are not causally important to the task, but that come up as important in the statistical measurements. Points in the bottom left triangle represent points that are causally important, but that do not show up as relevant in the statistical measures. The distance between the yellow and blue

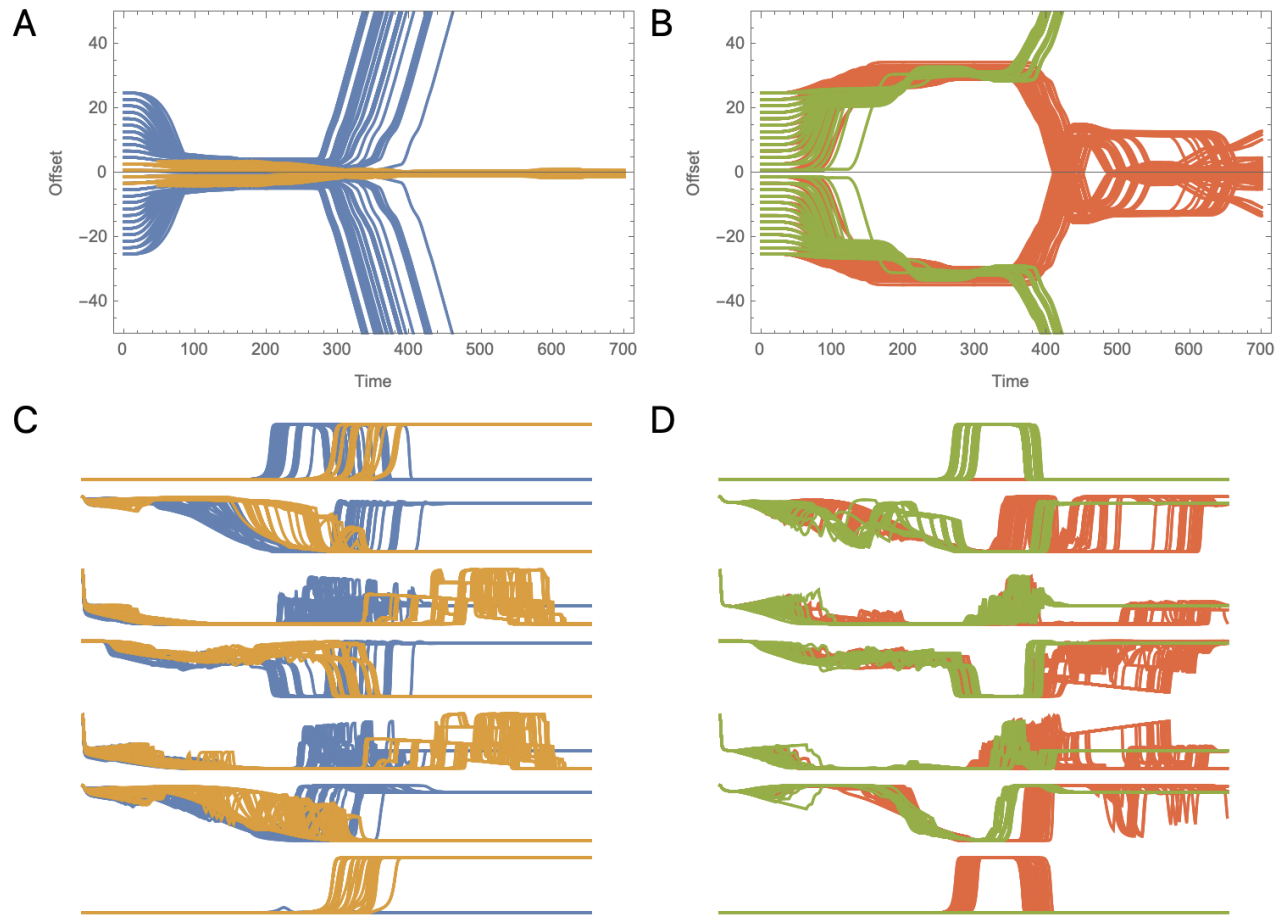


Figure 4: Behavior and neural activity across both tasks. (A and B) Behavioral traces during object-size discrimination and perceiving affordances tasks, respectively. Relative position of the agent in relation to the center of the falling objects over time. Traces are color coded depending on whether they need to be caught or avoided. In the object-size discrimination task, trials with the smaller circles are shown in yellow and traces with the larger circles are shown in blue. In the perceiving affordances task, trials with the smaller apertures are shown in green and those with the larger apertures are shown in red. (C and D) Neural activity from the seven recurrently interconnected interneurons driving behavior for each of the two tasks. Color coding follows the same pattern used for panels A and B.

points shown by the gray line represent the functional variation across tasks.

We highlight here four key qualitative insights gained from this analysis (Fig. 5). First, across all levels of analysis, the functional connectivities are different depending on the task being performed. This is true across all levels of analysis, including the informational lesion studies. This shows how the same nervous system, even without neuromodulation and synaptic plasticity, can have functionally different configurations, based on task engagement alone. This highlights the importance of studying nervous systems in the context of behavior. Second, the one-way lesion analysis reveals major differences in the directionality of interactions between components in the circuit. This, of course, high-

lights the inherent limitations of the two more established methods for estimating functional connectivity due to their symmetrical treatment of the relationships, Person's correlation and mutual information. Third, as can be appreciated by even a mere cursory look at how much darker the range of shades of the maps in the nFC column are in relation to those in the aFC column, all of the statistical measures overinfer the importance of the relationship between the variables relative to the actual role that those relationships play, as determined by the information lesions. As can be seen more easily in the third column, across all three measurements, lesions to the majority of connections have little or no effect on the behavior, despite the statistical measures inferring high levels of correlation, mutual information, and



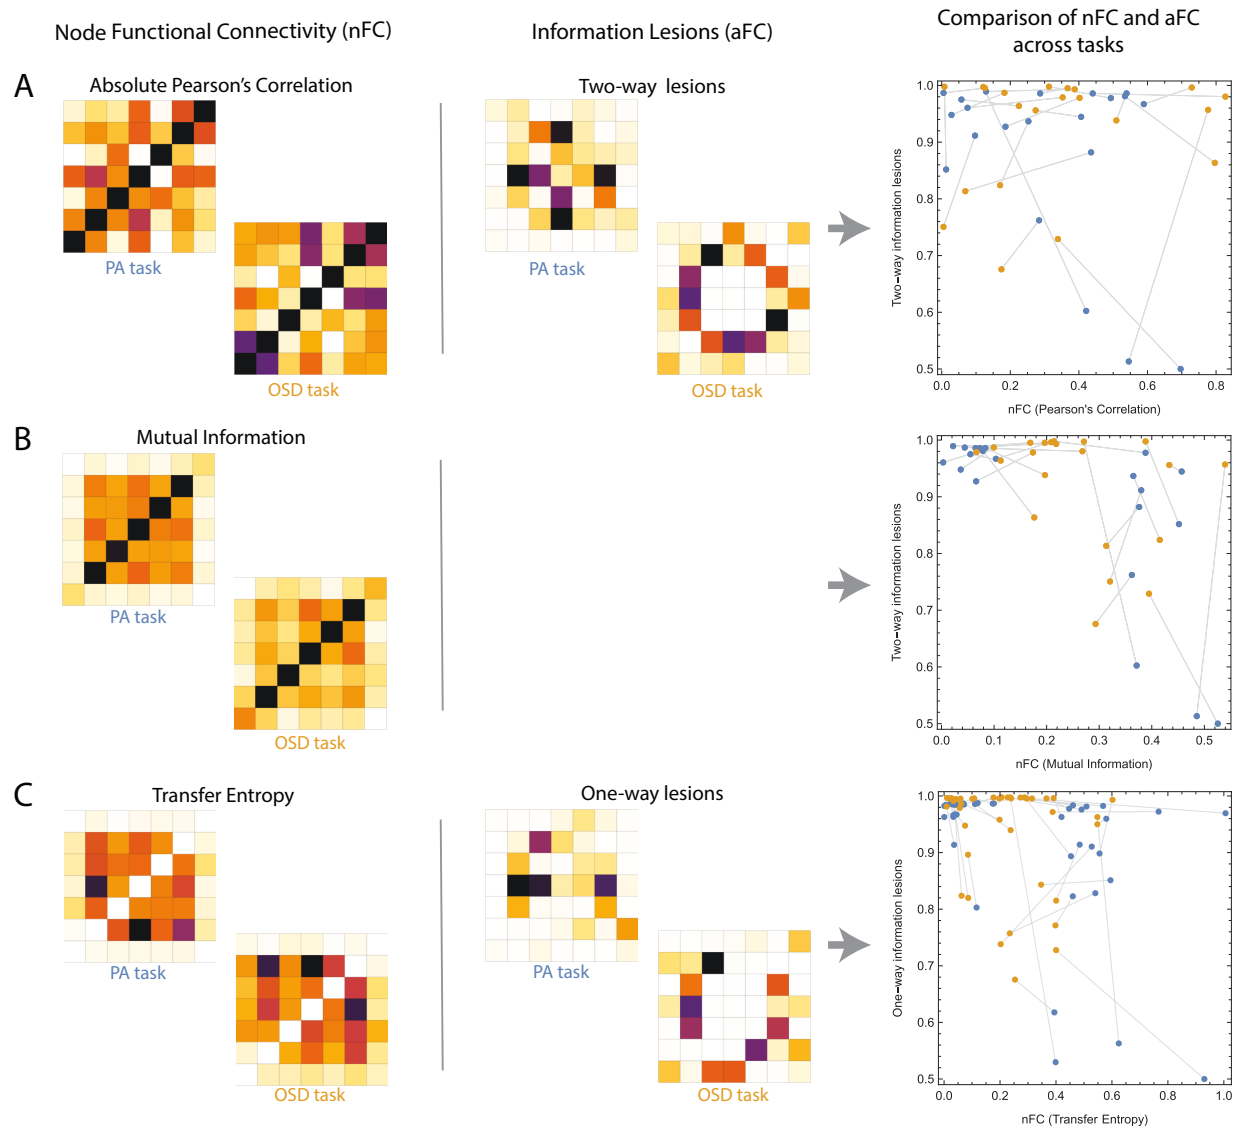


Figure 5: Comparing insights gained from the statistical measures of node functional connectivity based on neural activity to the ground-truth characterization from informational lesions analysis across tasks. (A) A comparison of the nFC using Pearson's correlation to the aFC using two-way lesions. (B) A comparison of the nFC using mutual information to the aFC using two-way lesions again. (C) A comparison of the nFC using transfer entropy to the aFC using one-way lesions. Across the first two columns, the matrix represents the pairwise interactions between the seven interneurons in the circuit. The shade of each cell represents the level of involvement of that pair in each task. The lighter the color represents little or no involvement, where white corresponds to no correlation, no mutual information, no transfer entropy, and no deficit in performance after a two-way lesion, or a one-way lesion, respectively. The darker shades represent increasing involvement, where black corresponds to the maximum level of correlation, mutual information, and transfer entropy for the statistical measurements, and a complete disruption in the performance of the task in the case of the two-way and one-way lesions. The last column depicts the comparison between statistical inference (nFC) and ground-truth (aFC). Each point represents a pair of neurons in the circuit (blue for the perceiving affordances task and yellow for the object-size discrimination task). The gray line connecting two dots represents the same pair of neurons across the different tasks. On the y-axis, the two-way and one-way lesions are defined such that the value represents the deficit in performance. So a performance of 0.95 after a lesion to the connection between neurons X and Y means that this link was not very important to the functioning of the circuit; on the other hand a resulting performance of 0.5 after a lesion would indicate that the connection between those pair of neurons is extremely important to the functioning of the circuit. See main text for interpretation of the results.

transfer entropy between them. Although this is expected to some degree, the full effect can be quantified here concretely. Fourth, despite the tendency for statistical measures to over-infer causality that we just discussed, our analysis of this agent allows us to see some clear examples of connections that are causally important that do not show up in the statistical measurements. These are, of course, much less common. They are particularly salient in the Pearson's correlational analysis and in the transfer entropy analysis.

## Discussion

In this paper, we set out to answer the question: What does functional connectivity tell us about the behaviorally functional connectivity of a multifunctional neural circuit? We used a computational neuroethology approach to begin to address this theoretical challenge. We evolved a dynamical recurrent neural network to be capable of performing multiple tasks, and then we analyzed its neural activity using traditional network neuroscience tools. While our analyses were performed on a neuron-to-neuron basis and functional connectivity is typically performed across brain regions, CTRNNs are universal function approximators and can model neural activity in brain regions thus enabling our analyses to scale. We then compared the results against a series of informational lesions as a way to reveal their degree of approximation to the ground-truth. Overall, our analysis reveals a large gap between the insights gained from statistical inference of the functionality of the circuits based on neural activity and the actual functionality of the circuits as revealed by mechanistic interventions.

It is important to note that the measures of functional connectivity being investigated in this paper are measures of a statistical relationship. They are neither measures of causal effect nor of how such a statistical relationship might relate to interventional impact on a task. However, these statistical methods are often used as tools to make claims about the relationship between neural activity and behavior. The goal of this paper is to examine the degree to which those measures of a statistical relationship estimate causal effect on behavior, as determined in this case through the interventional impact on task function. While there has been some work that had attempted to answer this question in the past (Ay and Polani, 2008; Lizier and Prokopenko, 2010; Chicharro and Ledberg, 2012), that work considered only neural circuits in a vacuum; here we extend this work to consider functional and complete brain-body-environment systems.

It is relatively straightforward to understand why a connection between two neurons may have a high correlation, a high mutual information, or a high transfer entropy, and yet not have a high interventional causality: The two neurons can be correlated for reasons other than their connection to each other. The opposite is also straightforward to understand (i.e., why a connection between two neurons may have a low correlation, low mutual information, or low

transfer entropy and nevertheless have a high interventional causality): The transformation of information between the two neurons may be such that the two neurons do not have similar informational profiles and yet they are still causally linked. Finally, it is only to the degree that the two neurons are causally linked, that lesioning the connection will have some effect on functional performance. In other words, a causal link is necessary for function, but not sufficient. There may be some connections that are causally linked, but that do not contribute to the circuit's function. Finally, it is important to note that all analyses in this work was done on dyads and the statistical and interventional methods alike would benefit from polyadic analysis.

## Future Work

We have three main directions of future work. (1) In this paper, we deliberately focus our analysis on a single circuit as a way to begin to gather intuition. One direction for future work is to perform a similar analysis across an ensemble of successful but different solutions, as a way to uncover the more general principles. This will involve examining methods for quantifying how close the different statistical methods approximate the causal relationships. (2) The neural network under consideration in our current analysis was fully recurrent. We would like to study the degree to which the structural connectivity of the neural circuit affects the usefulness of the functional connectivity. One direction for future work will be to systematically study the effect that the structure of the connectivity of the circuit has on the performance of the statistical measures of functional inference. As part of this analysis, we also plan to study structures that are grounded in available connectomes. (3) Finally, we deliberately focused our analysis on the most popular measures of node functional connectivity. One additional direction for future work is to study the wider range of measurements used in the network neuroscience literature. Crucially, we would like to use further refined versions of the computational neuroethology approach proposed here as an ideal testground for generating novel variations of statistical measurements to gain insights on the functional connectivity of these complex neural circuits.

## Acknowledgements

We are thankful for the reviewers' comments. This material is based upon work supported by the National Science Foundation under Grant No. 1845322.

## References

- Aimon, S., Katsuki, T., Jia, T., Grosenick, L., Broxton, M., Deisseroth, K., Sejnowski, T. J., and Greenspan, R. J. (2019). Fast near-whole-brain imaging in adult drosophila during responses to stimuli and behavior. *PLoS biology*, 17(2):e2006732.
- Ay, N. and Polani, D. (2008). Information flows in causal networks. *Advances in complex systems*, 11(01):17–41.



- Beer, R. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 421–429.
- Candadai, M. (2021). Information theoretic analysis of computational models as a tool to understand the neural basis of behaviors. *arXiv preprint arXiv:2106.05186*.
- Candadai, M. and Izquierdo, E. J. (2019). infotheory: A c++/python package for multivariate information theoretic analysis. *arXiv preprint arXiv:1907.02339*.
- Candadai, M. and Izquierdo, E. J. (2020). Sources of predictive information in dynamical neural networks. *Scientific reports*, 10(1):1–12.
- Chicharro, D. and Ledberg, A. (2012). When two become one: the limits of causality analysis of brain dynamics. *PLoS one*, 7(3):e32466.
- Datta, S. R., Anderson, D. J., Branson, K., Perona, P., and Leifer, A. (2019). Computational neuroethology: a call to action. *Neuron*, 104(1):11–24.
- Dayan, P., Abbott, L. F., et al. (2003). Theoretical neuroscience: computational and mathematical modeling of neural systems. *Journal of Cognitive Neuroscience*, 15(1):154–155.
- Fornito, A., Zalesky, A., and Bullmore, E. (2016). *Fundamentals of brain network analysis*. Academic Press.
- Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: a synthesis. *Human brain mapping*, 2(1-2):56–78.
- Ito, S., Hansen, M. E., Heiland, R., Lumsdaine, A., Litke, A. M., and Beggs, J. M. (2011). Extending transfer entropy improves identification of effective connectivity in a spiking cortical network model. *PLoS one*, 6(11):e27431.
- James, R. G. and Crutchfield, J. P. (2017). Multivariate dependence beyond shannon information. *Entropy*, 19(10):531.
- Lin, A., Witvliet, D., Hernandez-Nunez, L., Linderman, S. W., Samuel, A. D., and Venkatachalam, V. (2022). Imaging whole-brain activity to understand behaviour. *Nature Reviews Physics*, pages 1–14.
- Lizier, J. T. and Prokopenko, M. (2010). Differentiating information transfer and causal effect. *The European Physical Journal B*, 73(4):605–615.
- Maheswaranathan, N., Williams, A., Golub, M., Ganguli, S., and Sussillo, D. (2019). Universality and individuality in neural dynamics across large populations of recurrent networks. *Advances in neural information processing systems*, 32.
- Nguyen, J. P., Shipley, F. B., Linder, A. N., Plummer, G. S., Liu, M., Setru, S. U., Shaevitz, J. W., and Leifer, A. M. (2016). Whole-brain calcium imaging with cellular resolution in freely behaving *Caenorhabditis elegans*. *Proceedings of the National Academy of Sciences*, 113(8):E1074–E1081.
- Paninski, L. and Cunningham, J. P. (2018). Neural data science: accelerating the experiment-analysis-theory cycle in large-scale neuroscience. *Current opinion in neurobiology*, 50:232–241.
- Ramaswamy, V. (2019). An algorithmic barrier to neural circuit understanding. *BioRxiv*, page 639724.
- Randlett, O., Wee, C. L., Naumann, E. A., Nnaemeka, O., Schoppik, D., Fitzgerald, J. E., Portugues, R., Lacoste, A. M., Riegler, C., Engert, F., et al. (2015). Whole-brain activity mapping onto a zebrafish brain atlas. *Nature methods*, 12(11):1039–1046.
- Slocum, A., Downey, D., and Beer, R. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*, pages 430–439.
- Smith, S. M., Miller, K. L., Salimi-Khorshidi, G., Webster, M., Beckmann, C. F., Nichols, T. E., Ramsey, J. D., and Woolrich, M. W. (2011). Network modelling methods for fmri. *Neuroimage*, 54(2):875–891.
- Sporns, O. (2010). *Networks of the Brain*. MIT press.
- Stevenson, I. H., Rebesco, J. M., Miller, L. E., and Kording, K. P. (2008). Inferring functional connections between neurons. *Current opinion in neurobiology*, 18(6):582–588.
- Van Den Heuvel, M. P. and Pol, H. E. H. (2010). Exploring the brain network: a review on resting-state fmri functional connectivity. *European neuropsychopharmacology*, 20(8):519–534.
- Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., et al. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*.