Dynamical networks: finding, measuring, and tracking neural population activity using network theory

Mark D. Humphries

Faculty of Biology, Medicine, and Health, University of Manchester, Manchester, UK.

Contact: mark.humphries@manchester.ac.uk

Abstract

Systems neuroscience is in a head-long rush to record from as many neurons at the same time as possible. As the brain computes and codes using neuron populations, it is hoped these data will uncover the fundamentals of neural computation. But with hundreds, thousands, or more simultaneously recorded neurons comes the inescapable problems of visualising, describing, and quantifying their interactions. Here I argue that network theory provides a set of scalable, analytical tools that already solve these problems. By treating neurons as nodes and their interactions as links, a single network can visualise and describe an arbitrarily large recording. I show that with this description we can quantify the effects of manipulating a neural circuit, track changes in population dynamics over time, and quantitatively define theoretical concepts of neural populations such as cell assemblies. Using network theory as a core part of analysing population recordings will thus provide both qualitative and quantitative advances to our understanding of neural computation.

Neurons use spikes to communicate (Rieke et al., 1999). From this communication arises coding and computation within the brain; and so arises all thought, perception, and deed. Understanding neural circuits thus hinges critically on understanding spikes across populations of neurons (Pouget et al., 2013; Wohrer et al., 2013; Yuste, 2015).

This idea has driven a technological arms race in systems neuroscience to record from as many individual neurons as the same time as physically possible (Stevenson and Kording, 2011). Current technology, ranging from imaging of fluorescent calcium-binding proteins (Chen et al., 2013; Peron et al., 2015a,b) to large scale multi-electrode arrays and silicon probes (Buzsáki, 2004; Jun et al., 2017), now allows us to capture the activity of hundreds of neurons in a range of brain systems. These include such diverse systems as invertebrate locomotion, through zebrafish oculomotor control, to executive functions in primate prefrontal cortex. With the data captured, the key question for any system becomes: how do we describe these spike data? Visualise them? And how do we discover the coding and computations therein?

Here I argue that network theory provides a set of tools ideally suited to both describe the data and discover new ideas within it. Networks are simply a collection of nodes and links: nodes representing objects, and links representing the interactions between
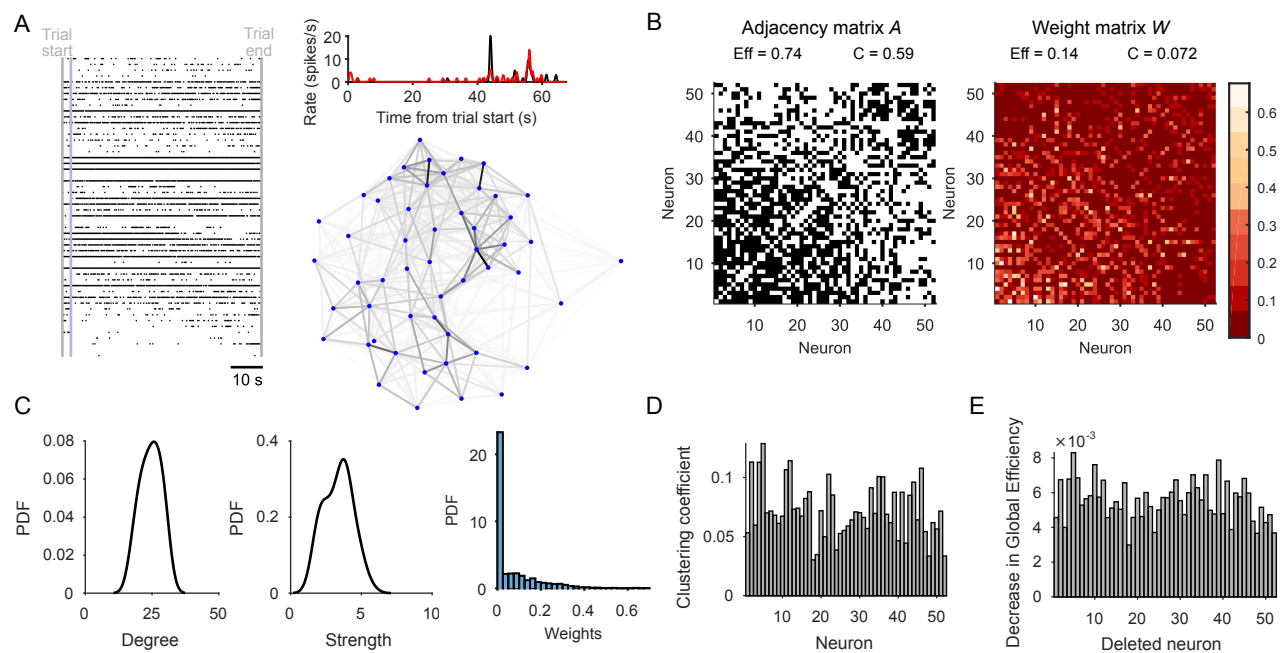
Figure 1: Quantifying neural population dynamics using network theory.
A) Schematic of turning neural activity time-series into a network. Left: a raster plot of 52 simultaneously recorded neurons in rat medial prefrontal cortex, during a single trial of a Y-maze navigation task. Blue line: reached end of goal arm. Grey lines: at, and returned to, the starting position. Right: the corresponding network representation: nodes are neurons, links indicate pairwise interactions, and their grey-scale indicates the strength of interaction. Top: Interactions here are rectified Pearson's R between pairs of spike-trains convolved with a Gaussian ($\sigma = 250$ ms); two example convolved trains are plotted here. B) Representations of the network in panel A: the adjacency matrix describes the presence (black) or absence (white) of links; the weight matrix describes the strengths of those links. Neurons are ranked by total link strength in descending order. Above each we give the global efficiency (Eff) and average clustering coefficient (C), respectively measuring the ease of getting from one node to another, and the density of links in the neighbourhood of one node. C) Distributions of node degree (total number of links per node), node strength (total weight of links per node), and link strength for the network in panel A. D) Network clustering fingerprint. A histogram of the weighted clustering coefficient for each neuron, measuring the ratio of weighted triangles to weighted triples in which that neuron participates: the higher the ratio, the more strongly connected is the neighbourhood of that neuron. Some neurons (e.g. 2, 5) have strongly connected neighbourhoods, implying a local group of correlated neurons. E) Network efficiency fingerprint, given by the decrease in the network's global efficiency after deleting each neuron in turn. Neurons that strongly decrease the efficiency (e.g. 3) are potential network hubs, mediating interactions between many neurons.

those objects. This representation can encapsulate a wide array of systems, from email traffic within a company, through the social groups of dolphins, to word co-occurrence frequencies in a novel (Newman, 2003). By abstracting these complex systems to a network description, we can describe their topology, compare them, and deconstruct them into their component parts.

2

Describing multi-neuron data as a network

A network description of multi-neuron recording data rests on two ideas: the nodes are the neurons, and the links are the interactions between the neurons (Figure 1A). (Strictly, the nodes are the isolated time-series of neural activity, with the usual caveats applied to the accuracy of spike-sorting for electrodes or image segmentation and stability for imaging; Harris et al., 2016). An immediate advantage of a network formalism is that it separates the details of choosing the interaction from the network topology itself.

We are free to choose any measure of pairwise interaction we like; and indeed that choice depends on what questions we want to ask of the data. Typical choices include cosine similarity or a rectified correlation coefficient, as these linear measures are familiar, easy to interpret, and not data-intensive. But with sufficient data we could also use non-linear measurements of interaction including forms of mutual information (Bettencourt et al., 2007; Singh and Lesica, 2010) and transfer entropy (Schreiber, 2000; Thivierge, 2014). We could fit an Ising model, so estimating "direct" interactions while factoring out other inputs (Yu et al., 2008). We could even fit a model to each neuron for the generation of its spike train, such as a generalised linear model (Truccolo et al., 2005; Pillow et al., 2008), and use the fitted weights of the inputs from all other neurons as the interaction values in a network (Gerhard et al., 2011). Whatever measure of interaction we use, the important distinction is between whether the interaction measurement is undirected (e.g. correlation coefficient) or directed (e.g. transfer entropy), and so whether we end up with an undirected or directed network as a result.

And we end up with a weighted network (Newman, 2004). While much of network theory, and its application in neuroscience, is focussed on binary networks in which a link either exists or it doesn't, any measurement of interaction gives us a weight for each link (Figure 1B). Thresholding the weights to construct a binary network inevitably loses information (Humphries, 2011; Zanin et al., 2012). Consequently, our multi-neuron recording data are best captured in a weighted network.

One immediate advantage of using a network formalism is that we then have access to a range of null models for testing the existence of meaningful interactions between neurons and changes to those interactions. These models define the space of possible networks obtained by some stochastic process. One common example is the configuration model (Chung and Lu, 2002; Fosdick et al., 2016), in which we assume connections between nodes are made proportional to the number of links they already have. This model, applied to neural time-series, is a null model for testing whether the existence of interactions between a pair of neurons is simply a result of those neurons having many interactions. Other null model networks include the exponential random graph model (Robins et al., 2007), or the stochastic block model and its variants (Newman and Martin, 2014). More work though is needed for appropriate weighted network null models (Rubinov and Sporns, 2011; Palowitch et al., 2016). Whatever null model network we chose, by using networks to describe neural recordings we open up a range of rigorous models for hypothesis testing.

What can we do with such "dynamical" networks of neurons? In the following I show how with them we can quantify circuit-wide changes following perturbations and

manipulations; we can track changes in dynamics over time; and we can quantitatively define qualitative theories of computational concepts.

Capturing circuit-wide dynamics and their changes by manipulations

Applying network theory to large-scale recordings of neural systems allows us to capture their complex dynamics in a compact form. The existing toolbox of network theory gives us a plethora of options for quantifying the structure of a dynamical network. We may simply quantify its degree and strength distributions, revealing dominant neurons (Figure 1C). We can assess the local clustering of the dynamical network, the proportion of a neuron's linked neighbours that are also strongly linked to each other (Watts and Strogatz, 1998; Figure 1D), revealing the locking of dynamics among neurons (Bettencourt et al., 2007; Sadovsky and MacLean, 2013). We can compute the efficiency of a network (Latora and Marchiori, 2001), a measure of how easily a network can be traversed (Figure 1E), revealing how cohesive the dynamics of population are - the higher the efficiency, the more structured the interactions amongst the entire population (Thivierge, 2014). We may define structural measures relative to a null model, such as quantifying how much of a small-world the dynamical network is (Yu et al., 2008; Gerhard et al., 2011). Our choice of quantifying measures depends on the aspects of dynamics we are most interested in capturing.

Having compactly described the dynamics, we are well-placed to then characterise the effects of manipulating that system. Manipulations of a neural system will likely cause system-wide changes in its dynamics. Such changes may be the fast, acute effect of optogenetic stimulation (Miesenböck, 2009; Boyden, 2015; Deisseroth, 2015); the sluggish but acute effects of drugs (Vincent et al., 2013); or the chronic effects of neurological damage (Otchy et al., 2015). All these manipulations potentially change the interactions between neurons, disrupting normal computation. By comparing the dynamical networks before and after the manipulation, one could easily capture the changes in the relationships between neurons.

There have been few studies examining this idea. Srinivas et al. (2007) used dynamical networks to quantify the changes to network-wide activity in hippocampus caused by the glutamate-injury model of epilepsy, suggesting a dramatic drop in network clustering in the epilepsy model. Vincent et al. (2013) used dynamical networks to quantify the potential neuroprotective effects of drug pre-conditioning in rat cortex in vitro, finding increased clustering and increased efficiency in the network, implying the drugs enriched the synaptic connections between groups of neurons. Quantifying manipulations using network theory is an under-explored application, rich in potential.

Efficiency: Reciprocal of the mean shortest path length between all pairs of nodes; path lengths are weighted. The higher the efficiency, the shorter the average path between a pair of nodes.

Small-world network: A network with both high clustering of nodes and high efficiency.
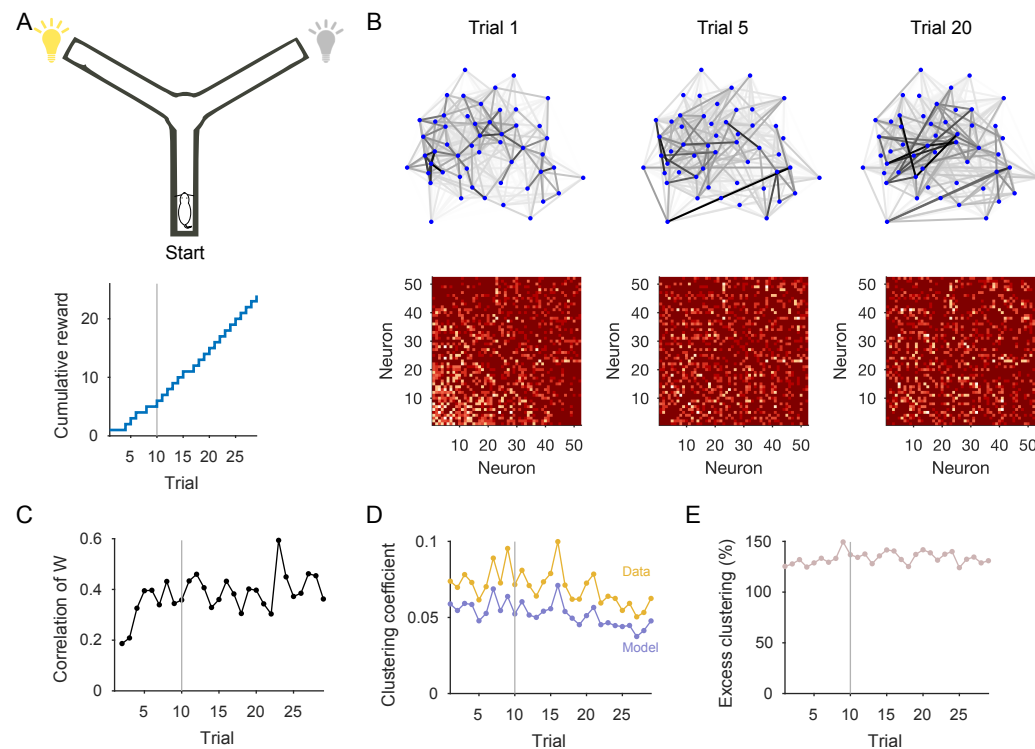
4

Figure 2: Tracking changes in neural population dynamics using network theory.
A) Recordings examined here are from one behavioural session of a Y-maze learning task. For this session, the rat had to reach the end of the randomly-cued arm to receive reward (schematic, top). This session showed evidence of behavioural learning (bottom), with a sustained increase in reward accumulation after trial 10 (grey line). A trial lasted typically 70 s, running from the rat leaving the start position through reaching the arm end and returning to the start position to initiate the next trial. B) Dynamical networks from trials 1, 5 and 20 of that session. The top row plots the networks, with nodes as neurons and greyscale links indicating the strength of pairwise interaction. The bottom row plots the corresponding weight matrix (ordered by total node strength in trial 1 throughout). The networks show a clear re-organisation of interactions between neurons during learning. C) Tracking network stability. The correlation between the weight matrix W at trial t and at trial $t-1$. The dynamical network rapidly increased in similarity over the first few trials. Grey line: behavioural learning trial. D) Clustering coefficient of the weighted network ('Data') on each trial; compared to the mean clustering coefficient over 20 null model weighted networks per trial ('Model'). E) Excess clustering in the data compared to the null model on each trial (data in panel D expressed as a ratio: $100 \times C_{\text{data}}/C_{\text{model}}$). The variation across trials in the data is well-accounted for by the null model, suggesting the average local clustering did not change over learning.

Tracking the evolution of dynamics

Neural activity is inherently non-stationary, with population activity moving between different states on a range of time-scales, from shifting global dynamics on time-scales of seconds (Zagha and McCormick, 2014), to changes wrought by learning on time-scales

5

of minutes and hours (Benchenane et al., 2010; Huber et al., 2012). For a tractable understanding of these complex changes, ideally we would like a way describe the entire population's dynamics with as few parameters as possible. A recent example of such an approach is population coupling, the correlation over time between a single neuron's firing rate and the population average rate (Okun et al., 2015). But with dynamical networks we can use the same set of tools above, and more, to easily track changes to the population activity in time.

Figure 2 illustrates the idea of tracking non-stationary activity with data from a study by Peyrache et al. (2009). Rats were required to learn rules in a Y-maze to obtain reward. I use here a single session in which a rat learned the rule "go to the cued arm" (Figure 2A); 52 simultaneously recorded neurons from medial prefrontal cortex were active in every trial of this session. As the rat learned the rule in this session, we might reasonably expect the population activity to evolve. Visualising trial-by-trial changes using dynamical networks (built as in Figure 1A) shows a stabilisation of the interactions between neurons over trials (Figure 2B). Quantifying this by correlating weight matrices on consecutive trials (Figure 2C), confirms there was a rapid stabilisation of neuron interactions at the start of this learning session. These analyses thus track potentially learning-induced changes in the population activity of prefrontal cortex.

Clustering coefficient: Ratio of weighted triangles to weighted triples - incomplete triangles - in the network.

We can also use these data to illustrate the benefits we accrue from the null models in network theory. Figure 2D plots the average clustering coefficient ; for the dynamical networks, and we can see that it varies across trials. We can compare this to a suitable null model; here I use a null model that conserves node strength, but randomly re-assigns the set of weights between nodes (Rubinov and Sporns, 2011). Plotting the average clustering coefficient for this null model on each trial shows that the clustering in the data-derived dynamical networks is well in excess of that predicted by the null model: the neurons are more densely connected locally than predicted by just their total interactions with all neurons.

But the null model also shows that the average local clustering does not change over learning. Plotting the ratio of the data and model clustering coefficients shows that it is approximately constant (Figure 2E), showing that trial-by-trial variation in clustering is largely accounted for by variations in the overall interactions between neurons (one source of these might be finite-size effects in estimating the interactions on trials of different durations). So we can conclude that changes over learning in this population of neurons reflected a local reorganisation (Figure 2B) and stabilisation (Figure 2C) of interactions, but which did not change the population-wide distribution of clustering.

The rich potential for tracking dynamics with the readily-available metrics of network theory has not yet been tapped. As just demonstrated, with dynamical networks we can track trial-by-trial or event-by-event changes in population dynamics. For long recordings of spontaneous activity, building dynamical networks in time-windows slid over the recorded data allows us to track hidden shifts underlying global dynamics (Humphries, 2011). On slower time-scales, we can track changes during development of neural systems, either using ex-vivo slices (Dehorter et al., 2011) or in vitro cultures (Downes et al., 2012; Schroeter et al., 2015). These studies of development have all shown how maturing

neuronal networks move from seemingly random connectivity to a structured network.

Other tools from network theory could be readily re-purposed to track neural population dynamics. The growing field of network comparison uses distributions of network properties to classify networks (Guimera et al., 2007; Onnela et al., 2012). A particularly promising basis for comparison is the distributions of motifs (or graphlets) in the networks (Przulj, 2007). Re-purposed to track changes in dynamical networks, by comparing motif distributions between time-points, these would provide tangible evidence of changes to the information flow in a neural system. Ongoing developments in temporal networks (Holme, 2015) – networks that include links between the same nodes at different time-points – and network-based approaches to change-point detection algorithms (Peel and Clauset, 2014; Barnett and Onnela, 2016; Darst et al., 2016) also promise powerful yet tractable ways to track neural population dynamics.

Motifs:
A specific pattern of connections between a small number of nodes, which includes at least one connection for every node. For example, in an undirected network, for 4 nodes there are 5 possible motifs.

Network theory quantitatively defines computational concepts of neural populations

The mathematical framework of networks can also provide precise quantitative definitions of important but qualitative theories about neural populations. A striking example is the theory of neural ensembles (Harris, 2005). An ensemble is qualitatively defined as a set of neurons who are consistently co-active (Harris, 2005), thereby indicating they code or compute the same thing. This qualitative definition leaves open key quantitative questions: what defines co-active, and what defines consistent?

The network theory concept of modularity provides answers to these questions. Many networks are modular, organised into distinct groups: social networks of friendship groups, or collaboration networks of scientists. Consequently, the problem of finding modules within networks in an unsupervised way is an extraordinarily fecund research field (Fortunato and Hric, 2016). Most approaches to finding modules are based on the idea of finding the division of the network that maximises its modularity $Q = \{$number of links within a module$\}$ - $\{$expected number of such links$\}$ (Newman, 2006). Maximising Q thus finds a division of a network in which the modules are densely linked within themselves, and weakly linked between them.

Applied to dynamical networks, modularity defines neural ensembles (Humphries, 2011; Billeh et al., 2014; Bruno et al., 2015): groups of neurons that are more co-active with each other than with any other neurons in the population, given the choice of pairwise interaction used. Figure 3 demonstrates this idea using an example recording of 94 neurons from the motor circuit of the sea-slug *Aplysia* during fictive locomotion (Bruno et al., 2015). The weight matrix and network view in Figure 3A clearly indicate some structure within the dynamical network. Applying an unsupervised module-detection algorithm finds a high modularity division of the dynamical network (Figure 3B). When we plot the 94 spike-trains grouped by their modules in the dynamical network, the presence of multiple ensembles is clear (Figure 3C).

With this modularity-based approach, we can also easily check how robust these ensembles are to the choice of time-scale of co-activity. When computing pairwise interactions,

7

we often have a choice of temporal precision, such as bin-size or Gaussian width (Figure 1A): choosing small values emphasises spike-time precision; large values emphasise co-varying firing rates. As shown in Figure 3D, we can also use Q to look for time-scales at which the population dynamics are most structured (see Humphries, 2011, for more examples): this view suggests a clear peak time-scale at which the ensembles are structured. Nonetheless, we can also see a consistent set of modules at all time-scales: the weight matrix W at the smallest and largest Gaussian width are similar (Figure 3E); and the majority of neurons are placed in the same group at every time-scale (Figure 3F). Modularity not only defines ensembles, but also lets us quantify their time-scales and find consistent structure across time-scales.

As a final step, we can now quantitatively define a Hebbian cell assembly (Holtmaat and Caroni, 2016). By definition, a cell assembly is an ensemble of neurons that become co-active because of synaptic changes between them during learning. Thus, by combining the ideas of tracking dynamical networks and of module detection, we can test for the formation of assemblies: if we find dynamical network modules that appear during the course of learning, then we have identified potential cell assemblies.

Outlook

The dynamics of neural populations are emergent properties of the wiring within their microcircuits. We can of course use network theory to describe physical networks of the microcircuit too (Humphries et al., 2006; Lee et al., 2016; Schroeter et al., 2017), gaining insight into the mapping from wiring to dynamics. But dynamical networks need not map to any circuit. Indeed while dynamical networks are constrained by their underlying physical connections, they can change faster than their corresponding physical networks. A clear example is with the actions of neuromodulators - these can increase or decrease the effective strength of connections between neurons and the responsiveness of individual neurons (Nadim and Bucher, 2014), so changing the dynamical network without changing the underlying physical network. More broadly, rapid, global changes in brain state can shift the dynamics of a neural population (Zagha and McCormick, 2014). Thus, dynamical networks describing the simultaneous activity of multiple neurons capture the moment-to-moment changes in circuit dynamics.

Our motivation for turning to network theory as a toolbox for systems neuroscience is rooted in the extraordinarily rapid advances in recording technology, now scaling to hundreds or thousands of simultaneously recorded neurons (Stevenson and Kording, 2011). Capturing whole nervous systems will require scaling by further orders of magnitude (Ahrens et al., 2012; Lemon et al., 2015). And here is where network theory has its most striking advantage: these tools have been developed to address social and technological networks of millions of nodes or more, so easily scale to systems neuroscience problems now and in the foreseeable future.

This is not a one-way street. Systems neuroscience poses new challenges for network theory. Most network theory studies concern a handful of static or slowly changing data networks. Neural populations have non-stationary dynamics, that change rapidly com-
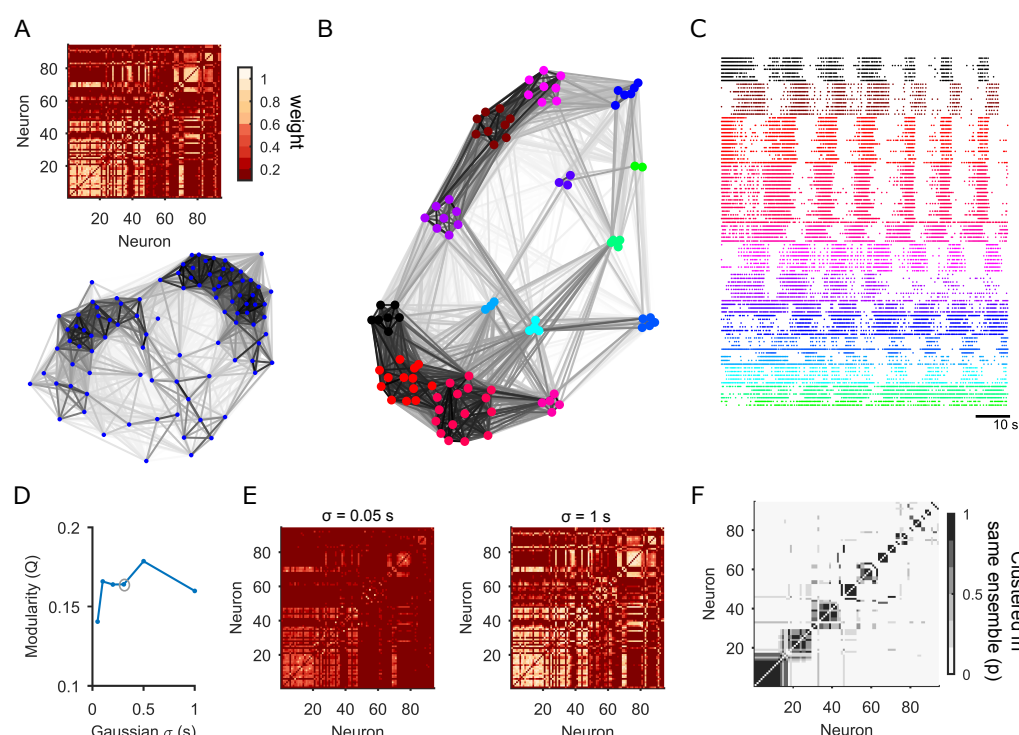
Figure 3: Defining and detecting neural ensembles using network theory.
A) A dynamical network of population dynamics during crawling in *Aplysia*. The weight matrix (top) and network view (bottom) for a simultaneous recording of 94 neurons during 90 seconds from the initiation of crawling (from the experimental protocol of Bruno et al., 2015). Weights are rectified Pearson's R between pairs of neurons convolved with a Gaussian of $\sigma = 0.306$ s. B) Modules within the dynamical network. Coloured nodes indicate different modules found within the dynamical network using an unsupervised consensus module-detection algorithm (Bruno et al., 2015). Placement of the modules reflects the similarity between them (Traud et al., 2009). C) Raster plot of the corresponding spike-trains, grouped according to the modules in panel B. The detection of multiple neural ensembles is evident. D) Dependence of the modular structure on the time-scale of correlation. Smaller Gaussian σ detects precise spike-timing; larger σ detects co-variation in firing rates. Circle: time-scale used in panels A-C. E) Weights matrices for the smallest and largest time-scale used for the Gaussian convolution. F) Stability of modules over time-scales. The confusion matrix showing for each pair of neurons the proportion of time-scales for which that pair was placed in the same module. The majority of neuron pairs were placed in the same module at every time-scale.

pared to the temporal resolution of our recordings. And systems neuroscience analysis requires quantitatively comparing multiple defined networks within and between brain regions, within and between animals, and across experimental conditions - stimuli, decisions, and other external changes. Bringing network theory to bear on challenges in systems neuroscience will thus create a fertile meeting of minds.

9

Supportive Information

Visualisations and analyses here drew on a range of open-source MATLAB (Mathworks, NA) toolboxes:

- Brain Connectivity Toolbox (Rubinov and Sporns, 2010): `https://sites.google.com/site/bctnet/`

- Network visualisations used the MATLAB code of Traud et al. (2009), available here: `http://netwiki.amath.unc.edu/VisComms`. This also the needs MatlabBGL library: `http://uk.mathworks.com/matlabcentral/fileexchange/10922-matlabbgl`. Mac OSX 64-bit users will need this version: `https://dgleich.wordpress.com/2010/07/08/matlabbgl-osx-64-bit/`

- Spike-Train Communities Toolbox (Humphries, 2011; Bruno et al., 2015): `https://github.com/mdhumphries/SpikeTrainCommunitiesToolBox`

Acknowledgments

References

Ahrens, M. B., Li, J. M., Orger, M. B., Robson, D. N., Schier, A. F., Engert, F., and Portugues, R. (2012). Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature*, 485:471–477.

Barnett, I. and Onnela, J.-P. (2016). Change point detection in correlation networks. *Scientific Reports*, 6:18893.

Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P. L., Gioanni, Y., Battaglia, F. P., and Wiener, S. I. (2010). Coherent theta oscillations and reorganization of spike timing in the hippocampal- prefrontal network upon learning. *Neuron*, 66(6):921–936.

Bettencourt, L. M. A., Stephens, G. J., Ham, M. I., and Gross, G. W. (2007). Functional structure of cortical neuronal networks grown in vitro. *Phys Rev E*, 75(2 Pt 1):021915.

Billeh, Y. N., Schaub, M. T., Anastassiou, C. A., Barahona, M., and Koch, C. (2014). Revealing cell assemblies at multiple levels of granularity. *J Neurosci Methods*, 236:92–106.

Boyden, E. S. (2015). Optogenetics and the future of neuroscience. *Nat Neurosci*, 18:1200–1201.

Bruno, A. M., Frost, W. N., and Humphries, M. D. (2015). Modular deconstruction reveals the dynamical and physical building blocks of a locomotion motor program. *Neuron*, 86:304–318.

Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nat Neurosci*, 7:446–451.

Chen, T.-W., Wardill, T. J., Sun, Y., Pulver, S. R., Renninger, S. L., Baohan, A., Schreiter, E. R., Kerr, R. A., Orger, M. B., Jayaraman, V., Looger, L. L., Svoboda, K., and Kim, D. S. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499:295–300.

Chung, F. and Lu, L. (2002). Connected components in random graphs with given expected degree sequences. *Annals of Combinatorics*, 6:124–145.

Darst, R. K., Granell, C., Arenas, A., Gómez, S., Saramki, J., and Fortunato, S. (2016). Detection of timescales in evolving complex systems. *Scientific Reports*, 6:39713.

Dehorter, N., Michel, F., Marissal, T., Rotrou, Y., Matrot, B., Lopez, C., Humphries, M., and Hammond, C. (2011). Onset of pup locomotion coincides with loss of NR2C/D-mediated corticostriatal EPSCs and dampening of striatal network immature activity. *Front Cell Neurosci*, 5:24.

Deisseroth, K. (2015). Optogenetics: 10 years of microbial opsins in neuroscience. *Nat Neurosci*, 18(9):1213–1225.

Downes, J. H., Hammond, M. W., Xydas, D., Spencer, M. C., Becerra, V. M., Warwick, K., Whalley, B. J., and Nasuto, S. J. (2012). Emergence of a small-world functional network in cultured neurons. *PLoS Comput Biol*, 8(5):e1002522.

Fortunato, S. and Hric, D. (2016). Community detection in networks: A user guide. *Physics Reports*, 659:1–44.

Fosdick, B. K., Larremore, D. B., Nishimura, J., and Ugander, J. (2016). Configuring random graph models with fixed degree sequences. *arXiv*, page 1608.00607.

Gerhard, F., Pipa, G., Lima, B., Neuenschwander, S., and Gerstner, W. (2011). Extraction of network topology from multi-electrode recordings: is there a small-world effect? *Front Comput Neurosci*, 5:4.

Guimera, R., Sales-Pardo, M., and Amaral, L. A. N. (2007). Classes of complex networks defined by role-to-role connectivity profiles. *Nat Phys*, 3(1):63–69.

Harris, K. D. (2005). Neural signatures of cell assembly organization. *Nat Rev Neurosci*, 6:399–407.

Harris, K. D., Quiroga, R. Q., Freeman, J., and Smith, S. L. (2016). Improving data quality in neuronal population recordings. *Nat Neurosci*, 19:1165–1174.

Holme, P. (2015). Modern temporal network theory: A colloquium. *Eur Phys J B*, 88:234.

Holtmaat, A. and Caroni, P. (2016). Functional and structural underpinnings of neuronal assembly formation in learning. *Nat Neurosci*, 19:1553–1562.

Huber, D., Gutnisky, D. A., Peron, S., O'Connor, D. H., Wiegert, J. S., Tian, L., Oertner, T. G., Looger, L. L., and Svoboda, K. (2012). Multiple dynamic representations in the motor cortex during sensorimotor learning. *Nature*, 484(7395):473–478.

Humphries, M. D. (2011). Spike-train communities: finding groups of similar spike trains. *J Neurosci*, 31:2321–2336.

Humphries, M. D., Gurney, K., and Prescott, T. J. (2006). The brainstem reticular formation is a small-world, not scale-free, network. *Proc Roy Soc B. Biol Sci*, 273:503–511.

Jun, J. J., Mitelut, C., Lai, C., Gratiy, S., Anastassiou, C., and Harris, T. D. (2017). Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction. *bioRxiv*, page 101030.

Latora, V. and Marchiori, M. (2001). Efficient behaviour of small-world networks. *Phys Rev Lett*, 87:198701.

Lee, W.-C. A., Bonin, V., Reed, M., Graham, B. J., Hood, G., Glattfelder, K., and Reid, R. C. (2016). Anatomy and function of an excitatory network in the visual cortex. *Nature*, 532:370–374.

Lemon, W. C., Pulver, S. R., Höckendorf, B., McDole, K., Branson, K., Freeman, J., and Keller, P. J. (2015). Whole-central nervous system functional imaging in larval Drosophila. *Nat Commun*, 6:7924.

Miesenböck, G. (2009). The optogenetic catechism. *Science*, 326(5951):395–399.

Nadim, F. and Bucher, D. (2014). Neuromodulation of neurons and synapses. *Curr Opin Neurobiol*, 29:48–56.

Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45:167–256.

Newman, M. E. J. (2004). Analysis of weighted networks. *Phys Rev E*, 70:056131.

Newman, M. E. J. (2006). Finding community structure in networks using the eigenvectors of matrices. *Phys Rev E*, 74:036104.

Newman, M. E. J. and Martin, T. (2014). Equitable random graphs. *Phys Rev E*, 90:052824.

Okun, M., Steinmetz, N. A., Cossell, L., Iacaruso, M. F., Ko, H., Barth, P., Moore, T., Hofer, S. B., Mrsic-Flogel, T. D., Carandini, M., and Harris, K. D. (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature*, 521:511–515.

Onnela, J.-P., Fenn, D. J., Reid, S., Porter, M. A., Mucha, P. J., Fricker, M. D., and Jones, N. S. (2012). Taxonomies of networks from community structure. *Phys Rev E*, 86:036104.

Otchy, T. M., Wolff, S. B. E., Rhee, J. Y., Pehlevan, C., Kawai, R., Kempf, A., Gobes, S. M. H., and Ölveczky, B. P. (2015). Acute off-target effects of neural circuit manipulations. *Nature*, 528:358–363.

Palowitch, J., Bhamidi, S., and Nobel, A. B. (2016). The continuous configuration model: A null for community detection on weighted networks. *arXiv*, page 1601.05630.

Peel, L. and Clauset, A. (2014). Detecting change points in the large-scale structure of evolving networks. *arXiv*, page 1403.0989.

Peron, S., Chen, T.-W., and Svoboda, K. (2015a). Comprehensive imaging of cortical networks. *Curr Opin Neurobiol*, 32:115–123.

Peron, S. P., Freeman, J., Iyer, V., Guo, C., and Svoboda, K. (2015b). A cellular resolution map of barrel cortex activity during tactile behavior. *Neuron*, 86:783–799.

Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I., and Battaglia, F. P. (2009). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat Neurosci*, 12:916–926.

Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., and Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454:995–999.

Pouget, A., Beck, J. M., Ma, W. J., and Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nat Neurosci*, 16(9):1170–1178.

Przulj, N. (2007). Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23:e177–e183.

Rieke, F., Warland, D., de Ruyter van Stevninck, R., and Bialek, W. (1999). *Spikes: Exploring the neural code*. Cambridge MA: MIT Press.

Robins, G., Pattisona, P., Kalisha, Y., and Lushera, D. (2007). An introduction to exponential random graph (p*) models for social networks. *Social Networks*, 29:173–191.

Rubinov, M. and Sporns, O. (2010). Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, 52:1059–1069.

Rubinov, M. and Sporns, O. (2011). Weight-conserving characterization of complex functional brain networks. *Neuroimage*, 56:2068–2079.

Sadovsky, A. J. and MacLean, J. N. (2013). Scaling of topologically similar functional modules defines mouse primary auditory and somatosensory microcircuitry. *J Neurosci*, 33:14048–60.

Schreiber, T. (2000). Measuring information transfer. *Phys Rev Lett*, 85:461–464.

Schroeter, M., Paulsen, O., and Bullmore, E. T. (2017). Micro-connectomics: probing the organization of neuronal networks at the cellular scale. *Nat Rev Neurosci*, page in press.

Schroeter, M. S., Charlesworth, P., Kitzbichler, M. G., Paulsen, O., and Bullmore, E. T. (2015). Emergence of rich-club topology and coordinated dynamics in development of hippocampal functional networks in vitro. *J Neurosci*, 35:5459–5470.

Singh, A. and Lesica, N. A. (2010). Incremental mutual information: a new method for characterizing the strength and dynamics of connections in neuronal circuits. *PLoS Comput Biol*, 6:e1001035.

Srinivas, K. V., Jain, R., Saurav, S., and Sikdar, S. K. (2007). Small-world network topology of hippocampal neuronal network is lost, in an in vitro glutamate injury model of epilepsy. *Eur J Neurosci*, 25:3276–3286.

Stevenson, I. H. and Kording, K. P. (2011). How advances in neural recording affect data analysis. *Nat Neurosci*, 14:139–142.

Thivierge, J.-P. (2014). Scale-free and economical features of functional connectivity in neuronal networks. *Phys Rev E*, 90:022721.

Traud, A. L., Frost, C., Mucha, P. J., , and Porter, M. A. (2009). Visualization of communities in networks. *Chaos*, 19:041104.

Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., and Brown, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol*, 93:1074–1089.

Vincent, K., Tauskela, J. S., Mealing, G. A., and Thivierge, J.-P. (2013). Altered network communication following a neuroprotective drug treatment. *PloS One*, 8:e54478.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393:440–442.

Wohrer, A., Humphries, M. D., and Machens, C. (2013). Population-wide distributions of neural activity during perceptual decision-making. *Prog Neurobiol*, 103:156–193.

Yu, S., Huang, D., Singer, W., and Nikolic, D. (2008). A small world of neuronal synchrony. *Cereb Cortex*, 18:2891–2901.

Yuste, R. (2015). From the neuron doctrine to neural networks. *Nat Rev Neurosci*, 16:487–497.

Zagha, E. and McCormick, D. A. (2014). Neural control of brain state. *Curr Opin Neurobiol*, 29:178–186.

Zanin, M., Sousa, P., Papo, D., Bajo, R., Garca-Prieto, J., del Pozo, F., Menasalvas, E., and Boccaletti, S. (2012). Optimizing functional network representation of multivariate time series. *Sci Rep*, 2:630.