
8 Analyzing Emergence in Biological Neural Networks Using Graph Signal Processing

*Kevin Schultz, Marisel Villafaña-Delgado,
Elizabeth P. Reilly, and Anshu Saksena*
Johns Hopkins University Applied
Physics Laboratory (JHUAPL)

Grace M. Hwang
Johns Hopkins University Applied
Physics Laboratory (JHUAPL)
Kavli Neuroscience Discovery Institute
Johns Hopkins University

CONTENTS

8.1	Introduction	172
8.2	Biological Neural Networks	173
8.2.1	The Neuron	174
8.2.2	Neural Networks	176
8.3	Graph Signal Processing.....	177
8.3.1	Introduction to Graph Theory	177
8.3.2	Introduction to GSP	178
8.3.3	GSP in Neuroscience	180
8.4	Emergence in the Short Term: Collective Activity.....	180
8.4.1	Case Study: Slow Switching Assemblies.....	182
8.4.2	Case Study: Fruit Fly Protocerebral Bridge	183
8.4.3	Summary	186
8.5	Emergence in the Long Term: Network Formation and Learning.....	186
8.5.1	Case Study: Structured Assembly Formation.....	187
8.6	Conclusion	188
	Acknowledgments.....	190
	References.....	190

8.1 INTRODUCTION

Emergent behavior is found throughout nature, including in biological neural networks where individual neurons connect and interact to form complex thought processes and task performance. We define *emergence* of a system to be the existence of complex system-level behavior that is not present in the much simpler individual entities of the system. The behavior of individual entities is locally defined and not directly tied to the system-level behavior. For instance, a bird in a flock might fly with the objective of being within a certain distance of approximately the six nearest birds, and yet, when viewed as a whole, this flock of birds can form ever-changing patterns with no centralized control.

In this chapter, we consider two different, yet related forms of emergence in biological neural networks. The first is the coordinated activity of neurons such as activity “bumps” with sustained and/or localized activity or synchronizing oscillations. These can be both self-sustaining, in the sense that the activity persists in the absence of stimuli, or transient behaviors brought about by a particular sequence of stimuli and vanish or decay after. These phenomena are often studied via relatively simple dynamical models of neurons interacting through a *static* network topology. These emergent behaviors are generally quite sensitive to both the parameters of the neuronal dynamics and the network interconnections between neurons. How, then, can large, complex neural networks form in a way that is highly dependent upon the experiences (i.e., stimuli) of the organism in question? The answer to this leads to the second form of emergence we consider. The process of network formation in a neural network over time is known as plasticity or learning. In this process, the connections between neurons (i.e., synapses) are strengthened or weakened over time based on the activity between the neurons in question. Thus, we have a dependency loop between two time scales of scales of emergence: The structure of the network is key to the emergent collective activity from individual neurons, and correlations in the neuronal activity drive the formation of the network.

In line with the overall theme of this volume, it is instructive to consider just how much the capabilities of an organism’s brain far outstrip the capabilities of an individual neuron, or even a collection of neurons serving as a functional component (i.e., the sub-systems of the system of systems (SoS) we call the brain). As we will discuss below, an individual neuron is often modeled using a first-order differential equation and communicates with a single-bit channel (i.e., a spike train) to its neighbors in the network. Despite this apparent simplicity, modest collections of hundreds of thousands of neurons can form networks that serve essential navigational functions such as head direction and position estimation as part of the hippocampal formation in mammals, including humans (Poulter, Hartley, and Lever 2018). These are a few of the functions of the hippocampal formation, which itself interacts with many other brain regions and comprises a small fraction of the total number of neurons present in the brain. Thus, it is quite clear that the overall dynamics of the brain can be viewed as a complex SoS with emergent behaviors that are driven by the interactions of different functional components. These functional components, in turn, have emergent behaviors that are driven by the interactions of the individual neurons over their network interconnections.

Given that we have cast these emergent phenomena as a set of interacting processes on a network topology, we propose that graph signal processing (GSP) will serve as a useful analysis tool here. GSP builds on its roots in spectral graph theory (Chung and Graham 1997) and algebraic signal processing (Puschel and Moura 2008) to generalize techniques from classical signal processing to signals defined on irregular domains by graphs (Sandryhaila and Moura 2013; Shuman et al. 2013). Here, we propose that the outputs (or other state) of the neurons should be treated as the signal of interest with the network interconnections defining the graph structure. GSP techniques have been applied across various non-invasive neuroimaging techniques, including functional magnetic resonance imaging (fMRI) and electroencephalography. Goldsberry et al. (2017) and Medaglia et al. (2018) introduced GSP analysis for the joint study of structure and functional data from diffusion tensor imaging and fMRI, respectively, studying the alignment between functional activity and anatomical network in cognitive flexibility.

The study of emergence in biological neural networks is relevant to a broad range of application areas. Obviously, understanding the collective behavior of neurons is of fundamental importance in biology and medicine. Additionally, neural networks are the backbone of many advances in machine learning (Schmidhuber 2015) and are the inspiration for neuromorphic computing (Furber 2016), both of which hold further promise for continuing to revolutionize computation. Finally, to directly tie neural networks to a classic area of emergence, we note that neuronal dynamics have served as both a motivating analogy in swarm intelligence (Trianni et al. 2011) and also as a direct source of swarm dynamics (Monaco et al. 2020). To this last point, we note that graph theoretical techniques have been widely applied in the analysis of swarming dynamics (e.g., Tanner, Jadbabaie, and Pappas 2007), further reinforcing the use of GSP as an analytical tool for emergence.

In the following, we first introduce some basic models from computational neuroscience that will serve as the signals in our GSP analysis. Next, we review the basics of GSP and discuss in some detail the unique challenges presented by biological constraints. We then move to the discussion of short-term emergence in neural networks in the form of collective activity and analyze an example emergent phenomena in the form of spontaneous collective firing in a structured network. Next, we introduce the basics of network formation through plasticity and discuss some relevant examples of emergence that tie the collective dynamics of neurons to the formation of networks. We then consider a specific example of network formation that leads to a similar structured network previously considered and analyze its long-term behaviors using GSP. We conclude with discussion and lessons learned.

8.2 BIOLOGICAL NEURAL NETWORKS

The human brain is composed of about 85 billion neurons (nerve cells) and 85 billion glia cells. For decades, neurons were thought to be exclusively involved in the computation of the brain through their electrical properties. Emerging evidence reveal that chemicals, or more specifically neuromodulators, are involved in neuronal computation and that glia cells, in particular at least one form of glia cells, the microglial

can produce negative feedback similarly to the role of inhibitory neurons (Badimon et al. 2020). In this chapter, a simplified mathematical abstraction of a biological neuronal network is described in the context of emergence. Thus details of neuromodulators, glia cells, and non-linearities within a neuron exhibited at particular dendritic branches (Gidon et al. 2020) that reveal how a single neuron can itself be modeled as a neuronal network capable of achieving many logical computations – AND, OR, XOR – will be ignored. Instead, we introduce some high-level and general concepts about various systems and processes in biological neuronal networks. To achieve this, we pursue the discussion at the level abstraction appropriate for computational or theoretical neuroscience; that is, at the level of relatively simple equations and interactions. For a more in-depth introduction into computational neuroscience, we refer the readers to full texts such as Dayan and Abbott (2001) or Miller (2018). For an accessible overview of a self-organized viewpoint of cognition, see Buzsáki (2019).

8.2.1 THE NEURON

The neuron is arguably the fundamental processing unit of biological neuronal networks and artificial neural networks (ANN). However, to speak of “the neuron” does a disservice to the breadth of form and function of neurons that occur in biology. At its core, a biological neuron is a cell with branches of dendrites (inputs), connected to a soma or cell body (processing), which is connected to one axon (output) that can transmit neural activity (information) to a few thousand other downstream neurons via synapses. Dendrites, which are typically covered by synapses, can extend far from the soma by hundreds of microns to receive inputs via synapses from many upstream neurons. Dendrites have been shown to critically contribute to the non-linear computations performed by neurons (Poirazi and Papoutsi 2020; Gidon et al. 2020). The soma is connected to one axon, that can extend over 1 m in humans eventually branching out, to potentially transmit neural activity to other neurons. Each biological neuron has a state-dependent, adaptive, electrical threshold that when exceeded allows neural activity to propagate through the axon to other downstream neurons. This transmission is done via an electro-chemical process in the terminus of the axon that releases tiny vesicles containing neurotransmitters that can lead to a change in the postsynaptic membrane voltage. One can think of the adaptive threshold of a real biological neuron to be an emergent property within the neuron. In this chapter, single-compartment, point-like neuron models will be discussed in the context of emergence; these models are devoid of any intracellular compartments (e.g., dendrites, axons, vesicle) or components (e.g., neurotransmitters).

This chapter will also discuss typical ANN models, which assume dendrites to be passive linear receivers of neural activity, and thus all synaptic inputs are homogeneously summed within a point-like neural unit. This total neural input is then transformed by a non-linear threshold that represents the net effect of computation in biological neurons. Therefore, ANN models also exclude dendrites and axons. The weighted connection between neural units, however, is referred to as a “synapse,” despite the lack of dendrites or any other aspect of a biological synapses besides its strength. Another difference between real neurons and ANNs is that while artificial neural units can have positive and negative weights onto their targets, biological

neurons can only make one type of connection to downstream cells, either positive (excitatory neurons) or negative (inhibitory neurons). This division of connection valence according to cell type is known as Dale’s law (Strata and Harvey 1999). Finally, it should be noted that synapses in ANN models are inspired by chemical synapses as described above. The biological brain also has electrical synapses, based on a physical connection called a gap junction, that allow neurons to communicate directly by sharing membrane voltage. While electrical synapses have also been largely ignored by ANN models, further discussion is out of scope for this chapter. Again, a comprehensive review is beyond the scope of this text, and we will instead focus at a more abstract level.

A neuron can be modeled based on its electrical activity, arising from ion flow in and out of the neuronal cell membrane, which leads to a voltage potential difference that might lead to the generation of an action potential (spike) that will travel the length of the axon. The leaky integrate-and-fire (LIF) model is one of many ways to model a neuron (Dayan and Abbott 2001; Miller 2018), a relatively simple one that is nevertheless capable of producing emergent behaviors in the case studies below, while still remaining computationally tractable for larger network sizes. The “leaky-integration” portion of the LIF models the membrane potential V via dynamics

$$C \frac{dV(t)}{dt} = I(t) - \frac{V(t)}{R} \quad (8.1)$$

where C is the membrane capacitance, I is the neuron’s input, and R is the membrane resistance. The LIF model accumulates the input signal, and when V reaches some threshold potential V_i it “fires,” sending a spike of current to the downstream neurons (see next section) and resetting V to zero.

The LIF model can be generalized in a number of ways to add additional biological fidelity to neuronal dynamics. One common extension is to replace the linear integration of the input I with a non-linear term such as a quadratic or exponential rule. These non-linearities serve to essentially change the firing threshold based on the “shape” of the input $I(t)$, for example so that “fast” inputs trigger firing where slower inputs with similar area do not. Additionally, these non-linear models can be modified to introduce refractory periods, where the accumulation rate is dependent on the time since the last firing event. This latter behavior fundamentally limits the overall output rate of the neuron, potentially introducing stability to the overall network. Additional modifications for adaptation can be introduced to capture firing-rate patterns observed in nature, such as initial bursts followed by limited activity, increasing delays between spikes, delayed/transient responses, etc. We also note that the original LIF model and the extensions above have been introduced in terms of deterministic behaviors, and there are a number of mechanisms to introduce randomness to these models. One such simple mechanism is to have the firing and reset behavior of the LIF model to be based on some random firing process with the probability of firing proportional to the membrane potential V . In the context of more general emergent phenomena, the introduction of this sort of randomness introduced variability into the system that may prevent degeneracies in behavior. In summary, computational neuroscience offers a variety of modifications to the simple neuronal

dynamics of Equation 8.1 that adapt or regulate the behavior of an individual neuron that could have a profound impact on the collective behavior of the network.

While the analysis below will primarily focus on the basic LIF model of a neuron as described above, we now briefly discuss an alternative class of models of neuronal dynamics, so-called rate-based models. These model a spike firing rate $\nu(t)$ (rather than generating spikes) and have dynamics of the form

$$\frac{d\nu(t)}{dt} = -\nu(t) + F(I(t)) \quad (8.2)$$

where F is some non-linearity, and I is again the neuron's input. The fixed points of these systems correspond to the condition $\nu(t) = F(I(t))$, which should be immediately familiar to machine learning practitioners who deal with ANNs as the input-output relationship of an artificial neuron. As was the case with the LIF neuron, there are many variants (in the form of different non-linearities F) that can capture different firing patterns observed in biological neurons.

8.2.2 NEURAL NETWORKS

Biological neurons can communicate with electro-chemical signals via chemical synapses. In the context of a simplified neuron model described above, synapses are generally modeled as a set of “weights” w_{ij} between neurons. If V_i denotes the LIF model for the i th neuron in a network of N neurons, then we can adapt the above LIF dynamics to account for the network structure

$$C \frac{dV_i(t)}{dt} = I(t) + \sum_{j=1}^N w_{ij} S_j(t) - \frac{V_i(t)}{R} \quad (8.3)$$

where $S_j(t)$ is the impulsive “spike-train” output of neuron j which causes the voltage V_i to instantaneously change by w_{ij} when neuron j fires. In reality, a synapse will introduce additional noise, non-linearity, transport delay, etc., not captured by this linear relationship. Firing-rate models are adapted similarly, with the combination of exogenous input and network inputs occurring inside a non-linearity F .

At this level of abstraction, one can analyze a neural network using the branch of mathematics known as graph theory, which we discuss in the following section. Models of network processes that use graphs are often modeled as symmetric $w_{ij} = w_{ji}$, however, such a model is not considered biologically plausible as synapses are naturally directional (requiring a directed graph model). As noted above, Dale's law states that a given neuron is either *excitatory*, meaning its spiking should produce an increase in spiking on its post-synaptic neighbors, or *inhibitory*, which tends to suppress firing in its post-synaptic neighbors. From the perspective of our network model, this implies that our network contains both negative and positive weights, corresponding to inhibitory and excitatory neurons, again violating a common approach to modeling networks as consisting entirely of non-negative weights. Furthermore, as this is a property of the neuron and not the synapse, for a given neuron j , w_{ij} must

all be either non-positive (inhibitory) or non-negative (excitatory), inducing additional structure on our network model.

Taking a step back from the precise mathematical model described above, one can imagine the potential interactions between neurons (or indeed, agents in other SoS emergence scenarios) might be broadly impacted by the notion of excitatory and inhibitory interactions. Excitatory interactions result in increased activity in the network, whereas inhibitory interactions result in decreased activity in the network. These general concepts and the presence of non-linearities in the models suggest that there will likely be a delicate balance between the excitatory and inhibitory neurons that can generate the desired “positive” emergence in a neural network (i.e., the neural activity that supports its function). Any imbalances might lead to a lack of neuronal activity (due to too much inhibition) or saturation of neuronal activity (due to too little inhibition).

8.3 GRAPH SIGNAL PROCESSING

Over the past decade, the field of GSP has expanded across multiple applications and novel techniques have enabled analyses otherwise not possible. In this section, we provide the foundations for graph theory needed for GSP, followed by an overview of GSP, and will conclude with a review of graph spectral analysis and GSP in neuroscience.

8.3.1 INTRODUCTION TO GRAPH THEORY

First, we define a (undirected) graph $G = (V, E)$, where V is a set of vertices or nodes and $E = \{\{u, v\} : u, v \in V\}$ is a set of edges representing relationships among the vertices. Note that a graph is called simple if there are no self-loops ($\{u, u\} \notin E$ for all $u \in V$) and multiple edges do not exist between a single pair of vertices. For an edge $\{u, v\} \in E$, we say that u and v are adjacent and that v is a neighbor of u . Both vertices and edges may have attributes associated with them. For instance, a vertex attribute in a co-authorship graph may be a label (e.g., name of the author represented by the vertex) or a numerical value (e.g., how long the author has been publishing, how many papers they have published, etc.). An edge attribute describes the relationship. The most common edge attribute is a weight indicating the strength of the relationship. In the co-authorship network example, each edge might be weighted by the number of papers co-authored by two people. The degree of a vertex is the number of neighbors of the vertex. For $v \in V$, $deg(v) = |\{u : \{u, v\} \in E\}|$.

The adjacency matrix of a graph, indicated by A , represents the edge relationships among the vertices with $A_{ij} = 1$ if $\{v_i, v_j\} \in E$ and $A_{ij} = 0$ otherwise. For a weighted graph, $A_{ij} = w_{ij}$ where w_{ij} is the weight of the edge $\{v_i, v_j\}$. Note that if G is simple, then $A_{ii} = 0$ for all i . Additionally, A is symmetric when G is undirected. For a simple undirected graph, the Laplacian matrix is defined as $L = D - A$, where D is a diagonal matrix with $D_{ii} = deg(v_i)$, the degree of vertex v_i and A is the adjacency matrix. The Laplacian matrix is a rich representation of a graph, as it encodes several interesting properties of the graph such as the number of spanning trees, the number of connected components, and the overall strength of connectedness of the graph.

For undirected graphs with nonnegative, real-valued weights, both the adjacency and Laplacian matrices are real and symmetric, meaning they have n (not necessarily unique) real eigenvalues and one can find a corresponding set of n orthonormal eigenvectors, $L = U\Lambda U^T$. While the eigenvalues of the adjacency matrix sum to 0, the Laplacian matrix is positive semi-definite and thus has non-negative eigenvalues.

A graph is directed if the edges have directionality. We use parentheses to indicate a directed edge so the edge set is $E = \{(u, v), u, v \in V\}$. Then, u is referred to as the head and v is the tail. In-degree is the number of edges coming into a vertex, denoted $d^-(v) = |\{u : (u, v) \in E\}|$. Similarly, out-degree is the number of edges leaving a vertex, denoted $d^+(v) = |\{v : (u, v) \in E\}|$. Moreover, a signed directed graph is a graph in which its edges e_{ij} can take both positive and negative values. In addition to biological neural networks, signed directed graphs are observed in multiple applications, such as social networks.

8.3.2 INTRODUCTION TO GSP

The main motivation of GSP is to analyze signals over a graph, where signals now live on an irregular domain (Shuman et al. 2013). A graph signal $x \in \mathbb{R}^N$ is defined over a graph G (with N vertices) and the nodes of the graph form the domain of the signal. An example graph with corresponding signal is depicted in Figure 8.1. The field of GSP has focused on extending techniques from classical signal processing into signals defined over graphs. Note that now the emphasis is no longer on time, but

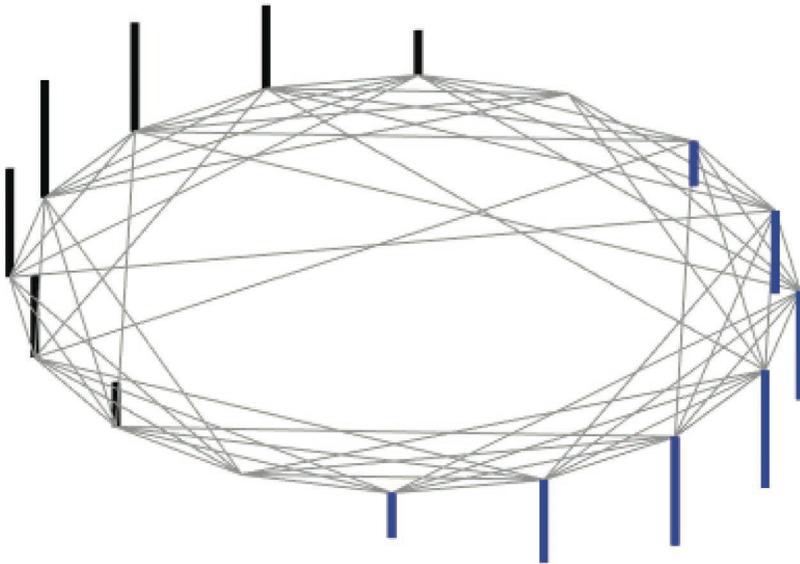


FIGURE 8.1 Illustration of graph signals over a Watts-Strogatz network. The thin lines are the edges between the nodes, and the graph signal is depicted by the thick vertical lines protruding from the top of the node for positive signal components or the bottom for negative signal components

on how the signal varies over the nodes of a graph. An important transform applied to graph signals is the graph Fourier transform (GFT), generally defined as

$$\hat{x} = U^T x, \quad (8.4)$$

where $x \in \mathbb{R}^N$ is the graph signal and $U \in \mathbb{R}^N \times \mathbb{R}^N$ is a matrix whose columns define the graph Fourier basis functions, or harmonics. Typically, these are the eigenvectors of the graph Laplacian, but the eigenvectors of the adjacency matrix can also be used as basis functions. In this chapter, the GFT of x , \hat{x} is defined over the eigenvalues of L , $\lambda_0 \leq \lambda_1 \cdots \leq \lambda_{N-1}$. Now the eigenvalues of the graph Laplacian define the frequencies of the signal. For undirected, unsigned graphs, the first eigenvector u_0 , corresponding to the smallest eigenvalue, is constant, an analog to a zero frequency component in classical signal processing. Eigenvectors corresponding to larger eigenvalues oscillate faster. This is a convenient property of the Laplacian spectrum for graph frequencies, which follows analog definitions as in classical Fourier analysis.

In multiple applications, the graph signals also vary over time, defining a new signal $X \in \mathbb{R}^{N \times T}$, i.e., the N nodes are discretely sampled at T time points. To analyze graph signals that vary over the graph and time, the joint vertex-time (JVT) transform extends the GFT (Grassi et al. 2017)

$$\hat{X}(l, k) = \frac{1}{\sqrt{T}} \sum_{n=1}^N \sum_{t=1}^T X_{n,t} u_{l,n}^* e^{-jvk t}, \quad (8.5)$$

where u_l is the l th eigenvector of the graph Laplacian and $e^{-jvk t}$ is the Fourier basis. Alternatively, in matrix form

$$\hat{X} = JVT\{X\} = U_G^\dagger X U_F^T \quad (8.6)$$

where U_G consists of the graph Laplacian eigenvectors, and U_F is the DFT matrix of appropriate dimension. The JVT essentially computes the classical Fourier transform over time of the GFT over the graph nodes.

Unlike the graphs used in many GSP applications, in biological neural networks, the graph model is both directed (due to the one-way directionality of the synapses) and signed (due to the presence of both excitatory and inhibitory neurons). From a spectral graph theory perspective, this presents several challenges as the graph Laplacian is no longer diagonalizable into an orthonormal basis for the transforms. These spectral conditions then impact the interpretation and intuition derived from the GFT. For directed graphs, multiple techniques have been proposed to obtain an orthonormal basis from the graph Laplacian or the adjacency matrix. Those methods involve either optimization approaches that impose certain constraints to derive the basis function, or approaches that propose novel representations of the graph into matrices that can produce an orthonormal basis, such as the Hermitian Laplacian Furutani et al. (2019). Here we employ the signed Hermitian Laplacian

$$L_q = D - \Gamma_q \odot A^{(s)}, \quad (8.7)$$

where D is the degree matrix of a symmetrized graph, $\Gamma_{qij} = e^{j2\pi q(w_{ij}-w_{ji})}$, \odot denotes elementwise multiplication, and $A^{(s)}$ is the symmetrized adjacency matrix, i.e., $A_{ij}^{(s)} = \frac{1}{2}(w_{ij} + w_{ji})$.

8.3.3 GSP IN NEUROSCIENCE

In recent years, network neuroscience has emerged as a powerful tool for the study of neuronal networks across multiple scales. In particular, graph spectral analysis has been used for graph comparison, graph embeddings, and structure-function analysis. One illustration of the use of graph spectral analysis for network comparison is the work of de Lange, de Reus, and Van Den Heuvel (2014) where the connectomes from the neuronal networks of the macaque, cat, and *C. elegans* were compared to model and empirical networks by using a similarity metric based on the spectral distance from a smoothed eigenvalue distribution from the normalized Laplacian. In another work, Raj et al. (2020) developed a spectral graph model based on the spectrum of the Laplacian from the structural connectome to derive a closed-form solution to the structure-function problem. The graph spectra has been also employed in the construction of spectral graph embeddings to determine the importance of cells in *C. elegans* (Petrovic et al. 2019). Finally, Aqil et al. (2021) developed a spatiotemporal framework of dynamical models in the human connectome based on the eigenvectors of the human connectome Laplacian to study structure-function relationships.

Earlier work in GSP for neuroscience applications focused mostly on macro-scale neuroscience. Various works focused on the study of alignment Medaglia et al. (2018) and the introduction of GSP for neuroimaging data (Huang et al. 2018; Goldsberry et al. 2017). GSP wavelets have also provided significant contributions to the study of macro-scale connectivity in the human brain (Leonardi and Van De Ville 2013), including novel ways to construct connectomes (Behjat et al. 2015). Other works focused on the study of fMRI temporal analysis using GSP techniques (Brahim and Farrugia 2020).

8.4 EMERGENCE IN THE SHORT TERM: COLLECTIVE ACTIVITY

Collective activity in neural networks can take place in many forms, including activity bumps, oscillations, and traveling waves. Activity bumps are when a subset of the neurons fire in response to an external (from the perspective of a given neural network) input signal. This type of behavior should be familiar to machine learning practitioners where, for example, inputs to convolutional neural networks produce feature extraction as activity bumps across the hidden layers (Lindsay 2021). Unsurprisingly, since these networks and related structures derive motivation from the visual cortex, activity bumps are prevalent there and other sensory receptive fields. Some forms of activity bumps maintain their activity even if the external input is removed, resulting in so-called self-sustaining activity bumps. This form of collective behavior is especially prevalent in the hippocampal formation where spatial computation and estimation is performed in the form of place, head direction, and grid cells, among others (Knierim and Zhang 2012). A leading theory of these stable activities is *attractor*

theory where the emergent bumps are controlled by these external inputs, yet the bumps remain active even when these inputs are removed due to “attraction” of the dynamics to the low energy states of the network. This theory was originally formulated as a mechanism for memory (Hopfield 1982; Amit and Treves 1989) that could store discrete patterns as fixed points of the network dynamics. However, as space is itself continuous, attractor theory was soon extended to so-called continuous attractors (Samsonovich and McNaughton 1997), where instead of stable fixed points, the dynamics of the network are attracted to stable sub-manifolds of equal energy.

The above emergent phenomena are analogous to collective behaviors in space, since the collective firing is occurring roughly simultaneously, in response to some external input (or persistently, in the case of self-sustaining activity bumps). There are also examples of emergent collective firing that are more time-oriented, such as oscillations. There are numerous emergent oscillations in the brain ranging from 0.02 to 600 Hz (Penttonen and Buzsáki 2003). These oscillations are important because they form a hierarchical framework for action potentials to traffic within and across neuronal circuits at many temporal scales (Buzsáki 2019; Monaco, Rajan, and Hwang 2021). When viewed together, these oscillation bands form a linear progression on a natural log scale, spanning ten frequency bands. These frequency bands can co-exist and interact with each other in the brain in the same or different structures giving rise to various brain states (e.g., task engagement or sleep). Many of these frequency bands are thought to be nested in which the phase of the slower oscillation modulates the amplitude of a faster oscillation, and in turn that phase of the faster oscillation modulates the amplitude of the even faster oscillation and so on. Many of these frequency bands have been observed across many species, and some have been given names. In rodent studies, entrainment of theta oscillations (4–10 Hz) is required to enable movement (Fuhrmann et al. 2015), while the frequency of theta oscillations modulate movement speed (McNaughton, Barnes, and O’Keefe 1983). In contrast, sharp wave ripples (100–200 Hz) are known to occur transiently during deliberate moments of immobility (Pfeiffer and Foster 2013) in spatial memory tasks. Monaco et al. (2021) proposed that these hierarchically nested oscillations are reentrant flows on recurrent networks that can form a new computational basis. Sharp wave ripples are an example of a traveling wave (that happens to be oscillatory), where a set of neurons fire in sequence.

In the following sections, we consider a pair of simulation examples from the computational neuroscience literature to illustrate the utility of GSP in understanding the collective behaviors of the neurons. These examples focus on the analysis of activity bumps as these are more analogous to the spatial motivation of GSP, as opposed to purely oscillatory behaviors that are readily identified by standard time-domain (Fourier) analysis. That said, the combination of GSP and standard Fourier analysis is a powerful tool for understanding joint spatial and temporal dynamics. To this end, in the latter example, we consider a network that is driven by an oscillatory input, resulting in oscillatory outputs that give a sense of the potential for the JVT analysis of neural networks. This behavior could also be viewed as a stationary traveling wave, and as such serves as an example for the ability of the JVT to reveal joint spatiotemporal structure in stationary emergent behaviors. Transient traveling waves are by their very nature non-stationary and would require additional tools

form non-stationary signal processing such as (graph) wavelets, which is beyond the scope of this chapter.

8.4.1 CASE STUDY: SLOW SWITCHING ASSEMBLIES

Schaub et al. (2015) used a basic LIF model of a neuron along with a series of structured network models to produce coordinated firing whose behavior is interpretable with the network structure. The first model we will consider from Schaub et al. (2015) consists of densely connected “blocks” of excitatory neurons with only sparse connections between these densely connected blocks. Unlike the excitatory neurons, the inhibitory neurons are connected uniformly at random, as are the connections between excitatory and inhibitory (and vice versa). An example network is shown in Figure 8.2a, which was actually generated using the learning rules discussed in the following section. In this particular instance, there are 100 excitatory neurons split into six densely connected blocks and 25 inhibitory neurons.

As found in Schaub et al. (2015), when the excitatory neurons are stimulated by a critical level of external stimuli, these blocks exhibit bursts of activity within a densely connected block that slowly (and chaotically) transitions between blocks. An example of this slow switching assembly (SSA) behavior is shown in Figure 8.2b, illustrating that activity is concentrated within one block at a time, and exhibits slow, random switches between blocks. The inhibitory neurons are essentially randomly activated throughout the process, yet must play a key role in regulating the overall emergent behavior of the neural network. In particular, Schaub et al. (2015) note

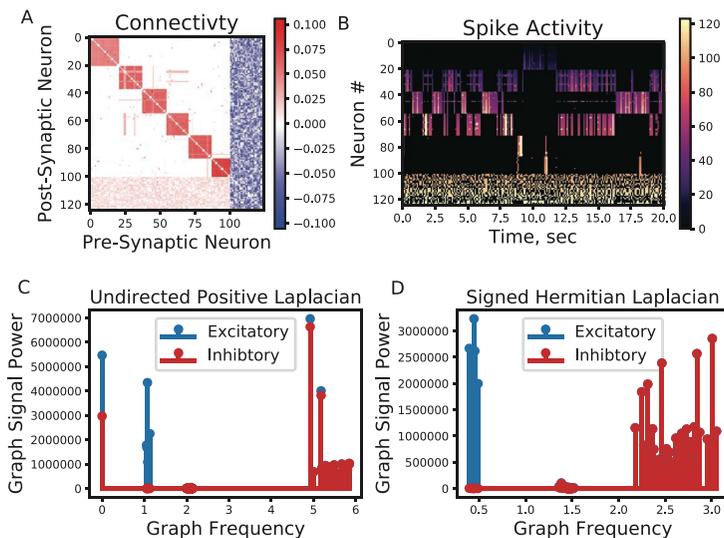


FIGURE 8.2 Slow-switching assembly simulations. (Adapted from Schaub et al. 2015.) (a) Adjacency matrix showing block structure. (b) Neuronal activity showing slow switching behavior between blocks. (c) GFT power using the undirected positive Laplacian. (d) GFT power using the signed Hermitian Laplacian.

that the inputs to the inhibitory networks must be slightly larger than the excitatory weights in order to maintain stability of the network.

In Schaub et al. (2015), analysis of this collective activity was considered in the context of the Schur decomposition of the network connectivity matrix. This linear algebraic technique could, in theory, be used to perform GSP analysis, since it produces a set of orthonormal basis vectors. However, the Schur decomposition is not unique, leading to a lack of reproducibility. As discussed above, since this network is both signed and directed, we need to consider GFTs that accommodate this additional structure. The first transform we consider uses the underlying undirected transform, and the graph signal power with respect to this transform is shown in Figure 8.2c. This transform captures most of the signal power of the excitatory network into a few harmonics, but there are more contributing harmonics than there are blocks in the network model. Additionally, the inhibitory and excitatory portions of the network are not totally separated, partially due to the constant harmonic, but also in the higher frequency harmonics.

In contrast to the underlying undirected transform, accounting for the signed and directed nature of the network model produces a graph Fourier power spectrum that clearly separates the excitatory and inhibitory portions of the network (see Figure 8.2d). Furthermore, the number of harmonics that capture the excitatory signal power is equal to the number of blocks in the excitatory portion of the network. The inhibitory portion of the network, on the other hand, is reasonably evenly split in power among 25 harmonics, consistent with the notion that they are seemingly firing at random. We will further explore the distribution of harmonics in this network in Section 5.1.

8.4.2 CASE STUDY: FRUIT FLY PROTOCEREBRAL BRIDGE

Another form of emergent collective firing in neural networks is self-sustained activity bumps, that is, localized regions of increased activity in a contiguous portion of the network. The output activity of the network may be guided by feed-forward inputs to the network that can manipulate the activity bump, but in the absence of such inputs the activity bump will be maintained. This sort of emergent behavior is prominent in neuronal circuits associated with navigation, where the inference of position and orientation must be maintained even in the absence of stimuli. Additionally, these circuits should be robust to noise such as random firings of both neurons within the network, and those modeled by the feedforward inputs. As such, seemingly chaotic coordinated firing as observed in the previous section represents a fundamentally different emergent behavior. To study this phenomenon through the lens of GSP, we use a simulation of the fruit fly protocerebral bridge (Kakaria and de Bivort 2017), which is believed to be responsible for an egocentric estimate of the fly's heading.

Unlike the SSA model from the previous example, the model of the fruit fly protocerebral bridge is structured in a fundamentally different way than the stochastic block model above (see Figure 8.3a). Here, the excitatory portion of the network results in ring-attractor dynamics with a corresponding triple of interlocking ring graphs (for full anatomic details see Kakaria and de Bivort (2017)). These ring attractor dynamics are exploited to maintain the heading estimate of the fruit fly. Additionally, unlike

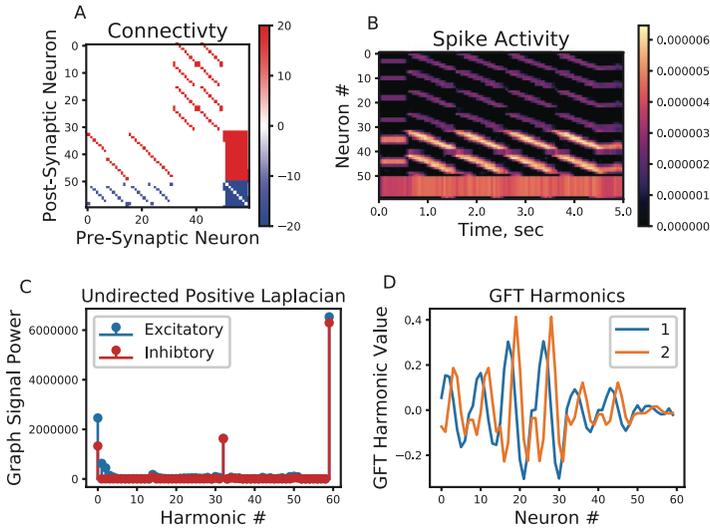


FIGURE 8.3 Fruit fly protocerebral bridge simulations. (Adapted from Kakaria and de Bivort, 2017.) (a) Adjacency matrix for the model. (b) Neural activity to a fixed angular velocity stimulus applied at 0.5 second for 4 seconds. (c) GFT power of the neural outputs. (d) GFT harmonics of the non-trivial active harmonics that are approximately sinusoidal and 90° out of phase on the excitatory neurons.

the SSA model, the weights of the synapses are chosen from similar, discrete values, and not the more random weights as above. The activity of the network is shown in Figure 8.3b. Unlike the SSA example above, this model expects a more structured feed-forward input that corresponds to the output of unmodeled upstream neurons. In the absence of this input signal, the network maintains a fixed activity bump corresponding to the last state estimate. Here, this corresponds to the times before 0.5 second and after 4.5 seconds. In the intervening times, a stimulus is applied that corresponds to a 1 Hz rotation, which causes the activity bumps on the excitatory neurons to rotate at that rate.

Using the undirected, positive Laplacian as the basis for a GFT we see considerable structure in the graph power spectrum (see Figure 8.3b). Harmonics 0, 32, and 59 capture the average contributions between the major functional components at each point in time. The next two contributing harmonics (1 and 2) are shown in Figure 8.3d. These harmonics resemble a pair of sinusoidal waves on the excitatory networks that are 90° out of phase. Such harmonics are characteristic of ring-like networks, and these harmonics contribute to the localized activity bumps that encode the heading of the fruit fly. In contrast, using the signed Hermitian Laplacian (not shown) appears to group the network into three regimes (0–31, 32–49, 50–59), and only has the ring-like harmonics on neurons (0–31).

Next, we further explore the interplay between GSP and emergence in neural networks by highlighting the ability of GSP to identify hidden patterns of coordinated activity in time and space. When the neurons are enumerated as in Figure 8.3b,

the coordinated firing activity is quite obvious, although one might question why the activity bump is repeated six times (or twice on each functional component). However, if the natural ordering were not known a priori, and instead we are presented with a random permutation of the neuronal indices, then the output is considerably harder to decipher (see Figure 8.4a). One might note that many of the neurons appear to have some periodic behavior, and standard Fourier analysis of the individual neurons indicates that many of the neurons do indeed have considerable power contributions at 1 Hz.

An important feature of the GFT is that it is “invariant” to permutations of the vertices, in the sense that individual harmonic vertex values will be permuted in the same way. Thus, the GFT power spectrum using the underlying undirected positive Laplacian is identical to that in Figure 8.3c. Armed with the knowledge of the structure from Fourier analysis in the time and vertex domains individually, we next considered the JVT transform (Figure 8.4c). This reveals that the 1 Hz power observed in the individual neurons in Figure 8.4b is strongly concentrated in just two harmonics, and these are of course harmonics 1 and 2. With the permutation of the neurons, these harmonics do not have the same readily apparent structure as in Figure 8.3d. Since these two harmonics have nearly identical graph frequency, one might be tempted to think of them in an analogous manner to the real and imaginary parts of a standard Fourier complex exponential harmonic. With this intuition, re-permuting the neuron index by the “phase” of the combined harmonic $U_1 + jU_2$ at

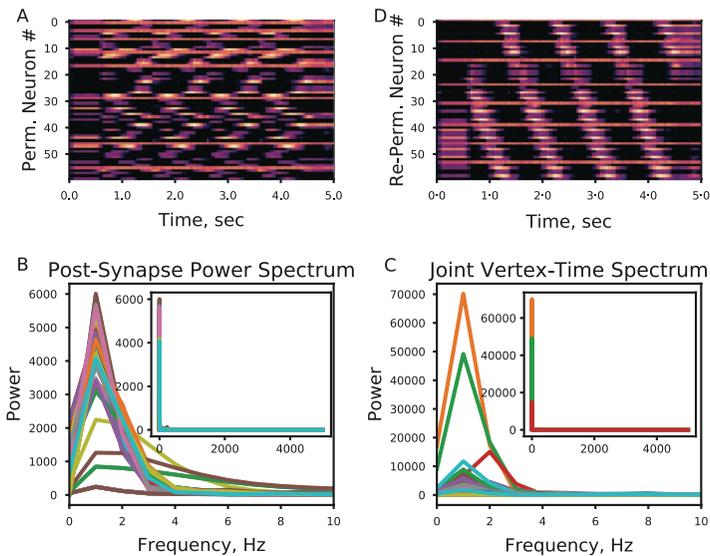


FIGURE 8.4 Permuting the index of the protocerebral bridge simulations. (a) Neuronal outputs when the neuron indices are permuted (i.e., the rows of Fig. 3b are permuted – the data is itself identical). (b) Power spectrum of neuronal outputs, illustrating concentration around 1 Hz. (c) JVT power spectrum, illustrating concentration of 1 Hz power in to two harmonics (note scale difference from b). (d) Re-permuting the neuron index using the phase between each neuron’s contribution to the dominant JVT harmonic, clearly revealing rotational input.

each neuron yields Figure 8.4d. This sorting completely unravels both the permutation and the original interleaved structure and shows how the activity bump travels smoothly across the excitatory portion of the network.

8.4.3 SUMMARY

In summary, neural networks exhibit many different forms of emergent collective activity that is coordinated spatially, temporally, or both. The synaptic network between the neurons is key in producing these collective behaviors, with the connectivity and weights ultimately responsible for generating the collective dynamics. Commonalities between the examples' network structure reveal a motif for potential consideration for wider SoS scenarios. This motif is that of structured, yet sparse, connectivity in the excitatory portion that demonstrates the “core” of some emergent behavior, coupled with dense connectivity involving a smaller inhibitory portion that serves to regulate the overall collective behavior. Additionally, we showed that GSP can be used to reveal low-dimensional hidden structure in time and space for these collective behaviors, but transforms that account for the directed nature and presence of inhibitory interactions may be required to fully reveal this structure.

8.5 EMERGENCE IN THE LONG TERM: NETWORK FORMATION AND LEARNING

The above discussion focused on emergent behaviors in networks of neurons where the network between the neurons is viewed as static in both the weights of the model as well as the connectivity. However, the structure of the network itself is fundamental in the overall function of the neural network. This is especially evident for self-sustaining activity bumps when viewed through the lens of attractor theory, as the network weights are a major component in the determination of the stable manifolds. As there is no centralized mechanism that governs the formation of the network, this formation process can itself be viewed as an emergent phenomena. The formation of the connections in the network and their strength (i.e., their weight in a model) is governed by a process known as plasticity, or more colloquially, learning. In machine learning applications, this learning process is often executed through a centralized gradient descent process that incrementally tunes the weights of the network from a pre-defined connectivity pattern or “architecture” in order to minimize some loss function.

In biological neural networks, the learning process is modeled using learning rules that are accomplished using only local information available to each neuron, individually. Perhaps the most common model of plasticity is Hebbian learning, which uses the reasoning that neurons that “fire together” should “wire together,” that is, coordinated firing should strengthen synaptic connectivity, and vice versa. The most common formulation of this approach is in terms of firing rate models, where the synaptic weights change via

$$\frac{dw_{ij}}{dt} = g(w_{ij}, v_i, v_j), \quad (8.8)$$

where g is a function of the current weight and the firing rates v_i of the neurons. The most basic form of Hebbian learning is $g(w_{ij}, v_i, v_j) = \eta_{ij} v_i v_j$, where η_{ij} is referred to as the learning rate.

While computationally tractable and intuitively appealing, this form of “pure” Hebbian learning has some shortcomings with respect to the network constraints observed in biological neural networks. First, it is clear that the learning rule above is symmetric, which will always result in undirected network models. Second, this approach to learning was originally intended for excitatory neurons only and does not account for inhibitory neurons. Finally, the dynamics of Hebbian learning often leads to instability, with exponential growth in the magnitudes of the weights. This latter failing can be addressed by various normalization techniques, leading to alternative learning rules such as Oja’s $g(w_{ij}, v_i, v_j) = \eta(v_i v_j - w_{ij} v_i^2)$, which asymptotically normalizes $\sum_j w_{ij}^2$ to 1.

The above learning rules are appropriate for firing rate models, with their rates η_i defined at all times. For spike-based models with their discrete spiking events, the learning rules should be dependent on the time between the pre-synaptic spike and the post-synaptic spike, so-called spike-timing-dependent plasticity (STDP). If we define t_k^i to be the time of the k th spike in the spike train S_i and t_l^j for the train S_j , a basic form of STDP is

$$\frac{dw_{ij}}{dt} = \sum_k \sum_l h(t_k^i - t_l^j), \quad (8.9)$$

where $h(t)$ is called the learning window. A basic form of h is

$$h(t) = \begin{cases} A_+ \exp\left(-\frac{t}{\tau_+}\right), & \text{for } t > 0, \\ -A_- \exp\left(\frac{t}{\tau_-}\right), & \text{for } t < 0, \end{cases} \quad (8.10)$$

with A_{\pm} , τ_{\pm} positive, and A_{\pm} may depend on w_{ij} . As with the rate-based Hebbian learning above, there are many variants of the STDP rules that capture different biologically observed phenomena. In particular, STDP can serve as both a model for directed connectivity as well as for a method of plasticity for inhibitory neurons.

8.5.1 CASE STUDY: STRUCTURED ASSEMBLY FORMATION

In Triplett, Avitan, and Goodhill (2018), a Hebbian-like learning rule was used to evolve the excitatory portion of a neural network from an initially uniformly randomly connected network to one that exhibits both the stochastic block structure and the slow switching behavior of Schaub et al. (2015). In a sense, this is essentially a generative mechanism for the dynamics of Schaub et al. (2015), although we note that the weights of any edges connected to the inhibitory network are fixed. There, for

ease of simulation, the LIF dynamics were simplified and time discretized, simulating the responses of the excitatory neurons S_i^E and inhibitory neurons S_i^I via

$$\begin{aligned} S_i^E(t+1) &= \Theta \left(\sum_j w_{ij}^{EE} S_j^E(t) - \sum_j w_{ij}^{EI} S_j^I(t) + \beta_i^E(t) - \gamma \right) \\ S_i^I(t+1) &= \Theta \left(\sum_j w_{ij}^{IE} S_j^E(t) - \sum_j w_{ij}^{II} S_j^I(t) + \beta_i^I(t) - \gamma \right), \end{aligned} \quad (8.11)$$

where w_{ij}^{EE} , w_{ij}^{EI} , w_{ij}^{IE} , and w_{ij}^{II} are the excitatory-to-excitatory, inhibitory-to-excitatory, excitatory-to-inhibitory, and inhibitory-to-inhibitory portions of the network, respectively, Θ is the Heaviside step function, γ is the activation threshold, and $\beta_i^E(t)$, $\beta_i^I(t)$ are random variables that drive spontaneous background activity. The model of Triplett, Avitan, and Goodhill (2018) uses a covariance learning rule, where the weights of the excitatory sub-network w_{ij}^{EE} are updated via

$$\Delta w_{ij}^{EE}(t) = \eta (S_i^E(t) - S_i^E(t)) (S_j^E(t) - S_j^E(t)), \quad (8.12)$$

where $\langle S_i^E(t) \rangle$ is the running average of $S_i^E(t)$ and η is the learning rate. Furthermore, the weight updates are constrained to prevent negative excitatory weights normalized to maintain a constant sum of weight for each excitatory neuron. This latter step is needed to prevent the ‘‘rich-get richer’’ phenomenon where all of the synaptic weight accumulates in a single edge.

As shown in Triplett, Avitan, and Goodhill (2018), the combination of these neuronal dynamics and learning rule results in the formation of strongly connected block models with slow-switching behavior in the vein of (Schaub et al. 2015). In fact, a simulation of this process was used to generate the network connectivity used in Section 8.1, that was then used as the connectivity for the higher-fidelity dynamics of (Schaub et al. 2015). Figure 8.5a shows how the number of strongly connected components evolves over time, increasing monotonically from the single initial component to six components. In the context of the GFT, the evolution of this network can be interpreted in terms of the associated graph frequencies, see Figure 8.5b. This figure shows that the formation of strongly connected blocks in the excitatory subnetwork are associated with the appearance of a low-frequency graph harmonic in the GFT. Of course, this result does not hold for the undirected transform, only the transform that accounts for the signed and directed nature of the graph.

8.6 CONCLUSION

In conclusion, biological neural networks exhibit many different forms of emergent coordinated firing activity that evolve both temporally and spatially across the network. Furthermore, the very process of network formation is itself an example of an emergent behavior on a much longer time scale. As a process that takes place on a natural network structure, tools from GSP can be applied to analyze these emergent

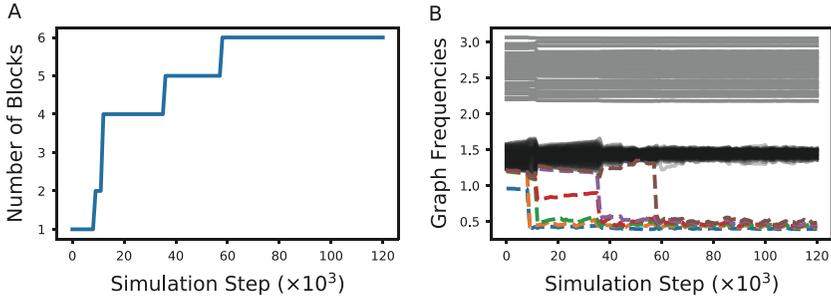


FIGURE 8.5 Learning induced block formation. (Adapted from Triplett, Avitan, and Goodhill, 2018.) (a) Number of blocks over time, eventually resulting in the network shown in Fig. 2a. (b) Evolution of graph frequencies over time, using signed Hermitian Laplacian. High-frequencies ($\lambda > 2$) cover the inhibitory neurons, whereas lower frequencies consist primarily of excitatory neurons. The dashed-colored lines are those that eventually map to a block in the model, which forms when the frequencies drop below the main group of frequencies.

behaviors, revealing insight into these processes. Given the rich history of interconnections between signal processing and control theory, this also points toward a future potential capability to engineer desired emergent behaviors. Looking beyond biological neural networks to the perspective of general SoS, biological neurons and neural networks have a number of interesting features to motivate the design of SoS, and GSP is a natural tool to consider these emergent SoS behaviors. Specific facets discussed in this chapter include:

- Collective emergent behaviors in neural networks assumes many forms beyond standard consensus and synchronization behaviors and may serve as inspiration for a number of other application areas. Many problems can presumably find neuromorphic or neuromimetic solutions once the appropriate neural circuit and conversion to problem domain is identified, see e.g., the swarming approach considered in Monaco et al. (2020).
- The structure of the neural network is key to the presence of emergence, and this structure must form organically, itself an example of emergence on a longer time scale than the shorter dynamical time scale.
- Neurons and the synapses that interconnect them have a number of regulatory processes that allow for interesting collective behaviors to emerge. These include
 - Excitatory and inhibitory neurons. In the examples studied here, what we would identify as the emergent behavior is observed primarily in the excitatory neurons, but the inhibitory neurons are needed to stabilize the behavior.
 - Auto-regulatory behaviors that vary the responsiveness of a neuron to inputs in order to prevent saturation.
 - Regulatory processes for the weakening and strengthening of synapses (network plasticity) to limit the rate of network change, maintain certain network features, and/or prevent degeneracies.

These processes and the motifs in the way they are employed could be valuable in understanding potential designs for emergence in other SoS, even those that are not explicitly neuro-inspired.

- GSP is a developing tool for analyzing dynamics and interactions on networks, but the peculiarities of biological neural networks (directed and signed) are an under-studied area.
- The proper choice of GFT can identify low-dimensional structure that is characteristic of emergent behavior and can also identify functional components in a neural network.
- The combination of graphical and standard Fourier analysis, leading to a JVT, can further reveal this low-dimensional behavior by detecting behaviors that evolve both temporally and spatially across the network.

ACKNOWLEDGMENTS

This work was supported by NSF NCS/FO 1835279 and JHUAPL IR&D funds. This material is based on work supported by (while serving at) the National Science Foundation. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- Amit, Daniel J., and Alessandro Treves. 1989. "Associative memory neural network with low temporal spiking rates." *Proceedings of the National Academy of Sciences of the United States of America* 86(20): 7871–7875.
- Aqil, Marco, Selen Atasoy, Morten L Kringelbach, and Rikkert Hindriks. 2021. "Graph neural fields: A framework for spatiotemporal dynamical models on the human connectome." *PLoS Computational Biology* 17(1): e1008310.
- Badimon, Ana, Hayley J Strasburger, Pinar Ayata, Xinhong Chen, Aditya Nair, Ako Ikegami, Philip Hwang, et al. 2020. "Negative feedback control of neuronal activity by microglia." *Nature* 586(7829): 417–423.
- Behjat, Hamid, Nora Leonardi, Leif Sörnmo, and Dimitri Van De Ville. 2015. "Anatomically-adapted graph wavelets for improved group-level fMRI activation mapping." *NeuroImage* 123: 185–199.
- Brahim, Abdelbasset, and Nicolas Farrugia. 2020. "Graph Fourier transform of fMRI temporal signals based on an averaged structural connectome for the classification of neuroimaging." *Artificial Intelligence in Medicine* 106: 101870.
- Buzsáki, Gyorgy. 2019. *The brain from inside out*. New York: Oxford University Press.
- Chung, Fan RK. 1997. *Spectral graph theory*. American Mathematical Soc. <http://dx.doi.org/10.1090/cbms/092>
- Dayan, Peter, and Laurence F Abbott. 2001. *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Computational Neuroscience Series. Cambridge: The MIT Press.
- de Lange, Siemon, Marcel de Reus, and Martijn Van Den Heuvel. 2014. "The Laplacian spectrum of neural networks." *Frontiers in Computational Neuroscience* 7: 189.
- Fuhrmann, Falko, Daniel Justus, Liudmila Sosulina, Hiroshi Kaneko, Tatjana Beutel, Detlef Friedrichs, Susanne Schoch, Martin Karl Schwarz, Martin Fuhrmann, and Stefan Remy. 2015. "Locomotion, theta oscillations, and the speed-correlated firing of

- hippocampal neurons are controlled by a medial septal glutamatergic circuit.” *Neuron* 86(5): 1253–1264.
- Furber, Steve. 2016. “Large-scale neuromorphic computing systems.” *Journal of Neural Engineering* 13(5): 051001.
- Furutani, Satoshi, Toshiaki Shibahara, Mitsuaki Akiyama, Kunio Hato, and Masaki Aida. 2019. “Graph Signal Processing for Directed Graphs Based on the Hermitian Laplacian.” In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 447–463.
- Gidon, Albert, Timothy Adam Zolnik, Pawel Fidzinski, Felix Bolduan, Athanasia Papoutsis, Panayiota Poirazi, Martin Holtkamp, Imre Vida, and Matthew Evan Larkum. 2020. “Dendritic action potentials and computation in human layer 2/3 cortical neurons.” *Science* 367(6473): 83–87.
- Goldsberry, Leah, Weiyu Huang, Nicholas F Wymbs, Scott T Grafton, Danielle S Bassett, and Alejandro Ribeiro. 2017. “Brain signal analytics from graph signal processing perspective.” In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 851–855. IEEE.
- Grassi, Francesco, Andreas Loukas, Nathanaël Perraudin, and Benjamin Ricaud. 2017. “A time-vertex signal processing framework: Scalable processing and meaningful representations for time-series on graphs.” *IEEE Transactions on Signal Processing* 66(3): 817–829.
- Hopfield, John J. 1982. “Neural networks and physical systems with emergent collective computational abilities.” *Proceedings of the National Academy of Sciences of the United States of America* 79(8): 2554–2558.
- Huang, Weiyu, Thomas AW Bolton, John D Medaglia, Danielle S Bassett, Alejandro Ribeiro, and Dimitri Van De Ville. 2018. “Graph signal processing of human brain imaging data.” In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 980–984. IEEE.
- Kakaria, Kyobi S, and Benjamin L de Bivort. 2017. “Ring attractor dynamics emerge from a spiking model of the entire protocerebral bridge.” *Frontiers in Behavioral Neuroscience* 11: 8.
- Knierim, James J, and Kechen Zhang. 2012. “Attractor dynamics of spatially correlated neural activity in the limbic system.” *Annual Review of Neuroscience* 35: 267–285.
- Leonardi, Nora, and Dimitri Van De Ville. 2013. “Tight wavelet frames on multislice graphs.” *IEEE Transactions on Signal Processing* 61(13): 3357–3367.
- Lindsay, Grace W. 2021. “Convolutional neural networks as a model of the visual system: Past, present, and future.” *Journal of Cognitive Neuroscience* 33(10): 2017–2031.
- McNaughton, Bruce L, Carol A Barnes, and JJEBR O’Keefe. 1983. “The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats.” *Experimental Brain Research* 52(1): 41–49.
- Medaglia, John D, Weiyu Huang, Elisabeth A Karuza, Apoorva Kelkar, Sharon L Thompson-Schill, Alejandro Ribeiro, and Danielle S Bassett. 2018. “Functional alignment with anatomical networks is associated with cognitive flexibility.” *Nature Human Behaviour* 2(2): 156–164.
- Miller, Paul. 2018. *An introductory course in computational neuroscience*. Cambridge: The MIT Press.
- Monaco, Joseph D, Grace M Hwang, Kevin M Schultz, and Kechen Zhang. 2020. “Cognitive swarming in complex environments with attractor dynamics and oscillatory computing.” *Biological Cybernetics* 114(2): 269–284.
- Monaco, Joseph D, Kanaka Rajan, and Grace M Hwang. 2021. “A brain basis of dynamical intelligence for AI and computational neuroscience.” *arXiv preprint arXiv:2105.07284*.
- Penttonen, Markku, and György Buzsáki. 2003. “Natural logarithmic relationship between brain oscillators.” *Thalamus & Related Systems* 2(2): 145–152.

- Petrovic, Miljan, Thomas AW Bolton, Maria Giulia Preti, Raphaël Liégeois, and Dimitri Van De Ville. 2019. "Guided graph spectral embedding: Application to the *C. elegans* connectome." *Network Neuroscience* 3(3): 807–826.
- Pfeiffer, Brad E, and David J Foster. 2013. "Hippocampal place-cell sequences depict future paths to remembered goals." *Nature* 497(7447): 74–79.
- Poirazi, Panayiota, and Athanasia Papoutsis. 2020. "Illuminating dendritic function with computational models." *Nature Reviews Neuroscience* 21(6): 303–321.
- Poulter, Steven, Tom Hartley, and Colin Lever. 2018. "The neurobiology of mammalian navigation." *Current Biology* 28(17): R1023–R1042.
- Puschel, Markus, and José MF Moura. 2008. "Algebraic signal processing theory: Foundation and 1-D time." *IEEE Transactions on Signal Processing* 56(8): 3572–3585.
- Raj, Ashish, Chang Cai, Xihe Xie, Eva Palacios, Julia Owen, Pratik Mukherjee, and Srikantan Nagarajan. 2020. "Spectral graph theory of brain oscillations." *Human Brain Mapping* 41(11): 2980–2998.
- Samsonovich, Alexei, and Bruce L McNaughton. 1997. "Path integration and cognitive mapping in a continuous attractor neural network model." *Journal of Neuroscience* 17(15): 5900–5920.
- Sandryhaila, Aliaksei, and José MF Moura. 2013. "Discrete signal processing on graphs." *IEEE Transactions on Signal Processing* 61(7): 1644–1656.
- Schaub, Michael T, Yazan N Billeh, Costas A Anastassiou, Christof Koch, and Mauricio Barahona. 2015. "Emergence of slow-switching assemblies in structured neuronal networks." *PLoS Computational Biology* 11(7): e1004196.
- Schmidhuber, Jürgen. 2015. "Deep learning in neural networks: An overview." *Neural Networks* 61: 85–117.
- Shuman, David, Sunil Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. 2013. "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains." *IEEE Signal Processing Magazine* 3(30): 83–98.
- Strata, Piergiorgio, and Robin Harvey. 1999. "Dale's principle." *Brain Research Bulletin* 50(5): 349–350.
- Tanner, Herbert G, Ali Jadbabaie, and George J Pappas. 2007. "Flocking in fixed and switching networks." *IEEE Transactions on Automatic Control* 52(5): 863–868.
- Trianni, Vito, Elio Tuci, Kevin M Passino, and James AR Marshall. 2011. "Swarm cognition: An interdisciplinary approach to the study of self-organising biological collectives." *Swarm Intelligence* 5(1): 3–18.
- Triplett, Marcus A, Lilach Avitan, and Geoffrey J Goodhill. 2018. "Emergence of spontaneous assembly activity in developing neural networks without afferent input." *PLoS Computational Biology* 14(9): e1006421.