MIT Libraries | DSpace@MIT

MIT Open Access Articles

Reading the Book of Memory: Sparse Sampling versus Dense Mapping of Connectomes

Massachusetts Institute of Technology

Perspective

Reading the Book of Memory: Sparse Sampling versus Dense Mapping of Connectomes

H. Sebastian Seung[1,*]
[1]Howard Hughes Medical Institute, Brain and Cognitive Sciences Department, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
*Correspondence: seung@mit.edu
DOI 10.1016/j.neuron.2009.03.020

Many theories of neural networks assume rules of connection between pairs of neurons that are based on their cell types or functional properties. It is finally becoming feasible to test such pairwise models of connectivity, due to emerging advances in neuroanatomical techniques. One method will be to measure the functional properties of connected pairs of neurons, sparsely sampling pairs from many specimens. Another method will be to find a "connectome," a dense map of all connections in a single specimen, and infer functional properties of neurons through computational analysis. For the latter method, the most exciting prospect would be to decode the memories that are hypothesized to be stored in connectomes.

In constructing a neural network model of brain function, it is standard to start from a mathematical description of spiking and synaptic transmission, make assumptions about how neurons are connected by synapses and then numerically simulate or analytically derive the activity patterns of the network. Success is declared if the model's activity patterns reproduce those measured by neurophysiologists.

Initially, the model neurons used in such networks were highly simplified to the point of being naive. But they have become more sophisticated over the years, incorporating findings about intrinsic and synaptic currents in neurons. In contrast, many assumptions about neural connectivity have been used by theorists for decades without revision, because they have been difficult to test empirically.

It has been popular to assume that the connectivity between any pair of neurons is a function of variables associated with the neurons. These variables, which I dub *cell labels*, are attributes of a neuron that can be measured without determining its connectivity directly. The cell label can include what neuroanatomists call *cell type*, which is defined classically by shape and location (Bota et al., 2003; Masland, 2004). In the retina, photoreceptors make connections onto horizontal cells, a rule of connectivity based on cell type (Masland, 2001b). A cell label could also include some property that is determined by a neurophysiologist through activity measurements. For example, some models of the primary visual cortex assume that excitatory neurons with similar preferred orientations are connected (Somers et al., 1995; Ben-Yishai et al., 1995), so that the cell label is preferred orientation.

For testing such a pairwise model of neural connectivity, two standard neuroanatomical methods are available. *Sparse reconstruction* relies on light microscopy and sparse labeling of neurons, and *dense reconstruction* relies on electron microscopy and dense labeling. Both methods have been problematic.

Axons can be less than 100 nm in diameter (Shepherd and Harris, 1998), and dendritic spine necks can be even narrower (Fiala and Harris, 1999). Since 100 nm is less than the wavelength of visible light, these structures cannot be resolved with a light

microscope if they are entangled in a densely stained neuropil (but see Hell [2007] for exceptions to this rule). However, one can see a single neuron stained with dye, as long as the surrounding neurons are unstained and hence remain invisible. This trick was employed by Golgi, who invented a stain that marked a sparse subset of neurons in the brain.

Cajal used Golgi's stain to reconstruct the branching patterns of neurons. If two neurons made contact with each other, Cajal inferred that they were connected. However, he could not rigorously prove this inference, because he could not see synapses. Contact suggests that a connection exists, but a synapse must be identified to prove it. In short, connection = contact + synapse.

In the 1970s, neuroanatomists began to use electron microscopy for dense reconstruction of neurons. In principle, this imaging method has enough spatial resolution to see *all* of the axons and dendrites in a densely labeled neuropil. It is also possible to identify synapses through telltale markers such as vesicles. Most famously, electron microscopy was used to map every connection in the nervous system of the nematode *C. elegans* (White et al., 1986). For every synapse between two neurites, the presynaptic and postsynaptic neurons were identified by tracing the neurites back to their parent cell bodies.

Although the *C. elegans* nervous system is quite small (see http://wormatlas.org for about 300 neurons and 7000 synapses), mapping its connections consumed over a decade of effort. White et al. (1986) called the fruits of their labors a "reconstructed nervous system." Others dubbed it a "wiring diagram," comparing the branches of neurons with the wires of an electronic device. Today we use the term *connectome* to refer to the complete map of all connections in a brain or piece of brain (Sporns et al., 2005; Lichtman and Sanes, 2008). Because of the Herculean labor involved, dense reconstruction has not been extended to more complex connectomes than that of *C. elegans*.

To diagnose the problems succinctly, sparse reconstruction has yielded contacts rather than connections, while dense reconstruction has been too laborious to be practical. Fortunately, these deficiencies are being rectified by emerging

technical advances. The advent of genetic methods of fluorescent labeling has improved the confidence with which sparse reconstruction can identify synaptically coupled pairs of neurons (Smith, 2007; Luo et al., 2008). And the automation of sectioning, electron microscopy, and image analysis is making the finding of connectomes more efficient (Briggman and Denk, 2006; Smith, 2007).

Improved methods for determining neural connectivity will aid neuroscience in many ways. In the study of neural development, it will become possible to precisely specify the goals of the processes that wire up the brain and characterize what happens when these processes malfunction. We will learn the exact nature of "connectopathies," pathologies of neural connectivity that are associated with mental disorders (Catani and ffytche, 2005). The study of interspecies differences in connectivity will shed light on brain evolution (Striedter, 2006).

These are all exciting possibilities, but I will take a more parochial view, considering only the implications for pairwise models of connectivity. Testing these models will either provide evidence for neural network models or refute them conclusively if the assumptions about connectivity are inconsistent with empirical data. The latter possibility is perhaps more important for scientific progress, if the definition of science rests on the Popperian criterion of falsifiability.

It is obvious how to use sparse reconstruction to test pairwise models. Simply determine the cell labels of connected pairs of neurons. Doing this repeatedly will eventually yield good estimates of pairwise statistics. If the cell label is just classical cell type, it can be determined by neuroanatomical means from the shape and location of the neuron. If the cell label also contains functional properties, finding it will require measurements of activity by electrophysiology or optical imaging.

In this paper, I'd like to advance a thesis that is less obvious: pairwise models will be testable by dense reconstruction, *even if the cell labels are unknown*. It may be possible to treat cell labels as hidden variables and infer their values by computational analysis of connectomes. In general, the hidden-variable approach will use dense information about connectivity to amplify the utility of sparse information about activity for understanding neural networks.

The retina will be an important arena for testing the new methods of sparse and dense reconstruction. Its neurons are divided into five broad classes, which are further subdivided into classical cell types, estimated to number 50–60 in the mammalian retina (Masland, 2001b). According to one hypothesis, the connectivity between a pair of retinal neurons is a function of their cell types and locations. The hypothesis appears to be true in the outer plexiform layer of the retina. Current research is focused on the inner plexiform layer, where the connectivity is still mostly unknown.

Success in the retina will be followed by more ambitious projects. An exciting prospect is testing the old theory that memories are stored in neural connections. If the theory is true, then a connectome might be compared to a book in which memories are written. To read a book, we must see the letters on the page and also decode their meaning. Dense reconstruction will enable us to "see" connectomes. Computational methods that infer cell labels by treating them as hidden variables will

allow us to "decode" connectomes. I will describe possible methods of reading procedural memories from avian brain area HVC (Li and Greenside, 2006; Jin et al., 2007) and spatial memories from the CA3 region of the rodent hippocampus (Tsodyks and Sejnowski, 1995; Samsonovich and McNaughton, 1997; Battaglia and Treves, 1998), based on existing pairwise models of connectivity.

Advances in Techniques for Measuring Connectivity

In recent years, genetic tools have revolutionized sparse reconstruction (see reviews in Smith, 2007; Luo et al., 2008). Genetically defined cell types can be labeled with fluorescent proteins, a degree of control not possible with the Golgi stain. Marking synapses with fluorescent proteins provides more evidence whether contact between two neurons should be considered a connection. Transsynaptic tracers spread from a neuron to connected neurons by crossing synapses. Wickersham et al. (2007) used a genetically modified rabies virus to label the presynaptic partners of a sparse set of neurons in vitro. More work will be required to verify the accuracy of this approach, as tracing errors could result if the virus fails to cross certain kinds of synapses or spreads between neurons that are not synaptically connected. The challenge of delivering the virus to specific cells or cell types in vivo also remains. But it's safe to say that sparse reconstruction has already improved dramatically in its ability to measure neural connectivity.

In dense reconstruction, both image acquisition and analysis have been laborious. Electron microscopy is coupled with the technique of serial sectioning to create 3D volume images of neural tissue (Harris et al., 2006). The specimen is manually cut into extremely thin sections, and electron microscopy is applied to produce a sequence of 2D images. This sequence is equivalent to a 3D volume image, up to the distortions that are produced by the process of sectioning and imaging. In expert hands, the best section thickness is about 50 nm, which limits the spatial resolution in the direction of cutting. Since axons can be less than 100 nm in diameter, they can become difficult or impossible to trace accurately when axons are nearly parallel to the cutting plane (White et al., 1986; Hobert and Hall, 1999).

By eliminating manual handling of sections, it is possible to cut thinner than the classical limit of 50 nm and yet reliably yield long sequences of 2D images. In serial block face scanning electron microscopy (SBF-SEM), the ultramicrotome is placed inside the vacuum chamber of an electron microscope (Denk and Horstmann, 2004). Thin sections are repeatedly cut off and discarded without imaging them. Instead, the exposed face of the block is repeatedly imaged with backscattered electrons. The entire process requires no manual intervention and yields a vertical resolution of better than 30 nm (Briggman and Denk, 2006). Alternatively, the automatic tape-collecting lathe ultramicrotome (ATLUM) is an instrument that cuts serial sections and collects them on a tape (Kasthuri et al., 2007). Once the fragile sections are on the sturdy tape, they can be handled for SEM imaging with relative ease.

But improved methods of image acquisition are not enough to make finding connectomes routine. White et al. (1986) reconstructed *C. elegans* neurons by manually tracing their cross sections through the images using a primitive technology, pen

marks on photographic prints. Image analysis is becoming faster and more accurate through the use of computers. In an ongoing project to find the connectome of the *C. elegans* male, computer software displays images to human users and captures and organizes their annotations (S. Emmons, personal communication). This computer-assisted approach has sped up image analysis by roughly ten times, making it probable that the connectomes of the male and larval stages will finally be completed in the near future. The approach has also enabled the reconstruction of highly branched neurons in the male, which were difficult or impossible by the original pen-based methods.

Scaling up the computer-assisted approach to larger connectomes would require large armies of human operators. Alternatively, it may be possible to reduce human effort by automating most of the analysis. This will depend on advances in computer vision, a branch of artificial intelligence. In general, it has been difficult to make computers perform visual tasks with accuracy approaching that of humans, and the task of tracing neurons in 3D electron microscopic images is no exception. One promising approach is based on machine learning. Examples of the tracing tasks are collected from human experts, and a computer is trained to emulate these examples (Jain et al., 2007; Andres et al., 2008). Historically, the machine learning approach has prevailed in speech and image recognition. Most likely, it will also produce superior accuracy in the neuron tracing problem.

The automation of image acquisition and analysis are crucial for speeding up the process of finding connectomes. At the present time, the spatial resolution and general quality of image acquisition is a prime concern. Once these problems are solved, the focus will turn to image analysis. There is still a long way to go before it becomes practical to find connectomes more complex than that of *C. elegans*, but there is reason to be optimistic. In time, technical progress will give rise to a new field called "connectomics," dedicated to the high-throughput generation of data about neural connectivity and the mining of that data for knowledge about the brain (Lichtman and Sanes, 2008).

Pairwise Models of Connectivity

New methods for sparse and dense reconstruction of neural connectivity will find many applications in neuroscience. One important application will be to formulate and test pairwise models of connectivity, which associate a variable l_i with each neuron i. The probability of a connection from neuron j to neuron i is

$$\mathbf{Pr}[C_{ij} = 1] = \Phi(l_i, l_j) \qquad (1)$$

where $C_{ij} = 1$ means that the $i \leftarrow j$ connection exists. In effect, the function Φ specifies a set of rules of connection. I will call the variable l_i the *cell label* of neuron i. Since the point of the model is to relate connectivity to other quantities, cell labels include attributes of neurons other than their connectivity.

The cell label could include the location of the cell body, the shape of the dendrite, and the projection targets of the axon, all part of the classical notion of *cell type* (Masland, 2004). Revisionists would prefer to define cell type based on other attributes, like gene expression, intrinsic firing properties, and

synaptic properties (Markram et al., 2004; Nelson et al., 2006), which could also enter in the cell label.

In the classic approach, it's assumed that cell types have functional significance, even though they are defined morphologically (Masland, 2004). Neurophysiologists are expected to find that each cell type possesses a distinct and well-defined function. If indeed cell types have functional significance, then a pairwise model *indirectly* relates the connectivity of a neural network to its function.

Neurophysiologists describe the encoding of information in neural activity using concepts like the preferred orientation of a neuron in primary visual cortex. Such attributes are functional properties of neurons and could also be included in the cell label. When cell labels contain functional properties, a pairwise model *directly* relates the connectivity of a neural network to its function.

Cell type is generally considered to be a discrete quantity. In other words, neurons are assumed to fall into distinct clusters in some parametric space, though formal tests of this assumption have been scarce (Badea and Nathans, 2004; Kong et al., 2005). There is no such restriction on cell labels: both discrete and continuous attributes are allowed.

According to Equation 1, the rules of connection are probabilistic rather than deterministic, allowing for "sloppiness" in the connectivity. Here I have written a Bernoulli model in which the number of $i \leftarrow j$ connections is either 0 or 1. An alternative is a Poisson model that allows multiple connections, in which case Φ would specify the mean of the Poisson distribution. The model can also be generalized to apply to analog strengths of synapses, in addition to binary connectivity.

C. elegans

In the *C. elegans* nervous system, connectivity is trivially a function of classical cell type. Since neurons have unique names and can be identified based on its shape and location, every neuron can be regarded as its own cell type. In a pairwise model of connectivity, the cell labels are simply the names of the neurons. Alternatively, a slightly more compact model results from taking the cell label to be one of the 118 classes into which White et al. (1986) divided the 302 neurons of the hermaphrodite worm (a left/right attribute would also be necessary).

The cell types or classes of *C. elegans* neurons are structurally defined, but it turns out that they also have functional significance. Perusal of the online database http://www.wormatlas. org/neuroimageinterface.htm shows that researchers have successfully identified functions for many of the worm's neurons.

The touch avoidance response is a good example. The worm responds to a gentle touch near the head or tail by swimming in the opposite direction. Motor neurons important for this behavior were identified by their connections to the dorsal and ventral body muscles used in swimming. Lesion studies using laser ablation of neurons or their precursors showed that some motoneurons are required for backward swimming, and others are required for forward swimming (Chalfie et al., 1985). Candidate touch receptor neurons were identified by a genetic screen and confirmed by laser ablation studies (Chalfie and Sulston, 1981). Interneurons were identified on the basis of their connections to both sensory and motor neurons. Laser ablation studies

showed that some were specialized for forward motion, whereas others were specialized for backward motion (Chalfie et al., 1985).

In these studies, the researchers used the connectome to formulate hypotheses about neuronal function, which were then tested using laser ablation. This strategy for relating structure to function was later applied to work out the neural basis of other behaviors, such as feeding (Avery and Horvitz, 1989) and navigation (Gray et al., 2005).

An Invertebrate Retina

In more complex nervous systems than *C. elegans*, the number of cell types is much less than the number of neurons. This is especially the case in neural systems with a repeated, columnar organization. For example, the compound eye of the horseshoe crab *Limulus polyphemus* is composed of about 1000 repeated units called ommatidia, each of which receives visual input from a particular direction (reviewed by Ratliff, 1965; Hartline and Ratliff, 1972). Most ommatidia contain a single eccentric cell that sends action potentials through the optic nerve to the brain. The axon collaterals of eccentric cells appear to form an interconnected plexus that mediates inhibitory interactions. Hartline, Ratliff, and their collaborators devised a pairwise model of these inhibitory connections, based on neurophysiological measurements. According to the model, the interactions depend on the separation between cells, but not on their absolute locations.

If we draw a coordinate system on the surface of the eye, then the location of the ith cell is specified by a two-dimensional vector $\mathbf{x}_i$. The average strength of the inhibition between cells i and j is modeled as

$$\mathbf{E}[W_{ij}] = w(\mathbf{x}_i - \mathbf{x}_j) \qquad (2)$$

where w is a scalar-valued function of a vector argument, W_{ij} is the strength of the inhibitory $i \leftarrow j$ coupling, and E denotes a statistical average. The function w is maximal for separations of a few hundred microns and vanishes for long distances (Barlow, 1969). As is typical of biological systems, real eyes only obey the idealized rule in a noisy or stochastic way. Real interaction strengths fluctuate around the average value set by the rule.

As mentioned earlier, cell type is generally considered to be a discrete quantity. It includes the name of the brain area to which the neuron belongs. However, the precise location of the neuron within a brain area is generally not included, because it varies continuously. By this convention, the locations in Equation 2 are cell labels, but they are not cell types. This is our first example of a rule of connection that depends on an attribute other than cell type.

The receptive fields of eccentric cells are composed of a central region surrounded by an annular region, and these two regions respond antagonistically to light. Hartline and Ratliff constructed a computational model in which the width of the surround is directly related to the spatial range of the inhibitory interaction in the pairwise model of Equation 2. This is still one of the most outstanding examples of relating the connectivity of a neural network to its function.

The Vertebrate Retina

When Hartline and Ratliff modeled the *Limulus* retina, their task was simplified by considering only a single cell type, the eccentric cell. The vertebrate retina is more complex and has been more difficult to understand. Its neurons are divided into five broad classes: photoreceptor, horizontal, bipolar, amacrine, and ganglion. A simple pairwise model of connectivity is based on labeling each cell with its location and its membership in one of the five classes. Cells connect to nearby cells, with a spatial dependence as in Equation 2. In addition, connections respect rules based on classes. Photoreceptors make chemical synapses onto horizontal cells, horizontal cells make electrical synapses onto each other, and so on. The rules of connection between photoreceptors and horizontal cells were crucial for understanding how the center-surround receptive fields of photoreceptors are generated by the outer plexiform layer (Naka and Rushton, 1967; Mahowald and Mead, 1988).

While this five-class model is useful, there is room for improvement. The five classes have been further subdivided into a larger number of cell types. Although everyone agrees that photoreceptors are divided into rods and cones, the proper division of other broad classes into specific cell types is still not settled. The task of cataloging retinal cell types is not trivial, given that it is not even clear how to rigorously define a cell type. One approach follows in the tradition of Cajal, which is to observe and categorize the shapes of large numbers of neurons using sparse reconstruction. Based on this approach, the number of cell types in the mammalian retina has been placed at about 60 (Masland, 2001a). A newer approach is to define cell types by gene expression and to fluorescently label them with genetic markers (Kim et al., 2008).

The outer retina, containing the photoreceptors and horizontal cells, has relatively few cell types and its connectivity is known. In contrast, the inner retina is estimated to contain over two dozen types of amacrine cell in mammals (Masland, 2001b), and little is known of their connectivity. This could be investigated through sparse reconstruction of connected pairs of neurons. The cell types of each pair would be determined from their shapes. Extensive sampling of connected pairs from many retinas would eventually lead to a pairwise model of connectivity. Such statistical sampling would require considerable labor, especially if some types of connections are rare. The yield of connected pairs per retina will be low, if cells are sparsely labeled at random, but could potentially be increased by using transsynaptic tracers.

An alternative is to perform dense reconstruction of a single retinal connectome. This would yield a complete set of reconstructed single neurons, which could be used to build a catalog of cell types. It would also provide a complete set of connected pairs. Therefore, exhaustive analysis of a single retina by dense reconstruction could potentially be more efficient than sampling from many retinas by sparse reconstruction. Whether this efficiency is realized will depend on emerging technological advances in the methods of connectomics.

Functional Properties as Cell Labels

The idea that rules of connectivity should be based on cell type dates back at least to Cajal (Masland, 2004; Bota and Swanson, 2007). Cell type and location may turn out to be sufficient for
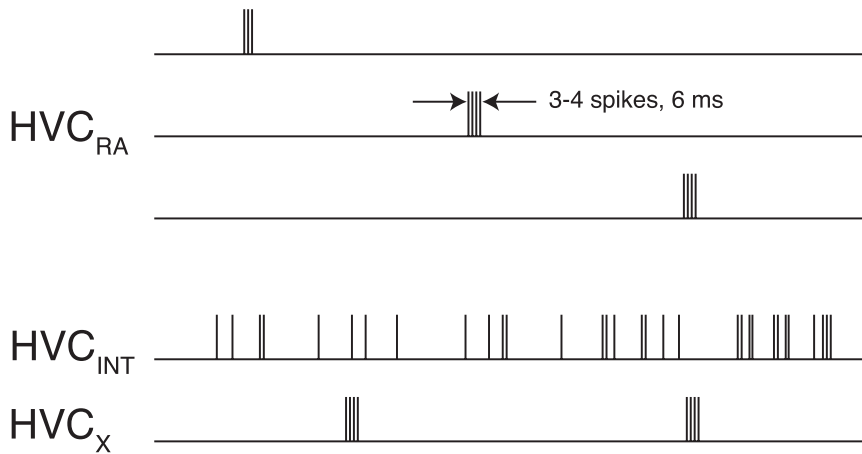
Figure 1. Cartoon of Spike Trains of HVC Neurons during a Song Motif
RA-projecting neurons generate a single burst of spikes. Interneurons fire at many times. X-projecting neurons generate a few bursts.

The zebra finch sings a single, highly stereotyped song that consists of repetitions of a motif, typically 0.5–1 s in duration. Lesion studies indicate that avian brain area HVC plays an important role in the production of birdsong (Nottebohm et al., 1976). Of the several cell types found in HVC (Dutar et al., 1998), the one that projects to nucleus RA is of prime importance for song production, because RA drives the motor neurons that control vocalization (Nottebohm et al., 1976).

understanding most of the structure and function of the retina. But in other neural circuits, and even in the retina to some extent, it will also be essential to include functional properties in cell labels.

Neurons of the same cell type may differ in their functional properties, and in their connectivity too. In the retina, it is likely that the connectivity of direction-selective ganglion cells will depend on their preferred directions. If so, the cell label should include the preferred direction, but this is impossible to predict from the shape of the ganglion cell, at least so far (Oyster et al., 1993). In the primary visual cortex, the connectivity of orientation-selective neurons may depend on their preferred orientations (Ben-Yishai et al., 1995; Somers et al., 1995). If so, the cell label should include the preferred orientation, but this seems impossible to predict from classical cell type.

Especially in sensory areas, neurons with similar functional properties tend to have similar locations, a phenomenon sometimes known as topographic organization. For example, in the retina and other visual areas of the brain, neighboring neurons are responsive to similar locations in visual space. Also, neighboring neurons have similar preferred orientations in the primary visual cortex of cats and monkeys (although see Girman et al., 1999; Ohki et al., 2005; concerning violation of this rule in rodents). In these cases, it could be argued that the location of a cell already subsumes its functional properties.

Areas involved in memory are better examples of the necessity of including functional properties in cell labels. As will be explained below, in avian brain area HVC (Fee et al., 2004) and rodent hippocampal CA3 (Redish et al., 2001), functional properties of neurons appear uncorrelated with location. This makes intuitive sense, as connections based only on cell type and location would be too inflexible to store memories. Furthermore, connections based on functional properties would result naturally if memories are stored in connections by activity-dependent synaptic plasticity, which is a popular idea among neuroscientists.

The Synaptic Chain Model of HVC
Procedural memory is defined as the ability to store and recall sequences of actions or thoughts, such as the long motor sequences of a musical performance. The neural basis of motor sequences has been studied in a songbird called the zebra finch.

The spiking of RA-projecting HVC neurons has been recorded during song in zebra finches (Figure 1). An RA-projecting neuron emits exactly one burst of several spikes during a song motif (Fee et al., 2004; Hahnloser et al., 2002). The timing of the burst is remarkably precise, with a jitter of less than a millisecond relative to the song. The activation times of the neurons are distributed throughout the motif. This means that the population of RA-projecting neurons generate a highly stereotyped, precisely timed sequence of bursts. This burst sequence drives activity in RA, which in turn drives the motor neurons that control vocalization (Fee et al., 2004; Hahnloser et al., 2002; Leonardo and Fee, 2005).
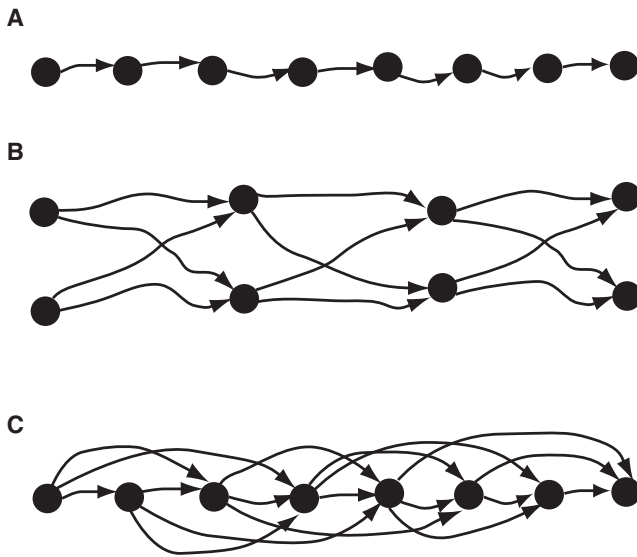
Projection neurons are known to send out axon collaterals within HVC (Katz and Gurney, 1981; Mooney, 2000), so they are thought to make excitatory connections with each other. It has been hypothesized that the connections between RA-projecting neurons are sequentially organized into a chain and that this connectivity is the cause of sequential activation (Li and Greenside, 2006; Jin et al., 2007). Some examples of chain-like networks are shown in Figure 2.

Suppose that the neurons of the chain are excitatory. If neurons at the left end are activated, the activity can propagate from left to right. The neurons will be activated sequentially, like a chain of falling dominoes. The synaptic chain has been proposed as an explanation for how humans are able to produce long sequences of movements, as in playing a piano or singing a song. It was first proposed by associationist philosophers over a century ago and later was mathematically formalized by neural network theorists (Amari, 1972; Abeles, 1991).

The synaptic chain can be formulated as a pairwise model by letting the activation time of a neuron be its cell label. There should be a high probability of connection only between pairs with a small and positive time difference. In other words, the probability of an $i \leftarrow j$ connection is expected to be

$$\mathbf{Pr}[C_{ij} = 1] = f(t_i - t_j) \qquad (3)$$

where f has a form like that shown in Figure 3 and C_{ij} is a binary random variable that indicates the presence of absence of an
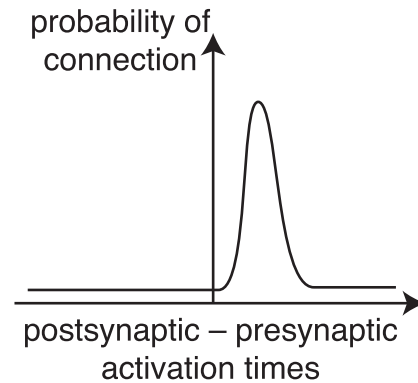
Figure 2. Examples of Synaptic Chains
(A) Each neuron makes a synapse onto a single neighbor on the right.
(B) Neurons are divided into groups, and each group makes synapses onto the next group.
(C) Each neuron makes a synapse onto three neighbors to the right (Amari, 1972; Abeles, 1991).

Figure 3. Connection Probability versus Difference of Activation Times in the Synaptic Chain Model of Equation 3

$t_i - t_j$ and check whether the histogram of the time differences resembles the function of Figure 3.
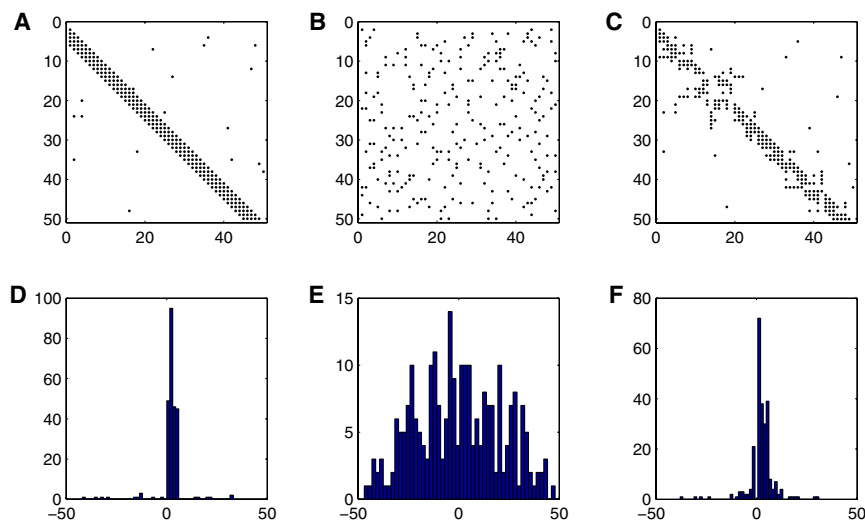
Less obvious is the fact that the synaptic chain model is testable by dense measurements of connectivity, even with no measurements of activity. If the activation times of the neurons are unknown, whether the network contains a chain is not obvious because the order of the neurons is scrambled. To find out whether there is a chain, we must try to unscramble the ordering. The activation times t_i are no longer directly observable but are hidden variables to be discovered through analysis.

Finding the activation times is an example of a problem known by computer scientists as *graph layout*, or *drawing* (Díaz et al., 2002), where "graph" is a mathematician's term for the connectivity of a network. The problem typically arises when a user would like to visualize a network as a low-dimensional picture. Most commonly, the nodes and connections are drawn as points and lines in two-dimensional space. An objective function is formulated that quantifies some notion of how "good" the graph looks. The locations of the nodes are chosen by optimizing the objective function. Many objective functions and optimization algorithms have been proposed for graph layout. Related algorithms have been applied to infer the spatial location of neurons in the *C. elegans* nervous system from their connectivity (Chen et al., 2006).

Finding the values of the t_i in Equation 3 is equivalent to laying out a directed graph in one dimension (Carmel et al., 2004). The problem is to arrange the RA-projecting neurons along a one-dimensional line so that these constraints are well-satisfied:

1. Connected neurons are close to each other.
2. Connections are pointed in the same direction.

If the neurons can be laid out in one dimension so that these constraints are well-satisfied, then it can be regarded as a synaptic chain. Furthermore, the resulting locations of the neurons could serve as predictions of their activation times. The above two constraints correspond to the assumption that the function f in Equation 3 is large only for small positive arguments. Note that the one-dimensional line, the layout space, is parametrized by activation time. This is an abstract space of functional properties, not a physical space of brain locations, since activation time is uncorrelated with location in HVC, as mentioned earlier (Fee et al., 2004).

$i \leftarrow j$ connection. If the chain were ideal, like those shown in Figure 2, the function f would vanish completely for a time difference that is either negative or large and positive. But f does not vanish in Figure 3, which is intended to simulate the "sloppiness" of biology. It is realistic to expect that a chain-like structure would be corrupted by some "wrong" connections that either go backward in time or jump too far forward in time.

One-Dimensional Directed Graph Layout
If the RA-projecting neurons were arranged along the length of HVC so that their activation times increase in an orderly fashion, then the cell label would be reducible to location. The synaptic chain model could be tested in vitro, because it predicts that connected pairs of RA-projecting neurons would be near each other in a brain slice preparation, and the direction of the connection would be the same as that of the increasing activation times. But in fact the activation times of RA-projecting neurons during song are apparently uncorrelated with their location in HVC (Fee et al., 2004).

The synaptic chain model could be tested by determining the cell labels directly through measurements of activation time during song. In the method of sparse reconstruction, one could use a transsynaptic tracer to mark the presynaptic partners of a single postsynaptic neuron and then measure the activation times of the marked neurons by optical imaging during song. The pairwise model predicts that the activation times of the presynaptic partners would typically be just a bit less than the activation time of the postsynaptic neuron. In the method of dense reconstruction, the activation time of each neuron could be measured by optical imaging prior to connectomic analysis. For every connected pair $i \leftarrow j$, compute the time difference

Figure 4. Finding Chain Structure in a Network

The connectivity matrices shown in the top panels are all of the same network, but with different orderings of the neurons.

(A) This ordering of the neurons makes the chain structure obvious in the connectivity matrix. Almost all the connections are $i \leftarrow j$ pairs in which $i - j$ is a small positive integer.

(B) With a random permutation of the neurons, the chain structure is hidden in the connectivity matrix.

(C) After sorting the neurons with the graph layout algorithm described in the Experimental Procedures, the chain structure is visible again.

(D–F) Histograms of $i - j$ for the connections corresponding to the three orderings of the neurons. If turned into probability distributions by normalization, they would represent the estimates of the probability of connection in Equation 3, given an ordering of the neurons. For the leftmost and rightmost orderings, the probability of connection is large only for small positive $i - j$. If no such ordering existed, there would be no chain structure in the network.

One-dimensional graph layout is an easy computational problem for an ideal chain like those shown in Figure 2, but difficult if the chain is corrupted by "inappropriate" connections. Many approaches are possible, and which approach eventually turns out to be best will depend on the exact nature of the HVC connectome data. Here I present a simple artificial example, just to illustrate how graph layout could be applied to connectomes. A network with chain structure was generated by using the probabilistic model of Equation 3. The exact procedure is described in the Experimental Procedures. For simplicity, it's assumed that the variables t_1 through t_n are a permutation of the integers 1 through n, where n is the number of neurons. Then any set of values for the t_i is equivalent to an ordering of the neurons.

The chain structure is obvious in the connectivity matrix of the network (Figure 4a) but becomes hidden if the neurons are randomly permuted (Figure 4b). A reordering of the neurons can be obtained by applying a graph layout algorithm based on a convex objective function described in the Experimental Procedures. This reordering is not as good as the original one because the algorithm is so naive, but it is good enough to make the chain structure evident in the connectivity matrix (Figure 4c).

The best estimate of the connection function f in Equation 3 is also shown in the figure. With the "good" orderings on the left and right, f is only large for small positive $i - j$. The existence of orderings with this property is evidence that the graph has chain structure.

The same analysis can be applied to a network with the same number of connections as the chain network, but with the connections placed completely randomly. In this case, the graph layout algorithm is unable to come up with a good ordering, i.e., one for which the estimated connection function is only large for small positive $i - j$ (data not shown).

The graph layout approach is possible only if a substantial fraction of the connections are known. Suppose some inappropriate connections are added to the ideal chains of Figure 2, to simulate the "sloppiness" of biology. Suppose also that some connections are randomly deleted, to simulate partial knowledge

of the connectome. If the fraction of deleted connections becomes large, it will become impossible to detect chain-like structure, due to the noise from the sloppy connections. Therefore, the HVC example demonstrates the utility of dense information about connectivity.

Strengths of Connection

Connectivity is generally defined as a binary property. Two neurons are either connected or they are not. But synaptic connections also have analog strengths, which brings us to a potential weakness of connectivity analysis. Consider the worst-case scenario of a network in which all pairs of neurons are connected, but most connections are so weak that they are nonfunctional. In this scenario, connectivity would provide almost no information about function. Let's reexamine the synaptic chain with this idea in mind.

I proposed a test of the synaptic chain hypothesis through an analysis of the connectivity between RA-projecting neurons in HVC. If graph layout algorithms fail to identify a chain structure, I asserted that the synaptic chain hypothesis will be falsified. In fact, that conclusion is a bit hasty. Imagine a chain of strong connections. Now add many weak connections in the "backward" direction. In this network there is a chain, but it cannot be identified by a pure connectivity analysis.

If HVC turns out to be like this, it will be important to distinguish between strong and weak connections. This could be possible in sparse and dense reconstruction. When a neuron makes multiple synaptic contacts onto another neuron, the strength of interaction is correlated with the number of synapses. Furthermore, synaptic strength as estimated by amplitude of postsynaptic calcium transients has been correlated with synapse size in vitro (Mackenzie et al., 1999; Matsuzaki et al., 2001). There is evidence that all components of synapses (spine, bouton, active zone, postsynaptic density) are correlated in size and are also correlated with properties like total vesicle number and docked vesicle number, which are thought to be related to release probability (Harris and Stevens, 1989; Pierce and Lewin, 1994; Schikorski and Stevens, 1997; Murthy et al., 2001). In

short, it should be possible to estimate synaptic strength from the same light or electron microscopic images used to determine connectivity.

Is Inhibition Less Specific than Excitation?

So far our discussion of HVC has focused on the RA-projecting neurons and how the excitatory synapses between them could be organized like a chain. But HVC also contains other neuron types (Dutar et al., 1998). There are X-projecting neurons, which send axons to area X in the anterior forebrain pathway, as well as inhibitory interneurons, which do not project out of HVC. Here we'll discuss the possible roles of inhibition.

Historically, many theories of neural networks have assumed that inhibitory connections are less selective than excitatory connections (see Orban [1984] for review). Inhibitory synapses are hypothesized to play a supporting role in a variety of ways Worgotter and Koch (1991), such as preventing runaway activity by holding excitation in check (Douglas et al., 1995), or sharpening a feature selectivity that is created by excitation (Somers et al., 1995; Ben-Yishai et al., 1995). Along the same lines, many theories of learning in neural networks focus on plasticity at excitatory synapses. The assumption is that inhibitory synapses are less specific, because they are not adjusted by plasticity.

In particular, network theories of sequence generation have taken different views of the interplay between excitation and inhibition. The excitatory chain model actually fell out of favor long ago in psychology, after the critique of Lashley (1951). Rumelhart and Norman (1982), following a suggestion by Estes (1972), modeled human sequence generation using a network with precisely organized inhibitory connections and a "fatigue" mechanism that limits how long a neuron can be active. Sequences were generated by the orderly release of neurons from inhibition. A related idea was used by Drew and Abbott (2003) in their disinhibitory chain model of HVC.

The validity of such models for human sequence generation remains untested by neuroscientific methods. But in zebra finch HVC, the neurophysiological evidence is against specificity of inhibition. As shown in Figure 1, the spiking of interneurons during song is very different from that of projection neurons. Interneurons are temporally unselective in their song-related activity, firing at many times during a song motif (Hahnloser et al., 2002). Therefore it seems implausible that inhibition plays an important role in determining the sequence in which RA-projecting neurons are activated.

It's known that interneurons receive synaptic input from RA-projecting neurons, based on paired recordings in vitro (Mooney and Prather, 2005). The lack of temporal selectivity could be explained if each interneuron receives synaptic input from a large number of randomly chosen RA-projecting neurons. This excitatory synaptic drive would cause each interneuron to fire throughout the song motif. While the connections from RA-projecting neurons to interneurons are likely to be nonspecific, the nature of the connections in the other direction is less clear. According to the paired recordings of Mooney and Prather (2005) random pairs of RA-projecting neurons are coupled with high probability by disynaptic inhibition, suggesting that the connections from interneurons to projection neurons are also

nonspecific. It seems likely that inhibition plays a supporting rather than starring role in sequence generation. Perhaps inhibition might serve to help prevent instability to runaway excitation (Jin et al., 2007).

Such global or unstructured inhibition is assumed by another model of sequence generation in which the memory of a sequence is stored in the afferent inputs to the network (Burgess and Hitch, 1999). The inhibitory connections lead to winner-take-all behavior, so that only one neuron or group of neurons can be active at a time. The order of activation in the sequence is not defined by local connections, but rather by a gradient in the strength of afferent input from outside the network. While active, each neuron or group of neurons suppresses the others. Eventually it fatigues, and the next neuron or group of neurons takes over. Although this model is not so robust for long sequences, it could become more plausible if analysis of the HVC connectome fails to find sequential structure in the local excitatory connections.

There are other network models of sequence generation; space does not permit an exhaustive discussion. The main point is that all models make assumptions about the organization of excitation and inhibition, and these assumptions could be tested by connectivity analysis.

Cognitive Maps and Hippocampal Connectivity

While birdsong may seem like a very specialized behavior, it can be viewed more generally as an example of procedural memory. The memory is laid down while a juvenile bird learns to imitate the song of its adult tutor. After learning is complete, each bout of song is recall of the memory, and sequential activity in HVC is the neural basis of memory recall.

From this viewpoint, the synaptic chain model of HVC is a specific example of a classic theory in neuroscience. It has long been hypothesized that memories are stored as synaptic connections, but experimental tests have been difficult. If the HVC connectome is obtained by dense reconstruction, and a chain structure is discovered through hidden variable inference, this would amount to a kind of "mind-reading." Predicting the sequential order of activation of HVC neurons from connectivity would be impressive, because it would be an example of reading out a memory stored in the connections of a network.

Could this approach be extended to other kinds of memory? The hippocampus has been of great interest to memory researchers. In humans it has been implicated in declarative memory (Eichenbaum, 2004). In rodents, the hippocampus appears to be dedicated to spatial memory (O'Keefe and Nadel, 1978). Neurophysiologists have recorded neural activity in the hippocampus in freely moving rats and found neurons called "place cells." For each of these neurons, there is a particular location in the environment at which spikes are generated. This location is called a "place field." The place fields of different neurons are distributed over many locations, so it has been hypothesized that the place cells constitute a cognitive map of the environment. If the rat is put into a novel environment, then the neurons again have place fields, but the locations of the new place fields appear to have no relationship with those in the previous environment. This has been taken as evidence that multiple cognitive maps are stored in the hippocampus.

How could these memories be stored? One possibility is the connections between excitatory neurons in the CA3 region of the hippocampus. In one proposal, neurons with nearby place fields in some cognitive map tend to be connected with each other (Tsodyks and Sejnowski, 1995). This is justified by invoking Hebbian plasticity and the fact that neurons with overlapping place fields tend to be coactive. In the formal model, the strength W_{ij} of the $i \leftarrow j$ synapse between excitatory neurons is given by (Samsonovich and McNaughton, 1997; Battaglia and Treves, 1998)

$$W_{ij} \propto C_{ij} \sum_{\mu = 1}^{P} K(|\mathbf{x}_i^\mu - \mathbf{x}_j^\mu|). \qquad (4)$$

The argument of the function K is the distance between the two-dimensional vectors $\mathbf{x}_i^\mu$ and $\mathbf{x}_j^\mu$, which are the locations of the place fields of neurons i and j in cognitive map μ. The function K decreases with distance, and the sum over μ includes P cognitive maps. The cell label of neuron i consists of the set of locations $\{\mathbf{x}_i^\mu\}_{\mu = 1}^{P}$ of its place fields in the cognitive maps. These place field locations appear to be uncorrelated with the physical location of the neuron within the hippocampus (Redish et al., 2001).

The elements of the connectivity matrix C_{ij} are either 0 or 1 and are generated by tossing a biased coin. Here the connectivity of the network is assumed to be completely random and static, an assumption about the hippocampus that dates back at least to Marr (1971). Only the strengths of synapses change during learning, and they store the cognitive maps. While this model is intriguing, it has been difficult to test experimentally. In the future, direct tests could become possible through sparse or dense reconstruction.

In the sparse reconstruction approach, one could mark connected neurons with transsynaptic tracers, measure the locations of their place fields, and compare with the model of Equation 4. This would be more challenging than the HVC experiments, because the hippocampus is a deep brain structure, making intracellular injection of tracers and optical imaging of activity more difficult.

The dense reconstruction approach would require not only the connectivity of CA3, but also the strengths of synapses as described previously. Computational analysis would be simplified by raising the rat in a limited number of environments, or even just a single one. In the case of a single cognitive map ($p = 1$), the model of Equation 4 would be tested by attempting to lay out the neurons in two dimensional space so that nearby neurons are more strongly connected to each other. Note that "nearby" refers to distance in an abstract functional space, rather than physical space, as locations of place fields in an environment seem uncorrelated with locations in the hippocampus (Redish et al., 2001). Solving this graph layout problem would yield predictions for the locations of the place fields of the neurons. This is almost the same as the graph layout problem for HVC, except for two differences. Here the layout is in two spatial dimensions instead of one temporal dimension. Also, the layout is not directional. The case of a single cognitive map could be relevant for a rat that has been kept in the same environment for an extremely long time.

If there are many ($p > 1$) cognitive maps stored in CA3 (Samsonovich and McNaughton, 1997; Battaglia and Treves, 1998),

then analyzing the connectome would be more complex. The problem could be formulated as the minimization of the cost function

$$f(\tau) = \begin{cases} 1, & \tau = 1 \text{ to } 5, \\ 0, & \tau = 0, \\ 0.01, & \text{otherwise} \end{cases} \qquad (5)$$

with respect to the place field locations x_i^μ and the scale parameter a. Finding a good way of solving this optimization problem would be an excellent topic of future research.

Model Selection

In my definition of pairwise models, I wrote that the cell label should contain attributes of a neuron that can be measured independently from its connectivity. Later on I explained that one can also infer cell labels instead of measure them directly. Taking one more step in this direction, one can even build pairwise models in which direct measurement of cell labels is impossible. The cell labels have no independent existence apart from the model; they are only invoked to help explain the observed connectivity.

In this approach, the cell label of a neuron would effectively be defined by its connectivity, emerging from computational analysis of a connectome. This idea would not be foreign to the classical neuroanatomists, who viewed morphological cell type as a proxy for connectivity (Masland, 2004). The shape of a neuron constrains its possibilities for connection, since the dendritic and axonal arbors of a neuron define a spatial region that other neurons must enter, if they are to be connected. One can go as far to say that morphological cell types have functional significance precisely because they are indirect measures of connectivity.

When cell labels are not grounded on independent measurements, there is a danger of "overfitting" a pairwise model to connectivity data. Any connectome could be modeled by giving each neuron its own unique cell label, and then encoding the entire connectome in the function Φ. The resulting model would trivially fit the data perfectly, but would not be any more compact than the connectome itself. This violates Occam's Razor, the principle that a theory or model should be a compact description of the data. According to modern statistical learning theory, compactness of a model is not an end in itself, but rather a means to accurate predictions (Vapnik, 2000). Overfitting is avoided by selecting a model with good validation performance on a test set. To apply this cross-validation method, divide the pairs of neurons in a connectome into distinct subsets for training and testing. After fitting a pairwise model to the data in the training set, use it to predict the data in the test set. The best model is the one with the best prediction performance on the test set.

Discussion

New neuroanatomical methods offer the promise of finally testing pairwise models of connectivity. The sparse reconstruction approach repeatedly identifies connected pairs of neurons and then measures their cell labels. Dense reconstruction reveals the entire connectome, from which the values of cell labels can be inferred if direct measurements are not available.

Which approach is better? Both approaches are viable; ultimately the question boils down to efficiency. In the short run, I expect that more scientific findings will come from sparse reconstruction. Dense reconstruction is still very laborious, as shown by the example of *C. elegans*. However, the efficiency of dense reconstruction will improve over the coming years. In genomics, the cost of DNA sequencing has dropped exponentially over the past three decades, at a rate similar to Moore's Law of the semiconductor industry (Shendure et al., 2004). If neuroscience is fortunate, there will also be exponential progress in connectomics.

Although sparse reconstruction will also improve in efficiency, I see less upside potential. Because the approach obtains a small amount of information from each specimen, it depends on processing a large number of specimens. Large sample sizes may be required to find the connectivity of rare cell types. Sampling will be even more problematic if staining methods undersample or miss some cell types altogether. For these reasons, sparse reconstruction will always be a laborious endeavor.

Using the HVC and CA3 examples, I argued that connectomes will be useful even if they are accompanied with no information about activity from the same sample. This extreme case was considered for the sake of argument, but a connectome could also be accompanied by sparse information about activity. This extra information would only enhance the ability of the hidden variable approach to "fill in the blanks," i.e., infer the rest of the information in cell labels. In this way, dense information about connectivity could be used to amplify the utility of sparse information about activity.

It should be kept in mind that pairwise models of connectivity are the most amenable to testing by the sparse reconstruction approach. Only one connection per connectome need be sampled, which is very sparse indeed. But one can imagine higher-order models in which connections are no longer independent when conditioned on cell labels. To test such models by sparse reconstruction, larger groups of connected neurons would have to be sampled, which would become more difficult. For testing sufficiently complex models of connectivity, dense reconstruction would become the only practical method.

In restricting the discussion to specific pairwise models of connectivity, I have described a more hypothesis-driven style of research. But dense reconstruction would also allow a more exploratory style, with fewer preconceptions constraining the connectivity models that are investigated. Sparse reconstruction has an important advantage: the activity of neurons can be measured after finding their connectivity, rather than before. In some cases, knowing connectivity in advance could be important for designing the right experiment involving activity.

Experimental Procedures

In Figure 4, a network with 50 neurons was generated using Equation 3. The values of the cell labels were set at $t_i = i$, running from 1 to 50. The connection function was set to

$$f(\tau) = \begin{cases} 1, & \tau = 1 \text{ to } 5, \\ 0, & \tau = 0, \\ 0.01, & \text{otherwise.} \end{cases}$$

In words, there were rightward connections for pairs of neurons up to a distance five apart on the chain. Other random connections were added with 1% probability to model the kind of "sloppiness" that might exist in a biological network. There were no self-connections.

Graph layout algorithms are often based on convex optimizations (Carmel et al., 2004). These simple algorithms require little computation time, but may produce inferior results. They can be used as initial conditions for more sophisticated algorithms, if better results are needed. To produce the ordering used to display the connectivity matrix in Figure 4C, the cost function

$$Q = \frac{1}{2} \sum_{ij} C_{ij} (t_i - t_j - 1)^2 \qquad (6)$$

was minimized (Carmel et al., 2004). Assuming that $C_{ii} = 0$, the minimum of this quadratic function can be found by setting its derivatives

$$\frac{\partial Q}{\partial t_i} = \sum_j C_{ij}(t_i - t_j - 1) - \sum_j C_{ij}(t_j - t_i - 1)$$

to zero. The derivatives can be rewritten as

$$\frac{\partial Q}{\partial t_i} = \sum_j L_{ij} t_j - \sum_j (C_{ij} - C_{ji})$$

where

$$L_{ij} = \delta_{ij} \sum_k (C_{ik} + C_{ki}) - C_{ij} - C_{ji}$$

is the symmetrized graph Laplacian. So the minimum of Equation 6 satisfies the linear equations

$$L\mathbf{t} = \mathbf{d}^{in} - \mathbf{d}^{out}$$

where d_i^{in} and d_i^{out} are the indegree and outdegree of neuron i. Solving these equations produces a set of real values for t_i, up to a degeneracy corresponding to translating all times by the same amount. These real values were sorted to produce the ordering of the neurons used in Figure 4C.

This algorithm works reasonably well when there are fewer "sloppy" connections that corrupt the chain organization. However, it fails when these connections become more frequent (results not shown). The graph layout problem can be formulated as maximum likelihood inference of the cell labels in the probabilistic model (1). The connections are assumed to be independent when conditioned on the variables l_i, so the log-likelihood of a connectivity matrix C_{ij} is simply the sum of the log-likelihoods for the connections,

$$\log P(C|\{l_i\}) = \sum_{i,j} (C_{ij} \log \Phi(l_i, l_j) + (1 - C_{ij}) \log[1 - \Phi(l_i, l_j)]). \quad (7)$$

Optimizing this objective function should give superior results, but will be technically difficult because it is typically nonconvex. More research on the graph layout problem is clearly in order.
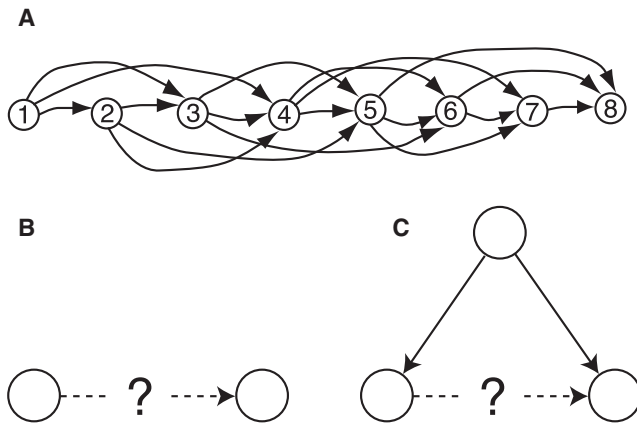
Figure 5. Hidden Variables Can Induce Statistical Dependencies between Connections Generated by the Pairwise Model of Equation 1

(A) The neurons are arranged in a chain and labeled from 1 to 8, indicating their locations. There is a rightward connection between two neurons if their labels differ by three or less.
(B) If the labels of the neurons are unknown, the connection probability is 18/56.
(C) If the neurons are known to both receive connections from a third neuron, then the connection probability rises to 1/2.

Application to real connectomes will require better algorithms that use nonconvex cost functions like Equation 7, rather than the convex cost function of Equation 6.

Hidden Variables and Dependencies between Connections

A pairwise model of connectivity contains a strong assumption called *conditional independence*. Suppose that your task is to predict whether a pair of neurons is connected, and you know nothing other than their cell labels. By a pairwise model, your prediction will not be helped by knowledge of the rest of the connections in the network. This might seem inconsistent with studies of *C. elegans* and visual cortex, which have found statistical dependencies between connections (Reigl et al., 2004; Song et al., 2005; Yoshimura et al., 2005). However, it turns out that the pairwise model can generate dependencies between connections, if some information in the labels is unknown. If the cell labels are allowed to be random variables, then they induce statistical dependencies between connections. In general, it is well-known in probabilistic modeling that hidden or latent variables can induce dependencies between observable variables.

For a specific example, consider the chain-like connectivity shown in Figure 5A. The connectivity can be modeled using Equation 1, if a label t_i is assigned to each neuron to represent its location on the chain. The labels take on the values 1 through 8 for the particular chain shown in Figure 5A. Neurons are connected from left to right if their separation is no greater than three. This is modeled by making the connection probability of Equation 3 depend only on the difference between the chain locations,

$$f(\tau) = \begin{cases} 1, & \tau = 1, 2, 3 \\ 0, & \text{otherwise.} \end{cases}$$

If the labels of all neurons are known, then all connections are statistically independent. Knowing whether one connection exists is irrelevant for knowing whether another one exists. On the other hand, suppose that the labels are unknown. Given a pair of neurons, there is no way to know where they are located in the chain. Then the connection probability for an arbitrary pair of neuron is 18/56, which is the ratio of the number of connections to the number of ordered pairs. But if both neurons in a pair are known to receive connections from a third node, then the separation between the pair must be no more than two. This implies that there must be an connection in one direction or the other. So the connection probability for a given direction rises to 1/2.

Here the connection probability changed when information about other connections was provided. This dependency arose because information about the labels was missing. Information about the connections from the third neuron changed the connection probability by reducing ignorance about the labels.

REFERENCES

Abeles, M. (1991). Corticonics: Neural Circuits of the Cerebral Cortex (Cambridge: Cambridge University Press).

Amari, S. (1972). Learning patterns and pattern sequences by self-organizing nets of threshold elements. IEEE Trans. Comput. *21*, 1197–1206.

Andres, B., Köthe, U., Helmstaedter, M., Denk, W., and Hamprecht, F.A. (2008). Segmentation of SBFSEM volume data of neural tissue by hierarchical classification. In Pattern Recognition: 30th DAGM Symposium, G. Rigoll, ed. (New York: Springer), pp. 142–152.

Avery, L., and Horvitz, H.R. (1989). Pharyngeal pumping continues after laser killing of the pharyngeal nervous system of C. elegans. Neuron *3*, 473–485.

Badea, T.C., and Nathans, J. (2004). Quantitative analysis of neuronal morphologies in the mouse retina visualized by using a genetically directed reporter. J. Comp. Neurol *480*, 331–351.

Barlow, R. (1969). Inhibitory fields in the *Limulus* lateral eye. J. Gen. Physiol. *54*, 383–396.

Battaglia, F.P., and Treves, A. (1998). Attractor neural networks storing multiple space representations: A model for hippocampal place fields. Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics *58*, 7738–7753.

Ben-Yishai, R., Bar-Or, R.L., and Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. Proc. Natl. Acad. Sci. USA *92*, 3844–3848.

Bota, M., and Swanson, L.W. (2007). The neuron classification problem. Brain Res. Rev. *56*, 79–88.

Bota, M., Dong, H.W., and Swanson, L.W. (2003). From gene networks to neural networks. Nat. Neurosci. *6*, 795–799.

Briggman, K.L., and Denk, W. (2006). Towards neural circuit reconstruction with volume electron microscopy techniques. Curr. Opin. Neurobiol. *16*, 562–570.

Burgess, N., and Hitch, G.J. (1999). Memory for serial order: A network model of the phonological loop and its timing. Psychol. Rev. *106*, 551–581.

Carmel, L., Harel, D., and Koren, Y. (2004). Combining hierarchy and energy for drawing directed graphs. IEEE Trans. Vis. Comput. Graph. *10*, 46–57.

Catani, M., and ffytche, D.H. (2005). The rises and falls of disconnection syndromes. Brain *128*, 2224–2239.

Chalfie, M., and Sulston, J. (1981). Developmental genetics of the mechanosensory neurons of *Caenorhabditis elegans*. Dev. Biol. *82*, 358–370.

Chalfie, M., Sulston, J.E., White, J.G., Southgate, E., Thomson, J.N., and Brenner, S. (1985). The neural circuit for touch sensitivity in Caenorhabditis elegans. J. Neurosci. *5*, 956–964.

Chen, B.L., Hall, D.H., and Chklovskii, D.B. (2006). Wiring optimization can relate neuronal structure and function. Proc. Natl. Acad. Sci. USA *103*, 4723–4728.

Denk, W., and Horstmann, H. (2004). Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. PLoS Biol. *2*, e329.

Díaz, J., Petit, J., and Serna, M. (2002). A survey of graph layout problems. ACM Comput. Surv. *34*, 313–356.

Douglas, R.J., Koch, C., Mahowald, M., Martin, K.A., and Suarez, H.H. (1995). Recurrent excitation in neocortical circuits. Science *269*, 981–985.

Drew, P.J., and Abbott, L.F. (2003). Model of song selectivity and sequence generation in area HVc of the songbird. J. Neurophysiol. *89*, 2697–2706.

Dutar, P., Vu, H.M., and Perkel, D.J. (1998). Multiple cell types distinguished by physiological, pharmacological, and anatomic properties in nucleus HVc of the adult zebra finch. J. Neurophysiol. *80*, 1828–1838.

Eichenbaum, H. (2004). Hippocampus cognitive processes and neural representations that underlie declarative memory. Neuron *44*, 109–120.

Estes, W.K. (1972). An Associative Basis for Coding and Organization in Memory. In Coding Processes in Human Memory, Chapter 7, A.W. Melton and E. Martin, eds. (Washington, DC: Winston), pp. 161–190.

Fee, M.S., Kozhevnikov, A.A., and Hahnloser, R.H.R. (2004). Neural mechanisms of vocal sequence generation in the songbird. Ann. N Y Acad. Sci. *1016*, 153–170.

Fiala, J.C., and Harris, K.M. (1999). Dendrite structure. In Dendrites, G. Stuart, N. Spruston, and M. Häusser, eds. (New York: Oxford University Press), pp. 1–34.

Girman, S.V., Sauve, Y., and Lund, R.D. (1999). Receptive field properties of single neurons in rat primary visual cortex. J. Neurophysiol. *82*, 301–311.

Gray, J.M., Hill, J.J., and Bargmann, C.I. (2005). A circuit for navigation in *Caenorhabditis elegans*. Proc. Natl. Acad. Sci. USA *102*, 3184–3191.

Hahnloser, R.H.R., Kozhevnikov, A.A., and Fee, M.S. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. Nature *419*, 65–70.

Harris, K.M., and Stevens, J.K. (1989). Dendritic spines of CA1 pyramidal cells in the rat hippocampus: serial electron microscopy with reference to their biophysical characteristics. J. Neurosci. *9*, 2982–2997.

Harris, K.M., Perry, E., Bourne, J., Feinberg, M., Ostroff, L., and Hurlburt, J. (2006). Uniform serial sectioning for transmission electron microscopy. J. Neurosci. *26*, 12101–12103.

Hartline, H.K., and Ratliff, F. (1972). Inhibitory interaction in the retina of Limulus. In Handbook of Sensory Physiology. Physiology of Photoreceptor Organs, Volume VII/2, M.G.F. Fuortes, ed. (Berlin: Springer-Verlag), pp. 381–447.

Hell, S.W. (2007). Far-field optical nanoscopy. Science *316*, 1153–1158.

Hobert, O., and Hall, D.H. (1999). Neuroanatomy: A second look with GFP reporters and some comments. Worm Breed. Gaz. *16*, 24.

Jain, V., Murray, J.F., Roth, F., Turaga, S., Zhigulin, V., Briggman, K.L., Helmstaedter, M.N., Denk, W., and Seung, H.S. (2007). Supervised learning of image restoration with convolutional networks. IEEE 11th International Conference on Computer Vision (ICCV), pp. 1–8.

Jin, D.Z., Ramazanoglu, F.M., and Seung, H.S. (2007). Intrinsic bursting enhances the robustness of a neural network model of sequence generation by avian brain area hvc. J. Comput. Neurosci. *23*, 283–299.

Kasthuri, N., Hayworth, K., Lichtman, J., Erdman, N., and Ackerley, C.A. (2007). New technique for ultra-thin serial brain section imaging using scanning electron microscopy. Microsc. Microanal. *13*, 26–27.

Katz, L.C., and Gurney, M.E. (1981). Auditory responses in the zebra finch's motor system for song. Brain Res. *221*, 192–197.

Kim, I.J., Zhang, Y., Yamagata, M., Meister, M., and Sanes, J.R. (2008). Molecular identification of a retinal cell type that responds to upward motion. Nature *452*, 478–482.

Kong, J., Fish, D.R., Rockhill, R.L., and Masland, R.H. (2005). Diversity of ganglion cells in the mouse retina: Unsupervised morphological classification and its limits. J. Comp. Neurol. *489*, 293–310.

Lashley, K.S. (1951). The problem of serial order in behaviour. In Cerebral Mechanisms in Behaviour (New York: Wiley), pp. 112–136.

Leonardo, A., and Fee, M.S. (2005). Ensemble coding of vocal control in birdsong. J. Neurosci. *25*, 652–661.

Li, M.R., and Greenside, H. (2006). Stable propagation of a burst through a one-dimensional homogeneous excitatory chain model of songbird nucleus HVC. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *74*, 011918.

Lichtman, J.W., and Sanes, J.R. (2008). Ome sweet ome: what can the genome tell us about the connectome? Curr. Opin. Neurobiol. *18*, 346–353.

Luo, L., Callaway, E.M., and Svoboda, K. (2008). Genetic dissection of neural circuits. Neuron *57*, 634–660.

Mackenzie, P.J., Kenner, G.S., Prange, O., Shayan, H., Umemiya, M., and Murphy, T.H. (1999). Ultrastructural Correlates of Quantal Synaptic Function at Single CNS Synapses. J. Neurosci. *19*, RC13.

Mahowald, M.A., and Mead, C.A. (1988). A silicon model of early visual processing. Neural Netw. *1*, 91–97.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. Nat. Rev. Neurosci. *5*, 793–807.

Marr, D. (1971). Simple memory: a theory for archicortex. Philos. Trans. R. Soc. Lond. B Biol. Sci. *262*, 23–81.

Masland, R. (2001a). Neuronal diversity in the retina. Curr. Opin. Neurobiol. *11*, 431–436.

Masland, R.H. (2001b). The fundamental plan of the retina. Nat. Neurosci. *4*, 877–886.

Masland, R.H. (2004). Neuronal cell types. Curr. Biol. *14*, R497–R500.

Matsuzaki, M., Ellis-Davies, G., Nemoto, T., Miyashita, Y., Iino, M., and Kasai, H. (2001). Dendritic spine geometry is critical for AMPA receptor expression in hippocampal CA1 pyramidal neurons. Nat. Neurosci. *4*, 1086–1092.

Mooney, R. (2000). Different subthreshold mechanisms underlie song selectivity in identified hvc neurons of the zebra finch. J. Neurosci. *20*, 5420–5436.

Mooney, R., and Prather, J.F. (2005). The HVC microcircuit: The synaptic basis for interactions between song motor and vocal plasticity pathways. J. Neurosci. *25*, 1952–1964.

Murthy, V.N., Schikorski, T., Stevens, C.F., and Zhu, Y. (2001). Inactivity produces increases in neurotransmitter release and synapse size. Neuron *32*, 673–682.

Naka, K.I., and Rushton, W.A. (1967). The generation and spread of S-potentials in fish (Cyprinidae). J. Physiol. *192*, 437–461.

Nelson, S.B., Hempel, C., and Sugino, K. (2006). Probing the transcriptome of neuronal cell types. Curr. Opin. Neurobiol. *16*, 571–576.

Nottebohm, F., Stokes, T.M., and Leonard, C.M. (1976). Central control of song in the canary, serinus canarius. J. Comp. Neurol. *165*, 457–486.

Ohki, K., Chung, S., Ch'ng, Y.H., Kara, P., and Reid, R.C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. Nature *433*, 597–603.

O'Keefe, J., and Nadel, L. (1978). The Hippocampus as a Cognitive Map (New York: Oxford University Press).

Orban, G. (1984). Neuronal Operations in the Visual Cortex (Berlin: Springer).

Oyster, C.W., Amthor, F.R., and Takahashi, E.S. (1993). Dendritic architecture of ON-OFF direction-selective ganglion cells in the rabbit retina. Vision Res. 33, 579–608.

Pierce, J.P., and Lewin, G.R. (1994). An ultrastructural size principle. J. Neurosci. 58, 441–446.

Ratliff, F. (1965). Mach Bands: Quantitative Studies on Neural Networks in the Retina (San Francisco: Holden-Day).

Redish, A.D., Battaglia, F.P., Chawla, M.K., Ekstrom, A.D., Gerrard, J.L., Lipa, P., Rosenzweig, E.S., Worley, P.F., Guzowski, J.F., McNaughton, B.L., et al. (2001). Independence of firing correlates of anatomically proximate hippocampal pyramidal cells. J. Neurosci. 21, RC134.

Reigl, M., Alon, U., and Chklovskii, D.B. (2004). Search for computational modules in the C. elegans brain. BMC Biol. 2, 25.

Rumelhart, D.E., and Norman, D.A. (1982). Simulating a skilled typist: a study of skilled cognitive-motor performance. Cogn. Sci. 6, 1–36.

Samsonovich, A., and McNaughton, B.L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. J. Neurosci. 17, 5900–5920.

Schikorski, T., and Stevens, C.F. (1997). Quantitative ultrastructural analysis of hippocampal excitatory synapses. J. Neurosci. 17, 5858–5867.

Shendure, J., Mitra, R.D., Varma, C., and Church, G.M. (2004). Advanced sequencing technologies: methods and goals. Nat. Rev. Genet. 5, 335–344.

Shepherd, G.M.G., and Harris, K.M. (1998). Three-dimensional structure and composition of CA3→ CA1 axons in rat hippocampal slices: implications for presynaptic connectivity and compartmentalization. J. Neurosci. 18, 8300–8310.

Smith, S.J. (2007). Circuit reconstruction tools today. Curr. Opin. Neurobiol. 17, 601–608.

Somers, D.C., Nelson, S.B., and Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. J. Neurosci. 15, 5448–5465.

Song, S., Sjostrom, P.J., Reigl, M., Nelson, S., and Chklovskii, D.B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. PLoS Biol. 3, e68.

Sporns, O., Tononi, G., and Kotter, R. (2005). The human connectome: a structural description of the human brain. PLoS Comput. Biol. 1, e42.

Striedter, G.F. (2006). Prècis of principles of brain evolution. Behav. Brain Sci. 29, 1–12.

Tsodyks, M., and Sejnowski, T. (1995). Associative memory and hippocampal place cells. Int. J. Neural Syst. 6, 81–86.

Vapnik, V.N. (2000). The Nature of Statistical Learning Theory (New York: Springer).

White, J.G., Southgate, E., Thomson, J.N., and Brenner, S. (1986). The structure of the nervous system of the nematode Caenorhabditis elegans. Philos. Trans. R. Soc. Lond. B Biol. Sci. 314, 1–340.

Wickersham, I.R., Lyon, D.C., Barnard, R.J.O., Mori, T., Finke, S., Conzelmann, K.K., Young, J.A.T., and Callaway, E.M. (2007). Monosynaptic restriction of transsynaptic tracing from single, genetically targeted neurons. Neuron 53, 639–647.

Worgotter, F., and Koch, C. (1991). A detailed model of the primary visual pathway in the cat: comparison of afferent excitatory and intracortical inhibitory connection schemes for orientation selectivity. J. Neurosci. 11, 1959–1979.

Yoshimura, Y., Dantzker, J.L., and Callaway, E.M. (2005). Excitatory cortical neurons form fine-scale functional networks. Nature 433, 868–873.