Neuron
Perspective

All My Circuits: Using Multiple Electrodes to Understand Functioning Neural Networks

Earl K. Miller[1,*] and Matthew A. Wilson[1]
[1]The Picower Institute for Learning and Memory and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
*Correspondence: ekmiller@mit.edu
DOI 10.1016/j.neuron.2008.10.033

Much of the work in systems neuroscience thus far has focused on the brain's parts studied individually. The past 20 years has seen the advent, rise, and application of multiple-electrode technology. This allows the study of the activity of many neurons simultaneously, which in turn has provided insight into how different neuron populations interact and collaborate to produce thought and action.

Introduction

Many of our current views of brain function center around hypotheses of interactions within and between different levels of networks: neurons, brain areas, systems. For example, the hippocampus is thought to consolidate memories in the cerebral cortex, top-down signals are thought to feedback to and modulate sensory cortex activity, etc. However, this is still mostly conjecture, inferred largely from indirect evidence such as anatomical connections and properties of the brain's individual parts studied in isolation.

This modular understanding stands to reason. Identifying and characterizing the brain's components is prerequisite to any integrated understanding of the whole, and technological limitations have largely restricted us to piecemeal investigation. But technical and methodological advancements over the past 20 years have led to increasing investigations on the network level. One new technique is human functional imaging. This provides a "big picture" of patterns of blood flow and, by examining their fluctuations, identifies putative large-scale, brain-wide networks (Dosenbach et al., 2007; Fox and Raichle, 2007).

Another recent advance is multiple-electrode neurophysiology, the implantation of up to 100 or more electrodes to study the activity of many neurons simultaneously, often in different brain regions. This adds to the long-standing single-electrode approach. It maintains the spatial and temporal precision needed to eavesdrop on brain function at the level of one of its elemental units, the neuron, but at the same time, it has the reach to examine neurons in a global context: the functioning of other neurons. This has led to new insights that would not otherwise have been possible. Here, we review this approach and some of its new insights into brain function.

Comparing Different Brain Areas

Multiple- and single-electrode approaches complement each other because each is well suited for different levels of investigation. With single electrodes, the focus is necessarily on the properties of each neuron. Investigators typically select only the most active neurons with properties of interest and tailor the experimental factors for those neurons. This is ideal for studying the unique characteristics of single neurons. With multiple electrodes, one loses this level of individualized detail. It is impractical and/or impossible to select individual neurons on the basis of a particular property on each of many electrodes and optimize conditions for each one simultaneously. But, in return, one gains greater sensitivity at comparing different neuron populations because the neurons can be compared under identical conditions. The painstaking nature of the single-electrode approach means that different investigators tend to study neurons in different brain areas with different experimental paradigms in different animals with different training histories. All of this can affect neural activity and thus confounds comparisons across areas, potentially producing spurious differences and/or obscuring real ones. Also, preselecting neurons for a particular property or response strength can normalize the neurons sampled and make neurons from different areas seem more alike than they actually are. With multiple electrodes, neurons typically are selected more randomly; any neuron encountered is studied. This approach is well suited for characterizing neuron properties at the population level: how a whole neuron population contributes to function, with a few neurons strongly activated and operating optimally, but many less strongly activated and operating under nonoptimal conditions.

Single-electrode studies have made important contributions to a global view of brain function. They have provided maps of sensory, memory, and motor functions and have allowed us to identify basic behavioral correlates and response properties of neurons across many regions of the brain. They can also uncover more shaded differences between brain areas if potential confounds are mitigated by using the exact same experimental paradigm. In an elegant series of studies, Romo and colleagues used the same somatosensory discrimination task to record individual neurons in multiple cortical areas. They found different degrees of strength and incidence of the sensory, memory, and motor signals. For example, de Lafuente and Romo (2006) found a progressive increase in the strength of correlation between activity and perceptual judgments from the parietal to frontal cortex. Likewise, Logothetis and colleagues recorded from individual neurons throughout visual cortex and found a stronger correlation between activity and perception in anterior areas (e.g., Sheinberg and Logothetis, 1997). This approach, however, is relatively rare because it is very labor intensive. Furthermore, because single-electrode data acquisition is slow, there can be

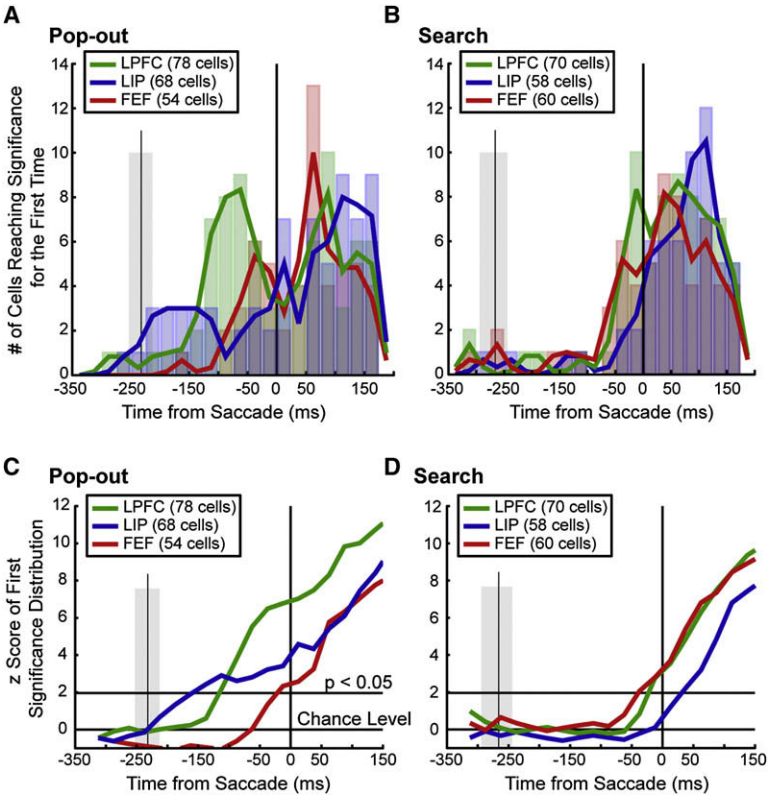

Figure 1. Timing of the Shift of Top-Down and Bottom-Up Attention

Bottom-up (pop-out) attention is shown in the left panels (A and C) and top-down (search) attention in the right panels (B and D) for the lateral prefrontal cortex (LPFC), frontal eye fields (FEF), and lateral intraparietal area (LIP).

(A and B) Distribution of times at which each neuron first began to carry significant information about the target location, relative to a saccade to the target (at 0 ms). Vertical black line indicates saccade; gray shaded regions indicate mean and ± one standard deviation of distribution of visual array onset.

(C and D) Normalized cumulative sum of the histograms shown in (A) and (B), respectively. During bottom-up task (pop-out), LIP neurons reflected the attention shift before LPFC and FEF neurons, whereas the opposite was found during top-down task (search). This suggests that top-down and bottom-up attention signals flow in opposite directions. (Reprinted from Buschman and Miller, 2007.)

differences in the animal's level of experience and ongoing performance between data collected from different brain areas, confounding comparisons. As a result, single electrodes may miss more subtle, but important, differences that can be better detected with the high fidelity of simultaneous recording from different brain regions under completely identical conditions.

Take, for example, a recent multiple-electrode study of top-down signals in monkey cortex. Top-down signals are derived from internal information, the current goal and the knowledge of the task demands needed to reach it. It is widely assumed that these signals arise in anterior cortical areas and flow to posterior cortex to modulate the processing of the bottom-up (sensory) signals. But there is almost no direct evidence for this—it is inferred from cortical connections and that more anterior cortical areas seem to have more complex, multimodal properties.

Buschman and Miller (2007) tested this by recording from up to 60 electrodes simultaneously implanted in frontal and parietal cortices. They found that bottom-up (automatic) shifts of attention to a salient sensory stimulus were registered with a shorter latency in the parietal than frontal cortex, first in the lateral intraparietal area (LIP), then in the lateral prefrontal cortex, and finally in the frontal eye fields, as if the bottom-up signals from the salient stimulus flowed through them in succession (Figure 1). By contrast, when monkeys had to search for a visual target based solely on memory, these top-down shifts of attention showed the opposite pattern, registering first in frontal cortex and then parietal cortex, suggesting that internal shifts of attention originate in the frontal cortex and are imposed on the posterior cortex.

Multiple electrodes can also give insight into the transformation of neural signals between areas. A powerful example is a study by Fyhn et al. (2007). The investigators found that spatial mapping properties of neuronal ensembles in the hippocampus were related to spatial phase relationships between neurons in the medial entorhinal cortex. This suggests that coordinated input from the entorhinal cortex converges onto hippocampal neurons to produce their place selectivity. Establishing the nature of such input-output transformations is essential to understanding the nature of neural computation. Because the spatial-behavioral correlates of neurons in these regions vary between animals and environments, this could only have been detected via simultaneous recording in both areas.

Multiple electrodes can also help compare learning-related neural activity across brain areas because there is often a high degree of variance in learning. Sometimes we may learn a new task quickly; other times more sluggishly. Multiple electrodes minimize this variance by allowing neurons to be compared under identical learning rates. For example, Pasupathy and Miller (2005) recorded simultaneously from the lateral prefrontal cortex (PFC) and the striatum of the basal ganglia (BG) during learning of new arbitrary rules (akin to "stop at red"). They found that learning-related changes in the striatum preceded those in the PFC. This suggests that the simple, arbitrary rules were first learned in the BG, which then trained slower learning mechanisms in the frontal cortex.

These are a few examples of how multiple electrodes can be used to compare and contrast neuron properties to gain insight into how information flows and is transformed between brain areas. Multiple electrodes can also capture more precise temporal dynamics, synchronous rhythms between neurons on the millisecond level. These rhythms are often not time-locked to external events—they rely on internal, not external, clocks—and thus can only be investigated by simultaneous recording from multiple brain sites. We discuss this next.

Synchronous Oscillations between Neurons

It has long been known that "brain waves" recorded from the human scalp exhibit a wide range of rhythmic oscillations (from

<1 to >100 Hz). There is a rich history linking changes in these oscillations to cognitive functions like attention and memory. Using multiple intracranial electrodes adds to this work because it can offer greater fidelity in detecting and localizing such activity.

Interest in this technique began to grow about 20 years ago with evidence for a role for synchronized oscillatory activity in perceptual organization (Gray et al., 1989; Eckhorn et al., 1988). Gray et al. (1989) found that neurons in visual cortex of cats were synchronized at about 40–60 Hz when they were activated by attributes of the same visual stimulus. By contrast, synchrony was lower or absent when different neurons were activated simultaneously by different stimuli, suggesting a role for neural synchronization in feature binding. The idea was that the average activity of neurons represents stimulus features, while synchronization between neurons binds those features together (Singer and Gray, 1995). These findings inspired hypotheses that synchronized oscillations play a role in consciousness (e.g., Crick and Koch, 1990; Llinás et al., 1998; Buzsáki and Draguhn, 2004).

Synchrony can be useful because it can enhance neural representations. Spikes arriving simultaneously at downstream neurons have a greater impact than unsynchronized spikes. This seems ideal for focal attention. Attention involves enhancing some stimulus representations at the expense of others. Evidence for this was reported by Desimone and colleagues. Fries et al. (2001) recorded local field potentials (LFPs) and spiking activity from recording sites in area V4 that had overlapping receptive fields. LFPs are often used to detect oscillations because they reflect coordinated activity across large numbers of neurons. When monkeys' attention was directed to a particular visual stimulus, neurons activated by the stimulus showed increased synchronized gamma band (30–90 Hz) oscillations and a reduction in low-frequency (<17 Hz) synchronization relative to V4 sites activated by an unattended stimulus. As the authors pointed out, this synchronization could effectively increase the gain at the postsynaptic targets of these neurons. When monkeys searched for a particular visual feature (e.g., a color), V4 neurons whose receptive field contained that feature showed gamma band oscillatory synchronization (Bichot et al., 2005).

Synchrony can also enhance neural processing by putting the brain and the external world in lockstep. Lakatos et al. (2008) presented monkeys with a stream of sequential visual and auditory stimuli. The exact timing of their presentation was jittered a bit, but their average within each stimulus stream was at about 1.5 Hz, and the two streams were out of phase with each other. When monkeys attended to the visual or auditory stream, LFPs and spikes in visual cortex synchronized to the rhythm of that stream and not to the rhythm of the unattended stream. Thus, attention can also amplify the gain of neural processing of stimuli by matching the brain's rhythms to those of the external world.

Oscillatory synchronization may also regulate communication between brain areas. When neurons are simultaneously depolarized, they are more susceptible to influence from each other. (Conversely, when they are out of phase, they are less likely to communicate.) Evidence for this comes from a number of studies. Bressler et al. (1993) found increased broad-band coherence between LFPs in visual, motor, and prefrontal cortex of monkeys during performance of a visual discrimination task. Roelfsema et al. (1997) observed that patterns of synchrony between areas depend on task and function. Cats were trained to respond to a change in a visual stimulus. During the task, visual cortex areas synchronized more strongly to other visual areas, and motor areas to other motor areas. Furthermore, changes in the visual stimulus were associated with increased synchrony between visual, but not motor, areas. This suggests task-dependent dynamic coupling between functionally related brain areas. Pesaran et al. (2008) trained monkeys to alternate between an instructed choice and free choice of one of three different visual targets and found increased correlations between spikes and LFPs between the parietal cortex and premotor cortex during the free choice. This suggested a free-choice circuit that coordinates activity between these areas to influence the choice. Saalmann et al. (2007) recorded simultaneously from area LIP and area MT, a more posterior visual cortical area that is earlier in the visual cortical pathway. These areas synchronized during a visual motion matching task, with the LIP leading MT. This, like the Buschman and Miller (2007) study cited above, suggests an anterior-to-posterior flow of top-down signals.

Communication between areas may also be regulated by details of how oscillations match up. Womelsdorf et al. (2007) found variability in the phase offset between synchronized oscillations from different visual cortical areas in cats and monkeys. The authors sorted oscillatory spiking and LFP activity according to their phase offset and found stronger synchrony between areas at specific offsets between their oscillations, suggesting that changes in the phase offset of the oscillations could dynamically weight the strength of connection between brain areas.

The frequency of the oscillations may also matter for communication. Buschman and Miller (2007), in their study of top-down versus bottom-up shifts in attention (see above), found increased synchrony between LFPs in the frontal and parietal cortex in two distinct frequency bands. There was a greater increase in the lower band (22–34 Hz) for top-down shifts of attention and a greater increase of synchrony in a higher band (35–55 Hz) for bottom-up shifts (Figure 2). The authors noted that lower-frequency synchrony is more forgiving of the exact timing of spikes and suggested that it may thus reflect the "broadcast" of the top-down signal on a wide anatomical scale. By contrast, the higher-frequency synchrony may reflect local interactions that enhance the representation of the salient stimulus for bottom-up shifts in attention. In any case, these results suggest that the brain can use different frequency bands, like two different spots on the radio dial, for communication related to two different cognitive functions.

Rhythmic synchrony is not only a cortical phenomenon. Siapas and Wilson (1998) found evidence for synchrony in aiding communication between the hippocampus and cortex. During slow-wave sleep, hippocampal ripples, which are a prominent LFP signature of offline hippocampal activity, were found to be correlated with the onset of neocortical spindles, a prominent LFP signature of offline neocortical activity (Siapas and Wilson, 1998; Sirota et al., 2003). Because sleep (as well as other brain states) involves neural processing that is largely governed by internally controlled variables, there is no behavioral measure that could have been used to detect this synchrony; it could only have been found by simultaneously recording activity in the hippocampus and neocortex. Siapas et al. (2005) also found neuron

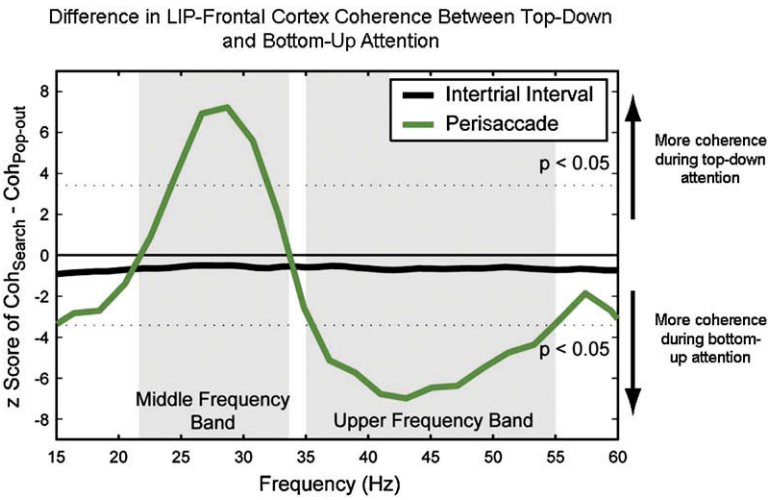

Figure 2. Differences in Local Field Potential Coherence between LIP and Frontal Cortex during Top-Down and Bottom-Up Attention

Coherence during the shift of attention (perisaccade, green) compared to a baseline from the intertrial interval (black) is plotted. Bottom-up coherence was subtracted from top-down coherence. Dotted lines indicate significance levels ($p < 0.05$, corrected for multiple comparisons). Differences above the upper dotted line indicate significantly more coherence during top-down than bottom-up attention, and differences below the lower dotted line indicate significantly more coherence during bottom-up than top-down attention. (Modified from Buschman and Miller, 2007.)

activity in the rat prefrontal cortex phase-locked to the theta oscillations (6–12 Hz) that are prominent in the hippocampus. Jones and Wilson (2005) further discovered that this phase-locking was selectively engaged during the choice phase of a spatial working memory task and was correlated with subsequent correct choice behavior. The correct-choice effect was not evident in the activity of single neurons alone.

Rhythmic synchrony can also occur between different subcortical structures. DeCoteau et al. (2007) recorded from the rat hippocampus and striatum during learning of a T maze. They found increased coherence between hippocampal and striatal theta rhythms that was stronger in rats that successfully learned the task. Interestingly, in the successful rats, these oscillations had an antiphase relationship in proportion to the level of learning. This suggests that learning may involve changes in the precise coordination of rhythmic activity between the hippocampus and striatum in addition to the well-known signal-neuron correlates.

These and other studies raise the possibility that anatomy offers the scaffolding for potential communication between areas. Synchronized, rhythmic activity between areas may regulate their effective connectivity, dynamically controlling communication by opening preferred channels between areas only when they need to communicate (Engel et al., 2001; Salinas and Sejnowski, 2001). Next, we illustrate how precise timing of neural activity may also play a role in representing information.

Detailed Timing of Activity within Neural Ensembles

It is widely accepted that information is encoded in the brain by distributed activity across ensembles of neurons both within and between brain areas. There is increasing evidence that within these ensembles, information is encoded in patterns of activity of neurons on the time scale of individual action potentials (milliseconds) as well as on the time scale of firing rate changes over time (seconds). Multiple electrodes afford the ability to examine more complex, dynamic, higher-order structures in neural activity in which correlations and timing are preserved across the large-scale ensembles that are believed to underlie the actual coding of information within the brain. The ability to decode behavioral neural correlates in the hippocampus was demonstrated by Wilson and McNaughton (1993), who used the activity of approximately 100 simultaneously recorded place cells to reconstruct behavioral trajectories during spatial exploration and identify rapid changes in neural coding during novel exposure to an environment.

The ability to identify patterns on the time scale of milliseconds was demonstrated by Lee and Wilson (2002). The investigators found that precise activity patterns of multiple hippocampal place cells during movement was evident during subsequent periods of non-REM sleep, as if the experience were being replayed in proper sequential order but at a compressed time scale (Figure 3). Previous work relied on the use of lower-order correlation methods to search for signatures of memory reactivation (Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Hoffman and McNaughton, 2002). Lee and Wilson were able to demonstrate the power of multiple-electrode methods to detect the higher-order structure of temporal relationships across neurons by explicitly comparing the statistical power of pairwise correlation methods with higher-order sequence detection measures. The reactivated patterns were barely detectable using simple pairwise correlations. By contrast, the use of higher-order structure increased the power of detection by five orders of magnitude. This allowed analysis of the content, structure, and significance of individual reactivation events lasting on the order of 100 ms. Buzsáki and colleagues found similar results using a template matching procedure to identify recurring spike sequences during awake and sleep states (Nadasdy et al., 1999).

This was extended in a study by Foster and Wilson (2006). They, like Lee and Wilson, found that precisely timed sequential patterns of hippocampal neural activity that were initially expressed as an animal ran down a simple linear track were replayed. These replayed events occurred not during sleep but rather when the animal stopped at the end of the track for food reward. But there was a critical difference: the patterns were replayed in reversed time order, that is, backward. Because many single neurons could be simultaneously sampled, these reverse replay events could be detected immediately following the first lap down the track. The ability to analyze activity patterns following a single novel exposure was critical in linking these events to memory of recent experience. The authors suggested that the reversal may be useful for linking memory patterns to the consequence of the behaviors they represent: in this case, a reward. Subsequent studies by Diba and Buzsáki (2007) found that both forward and reverse sequences could be replayed during quiet wakefulness, and Johnson and Redish (2007) were able

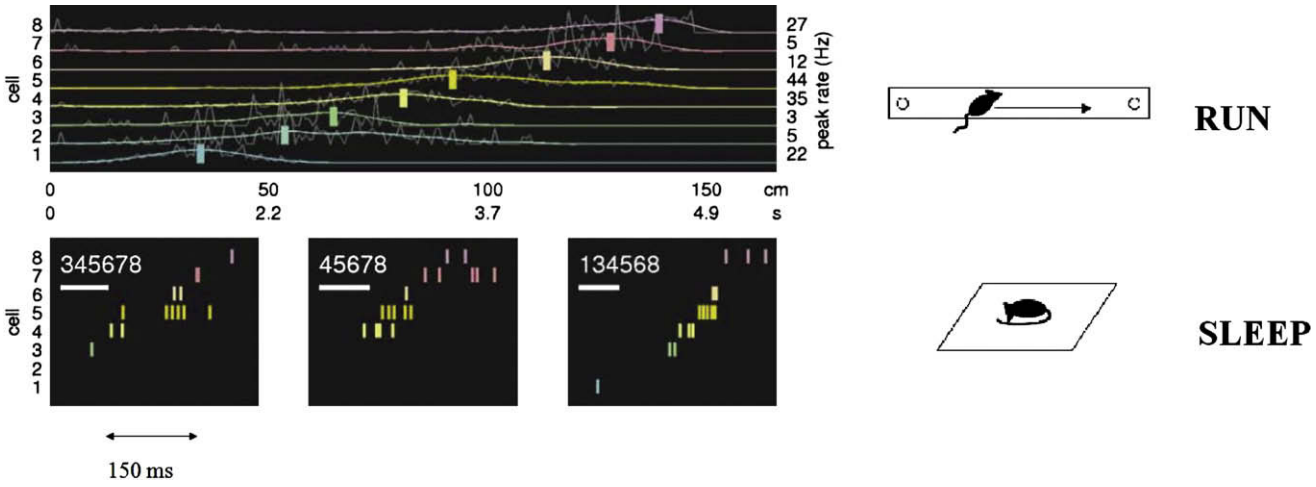

Figure 3. Replay of Neural Firing Patterns during Sleep

Upper panel shows the average firing rate of eight simultaneously recorded hippocampal place cells as a function of location of the place field (ticks indicate location of peak firing) along a linear track (shown to the right). As the animal walks along the track from left to right, the neurons fire in order from 1 to 8 over the course of the approximately 5 s that it takes to traverse the track. Lower panels display three examples of brief events that occur during subsequent slow-wave sleep in which the order of neuronal firing recapitulates the previous behavioral sequence, but in a compressed form. (Modified from Lee and Wilson, 2002.)

to correlate sequence patterns during running with future choice behavior. In each of these cases, the ability to simultaneously record ensembles of single neurons allowed evaluation of neural activity patterns that may contribute to discrete learning or behavioral events.

Replay of sequential patterns of activity at an intermediate time scale of about one second has also been observed in both the hippocampus and sensory cortex during periods of slow-wave sleep. Ji and Wilson (2007) used multiple electrodes to record from both hippocampus and visual cortex of rats and found replay of sequential patterns of activity associated with previous behavior in both areas. Furthermore, the replay in these brain areas was coordinated: it occurred when both areas were in an excited "up-state." The simultaneous monitoring of multiple neurons in each area allowed the content of reactivated sequences to be evaluated and compared across areas. Based upon the relative timing of events in both areas, the authors concluded that the visual cortex initiated the reactivation of memory sequences in the hippocampus that were followed by sequence memory reactivation in the cortex, thus providing a detailed picture of the dialog between the hippocampus and neocortex during slow-wave sleep.

The ability to identify higher-order temporal structure of neural ensembles at longer time scales during REM sleep was demonstrated in a study by Louie and Wilson (2001). The investigators used a template approach to match the activity patterns of many neurons in the hippocampus across entire REM episodes with corresponding patterns expressed during periods of running on a circular track. They concluded that behavioral sequences spanning minutes, expressed in the patterns of hippocampal place cells, were replayed during individual REM sleep episodes, each lasting a similar amount of time.

While these studies interpreted the structure of relative timing relationships with respect to previously measured behavioral correlates to establish their significance and to allow reconstruction of coded content, it is also possible to use the consistency of patterns independent of any obvious behavioral correlates to identify potentially meaningful internal representations. Buzsáki and colleagues have described such patterns as the product of cell assemblies (Harris et al., 2003; Pastalkova et al., 2008).

Summary

In this short review, we have focused on multiple-electrode studies that have yielded new insight into brain function. Multiple electrodes have also been used for the rapid online decoding of neural signals needed to control prosthetics (e.g., Carmena et al., 2003; Hochberg et al., 2006; Fetz, 2007; Mulliken et al., 2008). But all of this is just a step in the right direction. Understanding neural circuits will depend on more than just observing neural activity and correlating it with behavior. Ultimately, it will depend on establishing cause-and-effect relationships within the brain by direct perturbation of the circuits. Combining multiple-electrode technology with new molecular genetic techniques holds great promise. We will be able to precisely manipulate the activity of specific populations of neurons while recording from multiple electrodes in order to measure the consequences across the brain. This will no doubt provide exciting new insights into functioning neural circuits. We look forward to the progress to come over the next 20 years.

ACKNOWLEDGMENTS

The authors wish to thank Timothy Buschman, Markus Siegel, and Marlene Wicherski for valuable comments and discussions. E.K.M. thanks the NIMH and NINDS for support. M.A.W. thanks the NIMH for support. We both thank the Picower Institute for Learning and Memory and the RIKEN-MIT Neuroscience Research Center.

REFERENCES

Bichot, N.P., Rossi, A.F., and Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. Science *308*, 529–534.

Bressler, S.L., Coppola, R., and Nakamura, R. (1993). Episodic multiregional cortical coherence at multiple frequencies during visual task performance. Nature *366*, 153–156.

Buschman, T.J., and Miller, E.K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science *315*, 1860–1862.

Buzsáki, G., and Draguhn, A. (2004). Neuronal oscillations in cortical networks. Science *304*, 1926–1929.

Carmena, J.M., Lebedev, M.A., Crist, R.E., O'Doherty, J.E., Santucci, D.M., Dimitrov, D.F., Patil, P.G., Henriquez, C.S., and Nicolelis, M.A. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. PLoS Biol. *1*, E42.

Crick, F., and Koch, C. (1990). Toward a neurobiological theory of consciousness. Semin. Neurosci. *2*, 263–275.

DeCoteau, W.E., Thorn, C., Gibson, D.J., Courtemanche, R., Mitra, P., Kubota, Y., and Graybiel, A.M. (2007). Learning-related coordination of striatal and hippocampal theta rhythms during acquisition of a procedural maze task. Proc. Natl. Acad. Sci. USA *104*, 5644–5649.

de Lafuente, V., and Romo, R. (2006). Neural correlate of subjective sensory experience gradually builds up across cortical areas. Proc. Natl. Acad. Sci. USA *103*, 14266–14271.

Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. Nat. Neurosci. *10*, 1241–1242.

Dosenbach, N.U., Fair, D.A., Miezin, F.M., Cohen, A.L., Wenger, K.K., Dosenbach, R.A., Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E., et al. (2007). Distinct brain networks for adaptive and stable task control in humans. Proc. Natl. Acad. Sci. USA *104*, 11073–11078.

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., and Reitboeck, H.J. (1988). Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. Biol. Cybern. *60*, 121–130.

Engel, A.K., Fries, P., and Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. Nat. Rev. Neurosci. *2*, 704–716.

Fetz, E.E. (2007). Volitional control of neural activity: implications for brain-computer interfaces. J. Physiol. *579*, 571–579.

Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature *440*, 680–683.

Fox, M.D., and Raichle, M.E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nat. Rev. Neurosci. *8*, 700–711.

Fries, P., Reynolds, J.H., Rorie, A.E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. Science *291*, 1560–1563.

Fyhn, M., Hafting, T., Treves, A., Moser, M.B., and Moser, E.I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. Nature *446*, 190–194.

Gray, C.M., König, P., Engel, A.K., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. Nature *338*, 334–337.

Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. Nature *424*, 552–556.

Hochberg, L.R., Serruya, M.D., Friehs, G.M., Mukand, J.A., Saleh, M., Caplan, A.H., Branner, A., Chen, D., Penn, R.D., and Donoghue, J.P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature *442*, 164–171.

Hoffman, K.L., and McNaughton, B.L. (2002). Coordinated reactivation of distributed memory traces in primate neocortex. Science *297*, 2070–2073.

Ji, D., and Wilson, M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat. Neurosci. *10*, 100–107.

Johnson, A., and Redish, A.D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. J. Neurosci. *27*, 12176–12189.

Jones, M.W., and Wilson, M.A. (2005). Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. PLoS Biol. *3*, e402.

Lakatos, P., Karmos, G., Mehta, A.D., Ulbert, I., and Schroeder, C.E. (2008). Entrainment of neuronal oscillations as a mechanism of attentional selection. Science *320*, 110–113.

Lee, A.K., and Wilson, M.A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. Neuron *36*, 1183–1194.

Llinás, R., Ribary, U., Contreras, D., and Pedroarena, C. (1998). The neuronal basis for consciousness. Philos. Trans. R. Soc. Lond. B Biol. Sci. *353*, 1841–1849.

Louie, K., and Wilson, M.A. (2001). Temporally structured REM sleep replay of awake hippocampal ensemble activity. Neuron *29*, 145–156.

Mulliken, G.H., Musallam, S., and Andersen, R.A. (2008). Forward estimation of movement state in posterior parietal cortex. Proc. Natl. Acad. Sci. USA *105*, 8170–8177.

Nadasdy, Z., Hirase, H., Czurko, A., Csicsvari, J., and Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. J. Neurosci. *19*, 9497–9507.

Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. Science *321*, 1322–1327.

Pasupathy, A., and Miller, E.K. (2005). Different time courses for learning-related activity in the prefrontal cortex and striatum. Nature *433*, 873–876.

Pesaran, B., Nelson, M.J., and Andersen, R.A. (2008). Free choice activates a decision circuit between frontal and parietal cortex. Nature *453*, 406–409.

Roelfsema, P.R., Engel, A.K., König, P., and Singer, W. (1997). Visuomotor integration is associated with zero time-lag synchronization among cortical areas. Nature *385*, 157–161.

Saalmann, Y.B., Pigarev, I.N., and Vidyasagar, T.R. (2007). Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. Science *316*, 1612–1615.

Salinas, E., and Sejnowski, T.J. (2001). Correlated neuronal activity and the flow of neural information. Nat. Rev. Neurosci. *2*, 539–550.

Sheinberg, D.L., and Logothetis, N.K. (1997). The role of temporal cortical areas in perceptual organization. Proc. Natl. Acad. Sci. USA *94*, 3408–3413.

Siapas, A.G., and Wilson, M.A. (1998). Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. Neuron *21*, 1123–1128.

Siapas, A.G., Lubenov, E., and Wilson, M.A. (2005). Prefrontal phase-locking to hippocampal theta oscillations. Neuron *46*, 141–151.

Singer, W., and Gray, C.M. (1995). Visual feature integration and the temporal correlation hypothesis. Annu. Rev. Neurosci. *18*, 555–586.

Sirota, A., Csicsvari, J., Buhl, D., and Buzsáki, G. (2003). Communication between neocortex and hippocampus during sleep in rodents. Proc. Natl. Acad. Sci. USA *100*, 2065–2069.

Skaggs, W.E., and McNaughton, B.L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. Science *271*, 1870–1873.

Wilson, M.A., and McNaughton, B.L. (1993). Dynamics of the hippocampal ensemble code for space. Science *261*, 1055–1058.

Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. Science *265*, 676–679.

Womelsdorf, T., Schoffelen, J.M., Oostenveld, R., Singer, W., Desimone, R., Engel, A.K., and Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. Science *316*, 1609–1612.