9

Coordination in Circuits

Mayank Mehta, György Buzsáki, Andreas Kreiter,
Anders Lansner, Jörg Lücke, Kevan Martin,
Bita Moghaddam, May-Britt Moser,
Danko Nikolić, and Terrence J. Sejnowski

Introduction

What are the mechanisms underlying the emergence of mind from the activity of groups of neurons? This is a difficult question that has to be addressed at many levels of neural organization, all of which need to be integrated. The following discussion of the set of neural mechanisms, neural activity patterns, and animal behaviors sketches a few simple, but general and robust, neural mechanisms at all the different levels ranging from synapses to neurons, to networks and behavior, and is illustrated using experimental observations primarily from cortical and hippocampal activity patterns during behavior. The focus is on a few key mechanisms such as energy-efficient sparse codes of mental representations, need for synchrony among sparse codes for information transmission, and the contribution of recurrent connections between excitatory and inhibitory neurons in generating synchronous activity, oscillations, and competition among networks to facilitate fast and flexible behavior.

At the level of the mind, animals can perceive different components of a rapidly changing natural scene such as luminance, contrast, local features (e.g., lines), and global features (e.g., shapes). Similarly, when an animal navigates in the world, neurons can flexibly represent the position of the animal in a given environment, the composition of the environment, the head direction, the running speed, etc. These mental representations of the world are flexible and dynamic, determined and modulated by a range of environmental, behavioral, and neural parameters.

What are the mechanisms by which the brain generates these flexible mental representations? How do these mental representations across different brain regions interact to generate perception and decision making?

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

What Is the Right Anatomical Level of Investigation?

The flexible behaviors outlined above depend on common features of the neuroanatomical substrates underlying these mental representations. Historically, models of cortical circuits have been designed to capture a specific feature of the experimental data. An alternative strategy that leads to a "predictive connectivity" starts with the assumption that there is a basic ("canonical") circuit common across the entire neocortex (see Figure 9.1). This assumption justifies the use of the rich cache of structure, function, and neurochemistry to build more biologically realistic models. The models can then be challenged to provide an explanation of cortical activity patterns. To the extent that the simulations are successful, the model can quantitatively predict the connectivity pattern of the circuits in that area. The ability to test a prediction about structure is a radical departure from the traditional descriptive and anatomical methods of circuit analysis. In combination with new tools for tracing pathways and combining structure with function, this predictive structural modeling will not only greatly accelerate circuit analysis in neocortex, but will provide a far more

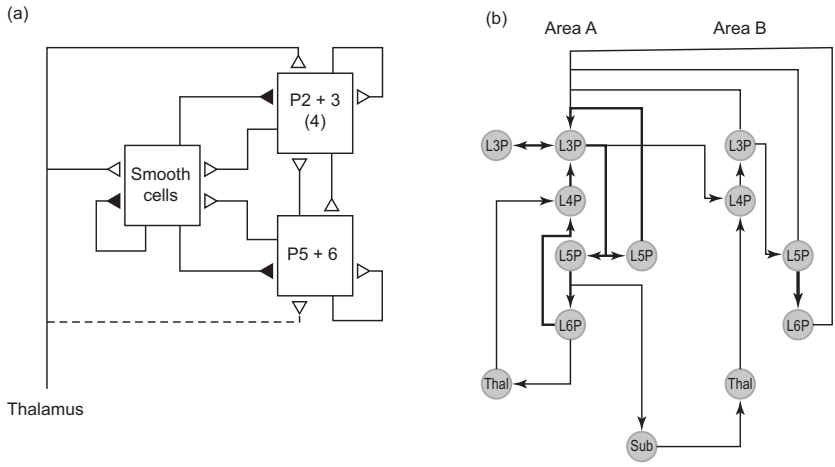

Figure 9.1 (a) Canonical circuit of neocortex. Three populations of neurons interact: the inhibitory, GABAergic population indicated by smooth cells; the excitatory population by a superficial layer population (P2+3 (4)); and a deep layer population (P5 + 6). The connections between them are indicated by edges and arrows. The functional weights of the connections are indicated by the thickness of the edges. (b) Graph of the dominant interactions between significant excitatory cell types in neocortex and their subcortical relations. The nodes of the graph are organized spatially; vertical dimension corresponds to the layers of cortex and horizontal to its lateral extent. Edges and arrows indicate the relations between excitatory neurons (P: pyramidal) in a local patch of neocortex, which are essentially those described originally by Gilbert and Wiesel (1983) and Gilbert (1983) for visual cortex. Thin edges indicate excitatory connections to and from subcortical structures and inter-areal connections. Thal: thalamus; Sub: other subcortical structures, such as the basal ganglia.

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

comprehensive and synthetic explanation of the computational strategies used in different cortical areas.

There are several common features of cortical anatomy shared across many brain regions. The neocortex is organized in a number of layers and columns. Major elements of the canonical cortical microcircuit of a column have been well described. For example, the excitatory, pyramidal neurons within layer 2/3 in a neocortical column are recurrently connected, as are the neurons within layer 5/6. The thalamic inputs arrive primarily in layer 4 whereas layers 5 and 6 send the output of a cortical column to other brain areas. Further, synaptic inputs are exquisitely organized on the extensive pyramidal neuronal dendrites, which have nonlinear properties. Also, there are neuromodulatory inputs specific to different layers, which could play a key role in information processing, as will be discussed later.

In addition to many other features of the canonical cortical circuit, there is one common feature in all of these layers, namely the presence of a wide variety of GABAergic inhibitory interneurons. While these inhibitory neurons comprise only about 20% of the neural population, they strongly control cortical activity because of their recurrent connections to excitatory neurons. Inhibitory synapses are often found near the soma, which can influence all excitatory inputs flowing from the dendrites to soma. Further, not only are the excitatory neurons recurrently connected to each other within and across layers, so are the inhibitory neurons. Such recurrently connected excitatory-inhibitory networks (denoted E-I networks) are ubiquitous: They are found not only in most parts of neocortex and hippocampus, but in many other structures as well.

In addition, cortical circuits receive powerful neuromodulatory inputs. Monoamine neuromodulators dopamine, norepinephrine, and serotonin are released by cells in discrete nuclei in the brainstem and midbrain that project heavily to basal ganglia and cortical regions. Most psychotherapeutic and psychoactive drugs, which have profound effects on cognition, act on receptors of monoamines. These include antidepressant and antipsychotic drugs as well as hallucinogens and stimulants such as amphetamine, cocaine, and methylphenidate. This suggests that monoamines are a critical component of neuronal machinery underlying perception and complex behaviors. The topography of their projections to cortical regions, as well as their targeted receptors, is quite diverse. For example, dopamine projections tend to be heavier to deeper cortical layers whereas norepinephrine projections are heavy in superficial layers. The receptor type and the signal transduction mechanisms used by these neuromodulators are diverse, with the exception of one of the subtypes of serotonin receptors—G-protein coupled receptors. The localization of these receptors is also specialized. For example, some subtypes of serotonin receptors are primarily localized on GABA interneurons. In addition, monoamine receptors, especially the dopamine receptors, are mostly localized extrasynaptically, suggesting that they produce slow and somewhat sustained effects on the state of cortical microenvironments. In the case of dopamine, the density of dopamine

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

transporters in cortical areas is sparse and thus the release of dopamine has the capacity to diffuse away from presynaptic sites and act more diffusely.

For generality, the following discussion will be focused on how generic E-I networks process information and how neuromodulators influence the process, keeping in mind that this is a simplification. The precise details, such as various intrinsic properties of the different cell types, their dendritic geometry, and the exact connectivity patterns within and across cortical columns, would need to be investigated in the future to understand the neuroanatomical basis of perception. The goal is to probe the system at progressively more detailed biological levels, while deciphering the emergent properties of the system at each level, which can be robustly tested both experimentally and theoretically.

How Do E-I Neural Networks Oscillate?

Converging evidence suggests that the E-I network is balanced; that is, the total amount of excitation and inhibition are comparable most of the time, even though the total amount of activity can vary over a wide range. In the absence of stimuli, under most conditions, cortical neurons are active at a low rate with an ensemble average of about 0.1 Hz, which varies systematically across layers. When stimuli arrive, a small fraction of the excitatory stimulus-responsive neurons increase their instantaneous firing rates to 10 or even 100 Hz. This increased activation of pyramidal neurons drives the feedback inhibitory neurons, which in turn briefly shut down the pyramidal neurons. This reduces the excitatory drive onto inhibitory interneurons which generates release from inhibition synchronously across a number of pyramidal neurons. Consequently, pyramidal neurons increase their spiking activity in synchrony. A key parameter governing the frequency of such E-I network synchronized oscillations is the time constant of the inhibitory $GABA_A$ receptors of 10–30 ms, resulting in about 30–100 Hz gamma frequency oscillations. Thus, oscillations in the gamma range can be a signature of cortical activation. Notably, this simple description for generating gamma oscillations applies only to excitatory neurons connected recurrently to inhibitory ones. The additional recurrent connections within the populations of neurons of the same type would profoundly influence the strength and frequency of oscillations. Synchronization of oscillations across different E-I networks is another, even more complex process.

Thus, in a simple scenario, stimulus-driven elevation in the firing of excitatory neurons can have two concurrent effects: (a) elevated firing of the excitatory and inhibitory neurons, and (b) synchronized oscillations. Notably, both the oscillation frequencies and the degree of synchronicity between oscillations influence neural information processing.

It is important to discuss the following four points: the range of oscillation frequencies, alternative mechanisms for generating oscillations, mechanisms that modulate oscillation frequency, and synchrony without oscillations.

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

Oscillations with frequencies ranging from 0.1–200 Hz have been commonly observed in several neocortical areas as well as in the hippocampus and the olfactory system. Here, the focus is on oscillations that occur on the timescales relevant for processing natural stimuli (i.e., less than about half a second), so that they can modulate neural processing. This means that the focus will be on frequencies greater than about 2 Hz. Synchronized oscillations of a variety of frequencies appear in numerous brain regions during perception, attention, working memory, motor planning, sleep, epilepsy and Parkinson's disease. For example, 4–12 Hz theta oscillations are prominent in the rodent and primate hippocampus during spatial exploration. They have been reported in visual, parietal, and prefrontal cortices during maintenance of information in working memory. Somewhat higher frequencies, 10–30 Hz or beta frequency oscillations, have been reported in the visual and motor cortices. The 40–120 Hz gamma oscillations are induced by visual stimuli in numerous visual cortical areas and prefrontal cortex, and they also occur in the hippocampus. In addition, bursts of 140–250 Hz ripple oscillations occur in the hippocampus during quiet wakefulness. The focus here is on the theta and gamma oscillations that appear in the neocortex and hippocampus during cognitive tasks.

The E-I network is not the only mechanism that can generate synchronized oscillations. Neurons are endowed with a variety of conductances and intrinsic mechanisms which can also make them respond rhythmically when a fixed amount of current or neuromodulators are applied, even when isolated from a network.

The key issue is: How do groups of oscillating neurons get synchronized? Invariably, this is achieved through their coupling with the rest of the network. For example, neurons in the reticular nucleus of thalamus oscillate in isolation, whereas these oscillations are synchronized through coupling between these neurons directly or through the thalamocortical loop. This mechanism is thought to generate sleep spindles. Similarly, septal neurons oscillate in isolation and are likely synchronized by their recurrent connection to the hippocampus, resulting in synchronous theta oscillations. Finally, even when the E-I network in a cortical column oscillates at gamma frequency, an important question is: How do oscillations of different cortical columns synchronize? In all these cases, further questions arise: How do these oscillators respond to a stimulus? Does an excitatory spike from another oscillator speed up the subsequent spike from a given oscillator or delay it? In other words, how do the oscillations change as a function of the phase at which inputs arrive from other oscillators? Thus, it is important to study how oscillations change as a function of the phase at which inputs arrive. Such dependence is called a phase resetting curve and has been investigated for a variety of physical and neural systems.

The frequency of neural oscillations can be modulated by several means. For example, neuromodulators can generate a threefold change in the effective time constant of $GABA_A$ receptors, resulting in a concomitant change in the frequency of gamma oscillations. Further, cholinergic levels alter spike

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

frequency adaptation of pyramidal neurons, which would influence the E-I balance and spike timing in an E-I network. The levels of neuromodulators change with behavioral state and attention, resulting in state-dependent modulation of amplitude, power, and synchrony of neural oscillations. For example, neuromodulators can raise the membrane potential of neurons. This can make it easier for the neurons to respond to a small amount of stimulation, and may make the E-I network more likely to oscillate.

Finally, two important features of oscillations need to be distinguished: synchrony and rhythmicity. Synchrony can occur without rhythmicity and vice versa. In particular, synchronous activation of groups of neurons occurs almost invariably, often without oscillations, when a strong stimulus is abruptly activated.

Neural synchrony is important for efficient and rapid transmission of information between brain regions. For example, individual neurons in a cortical column receive information from only about 100 thalamic neurons. To integrate this input and generate a spike, a cortical neuron typically needs to be depolarized by about 20 mV. Given the small amplitude ($\leq$1 mV) and short duration (~10 ms) of cortical excitatory, AMPAR-mediated postsynaptic potentials, this small amount of thalamic neurons can only activate an entire cortical column if the inputs are synchronized within a 10 ms time window.

There are several advantages of transmitting information using synchronous activity. First, only a small number of active neurons, or a sparse code, is sufficient to transmit information from one area to another, as opposed to asynchronous transmission which would require more activity. Given that spike generation consumes energy, sparse synchronous codes are energy efficient. Second, compared to the asynchronous systems, synchronous sparse codes can be brief, allowing the system to respond rapidly to changing stimuli. Finally, the synchrony-based codes allow the system to be flexible, requiring only small changes in the relative timings of groups of neurons to make one group drive the downstream neurons more effectively than through asynchronous codes.

Synchronous activity can be generated by two different mechanisms. Synchrony can be evoked by a transient stimulus or through dynamic interactions between internal temporally organized activity patterns, such as oscillations. The latter can generate synchronous activity across multiple cycles of oscillation. Subsequent sections will discuss computational advantages of this process.

Why Aren't Gamma Oscillations Always Observed during Behavior?

The E-I network is ubiquitous, and synchronized oscillations are a likely mode of the E-I networks, yet there are instances in which oscillations are not apparent. There are several reasons for this. It is often difficult to detect gamma

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

oscillations in the spike train of a single neuron because, even when modulated by gamma oscillations, neurons often do not spike at sufficiently high rates to be active on every gamma cycle. Also, neurons often join the population rhythm for only brief epochs. This probabilistic firing and rapidly changing neural assembly can appear to be nonrhythmic when analyzing the activity of single units in isolation. An alternative method for detecting synchronized oscillations of ensembles of neurons is through the measurement of the local field potential (LFP) and the analysis of its power spectrum. Notably, the power in the LFP spectrum decays inversely with the increase in frequency, which makes it more difficult to detect activity in the gamma band than in the lower frequencies. Here, analysis methods that compensate for this systematic tendency of power spectra can improve the ability to detect gamma oscillations. Further, the nature of the electrode used to measure the LFP influences the power of the measured oscillations: sharp, high-impedance electrodes integrate activity over a small pool of neurons, which may not be sufficient to detect synchronous oscillations above the electrical noise. Signals can be improved by using blunter electrodes, which additionally allow the detection of synchronous gamma activity in multiunit activity.

Presentation of visual stimuli within a receptive field, such as moving bars or gratings, are likely to generate synchronized gamma oscillations for at least several seconds in anesthetized animals. Similar oscillations may be more difficult to detect in behaving animals because the fixations last shorter as eyes move on average three times a second, moving the stimuli rapidly in and out of the receptive fields. This may augment the fluctuations of neural activity, resulting in rapid fluctuation of gamma power and frequency, and making detection by standard methods difficult. Time-frequency domain analyses may counter this problem by estimating the strength of gamma oscillations in smaller, relatively unperturbed windows of time.

Additionally, one should measure the gamma activity in a region that is likely to be critically involved in processing of the presented stimulus such that neurons are likely to be driven at high rates. A more strongly driven E-I circuit is more likely to be accompanied by strong gamma oscillations.

Finally, as discussed above, an E-I circuit will not always generate synchronous gamma oscillations. Nevertheless, information processing based on precise synchronization in sparse cortical circuits may take place. One possible reason is that synchronized activity patterns do not always follow limit-cycle attractors, characteristic of regular oscillations, but instead more irregular, maybe even chaotic attractors. As a consequence of the more broadband nature of these processes, auto-correlograms often show the familiar, a few milliseconds wide center peak flanked by troughs but lack satellite peaks. Another possibility is that synchronous events could occur through synfire chain mechanisms, which do not require regularly repeating activation of neurons. Both, chaotic attractors and synfire chains represent internal mechanisms of synchronization (induced synchronization) as they do not require precise locking to

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

stimulus events. Finally, synchrony can also be evoked by an external input. For example, a flashed visual stimulus or an auditory "click" can trigger synchrony. In these cases the synchronous activity is locked in time to the stimulus and can be detected in a peri-stimulus time histogram (PSTH).

How Do Gamma Oscillations Interact with Lower Frequency Oscillations?

Generating gamma oscillations in an E-I network requires a fair amount of activity, which costs energy and would be difficult to sustain continuously. One way to bring that network to an oscillatory state is by driving these neurons externally, as discussed above. Here, it is important to synchronize gamma oscillations across multiple interacting modules. This can be a difficult task when the modules are far apart, where the transmission delays, combined with complex phase resetting curves, may make it difficult to generate synchrony. This raises the question: Are there other, transmission-delay independent ways to increase the long-range synchrony of gamma oscillations?

One possibility to make gamma oscillations more prominent is to suppress the lower frequency oscillations, thereby increasing the signal-to-noise ratio. The power spectra of neural activity show $\sim 1/f$ dependence on frequency f, with a large amount of power in low frequency signals. A small suppression of low frequency signals can significantly improve the relative contribution of the gamma power, making the gamma oscillations more effective in modulating spiking activity. This enhancement of gamma efficacy induced by low frequency suppression has been reported in several sensory cortical areas during attention and voluntary movements. Further, cortical activity is modulated by synchronous activity in the delta band (0.5–3 Hz) during quiet wakefulness and sleep, and this low frequency activity disappears during active engagement in a task and in conjunction with an increase in the gamma power.

Mechanisms also exist under which lower frequency oscillations can facilitate synchronization of gamma power fluctuations across large distances. Neurons in the hippocampus oscillate synchronously at theta frequency and project to a majority of neocortical areas. During the phase of theta oscillation, in which the hippocampus is more active, it can activate the target neocortical neurons, thereby enabling gamma frequency oscillations in those neocortical areas. The reverse happens at the phases of theta oscillations where hippocampal neurons are less active. Thus, the power of neocortical gamma oscillations would be modulated by the phase of hippocampal theta oscillations. This has been observed in behaving animals, including humans. Similarly, the phase of lower frequency oscillations modulates the power of neocortical gamma oscillations during slow-wave sleep, with higher gamma power appearing during the more depolarized phase of slow oscillations.

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

Thus, lower frequency oscillations can modulate gamma oscillations in two different ways. First, removal of the incoherent lower frequency oscillations can enhance gamma power. Second, the coherent lower frequency oscillations can facilitate synchronous modulation of gamma power across neocortical areas. Neuromodulators can also influence this process, in some cases removing low frequency oscillations and in others, facilitating them. This suggests that some form of optimization may occur, adjusting the balance between low and high frequency oscillations to maximize the efficacy of information processing. As a first step toward understanding how oscillations influence information processing, this interaction is discussed at the level of single neurons, of ensembles of neurons, and ensembles of neuronal networks.

How Do Oscillations Influence the Neural Representation at a Single Neuron Level?

Neurons respond in a graded fashion to sensory stimuli by altering their average activity levels: optimal stimuli evoke greater amount of activity than suboptimal stimuli. This is the classic rate code. For example, hippocampal place cells change their firing rates as a function of the spatial location of an animal, and neurons in the primary visual cortex change their mean firing rates as a function of the orientation of a stimulus. Further, hippocampal place cell activity is modulated by synchronized theta rhythm, and visual cortical activity is modulated by gamma rhythm. Thus the neural responses are modulated by two very different forces: by stimuli anchored in physical space and by internally generated oscillations. The former contain information about the external world, the latter about internal processing and timing. How do the stimulus-driven responses interact with synchronized oscillations? Would such interaction serve any purpose?

In the simplest scenario, the neuron will simply sum up the two inputs and generate a spike when this input exceeds a threshold. Thus, when the stimulus-based input is low, the neuron would spike at only that phase of oscillation when the oscillatory input is high so that the total input is sufficient to reach spike threshold. On the other hand, when the stimulus-evoked input is high, the neuron can spike even at the phase of oscillation when the oscillatory input is minimal. Thus, an interaction between the input and oscillation generates a phase code: When the inputs are strong, neurons will respond at every phase of oscillation; when the inputs are weak, neurons can respond only at the peak of oscillation when the oscillatory drive is maximal. At intermediate values of input, the outcome is a combination of phases. Thus, interaction between rate-coded inputs and synchronized oscillations would generate a phase-coded output.

Such phase-code and rate-phase, or rate-latency transformation has been observed in the hippocampus and is called phase precession (i.e., the phase of

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

theta oscillation at which place cells spike varies systematically as a function of the animal's position). For example, on linear tracks, a place cell fires spikes near the trough of theta oscillation as the animal enters the place field, and the theta phase of spike precesses to lower values as the animal traverses farther in the place field. Phase precession has been observed in several parts of the entorhinal–hippocampal circuit, along with correlated changes in firing rates. Recent studies have shown a mathematically similar phase code in the visual system where neurons spike at an earlier phase of gamma oscillation when driven maximally by the optimal stimulus, but at later phases of gamma oscillation when driven by suboptimal stimuli. Further, rate-latency transformation can be detected in the structure of spatiotemporal receptive fields of direction-selective visual cortical neurons when probed using randomly flashed bars. Similar measurements in other structures are likely to detect similar phase codes. Further, when slowly rotating oriented bars are presented to visual cortex, they may generate gamma-phase progression, and so may single bars passing over a series of direction selective receptive fields in area MT or V1.

These phase codes have several computational and functional advantages. First, they enable the postsynaptic neuron to decode the stimulus parameters by simply measuring the phase of the oscillation at which the presynaptic neuron spiked. This phase or latency is clearly defined by the period of the oscillations. This is in contrast to a rate code where one has to specify arbitrarily the interval of time over which spike count has to be averaged to obtain an estimate of a rate code. Second, stimulus-evoked activation of groups of neurons that represent stimuli in a sequence several seconds long would generate a compressed version of the stimulus sequence within an oscillation cycle due to the rate-phase transformation, possibly allowing these stimuli to be bound together and perceived as a chunk. Third, this temporally precise sequence of activation of neurons would facilitate the induction of spike timing-dependent plasticity, thereby generating a permanent record of the group of coactivated neurons in terms of the strengths of synapses connecting them. This would not only involve strengthening of synapses, but also weakening of synapses, especially the ones that correspond to nonsequential activation.

This mechanism of rate-phase transformation can thus be used to learn temporal sequences that occur over a timescale of a second, even though synaptic plasticity mechanisms operate on timescales of milliseconds. Similar learning of sequences may occur in other scenarios as well, where oscillations are imposed by other means. For example, systematic movements of the eyes across a natural scene every third of a second could induce oscillations where sequentially perceived views of the scene are brought together to form a stable, coherent percept using short- and long-term synaptic plasticity mechanisms. In addition to the relative timing of spikes between the stimulus-selective excitatory neurons, inhibitory spikes and neuromodulatory inputs are likely to determine the pattern of synaptic modifications.

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

While these mechanisms would work well for learning sequences of events that occur over a period of about a second, it remains to be determined how sequences of events that occur several minutes or hours apart can be learned via hitherto unknown mechanisms. Further, the above discussion of neural responses assumed that they are fixed and can be described in terms of a receptive field. Next we question this assumption and discuss the possibility of dynamic receptive fields.

Are Neural Representations Static or Dynamic?

In a typical study of neural information processing, the experimenter measures the changes in neural activity in response to a variety of stimuli. The neurons may respond strongly to one set of stimuli and less so to others. The pattern of neural responses to stimuli defines the neuron's receptive field.

The notion of a receptive field guides our thinking on how single neurons represent information but also has several limitations:

1. There are an infinite number of possible stimuli, varying across many dimensions. Hence, it is difficult to find the stimulus or set of stimuli that drive the neuron optimally within a finite amount of time.
2. Internal variables, such as arousal and neuromodulatory state, modulate neural responses.
3. Not only the stimuli within the receptive field, but even stimuli outside the classical receptive field modulate the responses.
4. The responses of many neurons in the visual system are affected by the attentional level and the reward value of the stimulus.
5. Most importantly, during natural behavior, stimuli are not static and do not appear in isolation. Instead, a large number of visual stimuli typically appear simultaneously and the stimulus configuration changes rapidly.

As a consequence of these five influences, the classical receptive field of a neuron can change dramatically between situations. For example, transient inactivation of the somatosensory cortex or of a sensory organ generates a large reorganization of the sensory map—a process that occurs within a second. In the hippocampus, past experience can result in a complete reorganization of the spatial selectivity of place cells. This reorganization is called remapping. Remapping can occur even on short timescales (~minutes), not just over days. In addition, when stimuli are presented in a sequence, visual cortical neurons not only respond to the onset of the stimulus but the responses depend on the sequential position of the stimulus as well: some neurons fire maximally to the presentation of the first stimulus in the sequence, irrespective of the identity of that stimulus. Finally, recent experiments show that hippocampal neurons fire

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

in a sequence even when the animal is sleeping or is running without changing its position (i.e., in a fixed running wheel).

These results suggest that neural responses are dynamic and can change rapidly with changes in stimulus configurations and internal variables, such as past experiences. This should not be surprising. As discussed earlier, to be energy efficient, neuronal codes need to be sparse and synchronous. Depending on the connectivity state of the network, recent history of neural and synaptic activity, and the nature of stimuli, different groups of neurons and synapses may become synchronously active and may hence drive different downstream neurons. Short-term dynamics of neurons and synapses can play an important role in generating such dynamic receptive fields.

This raises the question of how the downstream neurons interpret the messages sent by dynamically changing upstream neurons. Clearly, the postsynaptic neurons not only respond to just one presynaptic neuron but to an entire ensemble. Thus, dynamic reorganization of neural responses should be coordinated across an ensemble of neurons. The following discussion sketches an oscillation-based mechanism of dynamic coordination of neural codes across ensembles.

Neural Attractors, Cell Assemblies, Synaptic Assemblies, Oscillations, and Dynamic Coordination

Information is thought to be represented by the activity patterns of groups of neurons. Fault tolerant, stable, content addressable, and associative representation of stimuli across an ensemble of neurons can be implemented using the Hopfield attractor dynamics. For fixed point attractors, there are large energy barriers between the different representations, represented by local energy minima, whereas the energy required to make transitions between stimuli along some other dimension may be negligible in the case of continuous attractors. The stability and convergence of attractor dynamics are achieved through iterative processing of information in a recurrent network of excitatory neurons. Recent studies show that attractor dynamics can work even in sparsely active E-I networks. Such networks may show attractor dynamics with or without oscillations. It remains to be seen if the oscillations can facilitate the attractor dynamics.

How can the attractor dynamics generate flexible and dynamic neural responses? The answer may lie in efficient networks with short-term dynamics. As discussed above, energetically it is efficient for a small group of neurons to fire a few spikes synchronously to drive the postsynaptic neuron. Estimates show that in a period of about 20 ms, a sufficient period for the postsynaptic neuron to integrate the inputs and fire a spike, only about 500 neurons may need to be coactive out of a population of 300,000 CA1 neurons. Similarly sparse representations of stimuli are also present within neocortical circuits,

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

given the similarly low mean firing rates of principal cells in various neocortical layers. Experiments *in vitro* indicate that within this time window, pyramidal cells can linearly integrate the activity of hundreds of presynaptic inputs and then discharge. The minimum number of presynaptic neurons may be even an order of magnitude smaller if presynaptic neurons terminate on the same dendritic segment and discharge within a time window <20 ms. Further, in the hippocampus, inhibitory interneurons respond much more effectively than principal cells and hence, a single action potential of a presynaptic principal cell may be sufficient to discharge an interneuron. Release from potent inhibition would generate synchronous computation in a population of target excitatory neurons.

The efficacy of a handful of neurons in driving the downstream neuron will therefore depend on several parameters. For example, if the synapses from these neurons are located near each other on the downstream neuron's dendrite, synchronous activation of these synaptic inputs may cooperate to generate a dendritic spike. This would increase the effective strength of such groups of synapses, thereby altering the structure of the attractor and increasing the ability of a small number of input neurons to drive the downstream neuron. Similarly, the amount of synchrony, within a 20 ms window, between the activation of these synapses would strongly influence their effective strength in driving the downstream neuron. Recent studies show that although neural responses, as a function of input strength, are threshold-linear in an asynchronous condition, neural responses are sigmoidal in the synchronous condition: low synchrony results in no response, and above some threshold amount of synchrony the result is a maximal response. Thus, small changes in the input synchrony may activate different sets of neurons. This can be rapidly reorganized by recent history, which would influence the synaptic strength via short-term plasticity, resulting in convergence to different attractors. This dynamics could explain the rapid reorganization of hippocampal and somatosensory maps with past history or with small changes in stimuli. Further, synchronous inhibition in an E-I network would synchronously release excitatory neurons from inhibition during gamma oscillations, thereby allowing the neurons to change rapidly their response to inputs in a dynamic fashion. Finally, neuromodulators could alter the efficacy and timing of these synapses, which would result in dynamic reorganization of neurons responsiveness to stimuli based on internal variables.

Neuromodulators act broadly on neural circuits, and they are typically thought to act on slow timescales. The influence of neuromodulators can be focalized and accelerated by the following hypothesized extracellular mechanism: A region of the brain with higher activity could contain a larger amount of glutamate in the extracellular medium, and the clearing of the neuromodulators (e.g., through glial processes) may be altered by the level of glutamate. This may result in rapid changes in influence of neuromodulators on neural ensembles.

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

In this scenario of efficient and synchronous networks, activity would propagate rapidly across processing stages, and the relevant parameter would be the group of coactive neurons within a gamma cycle or the roughly 20 ms taken to activate the postsynaptic neurons. This group of temporarily synchronous neurons is referred to here as a cell assembly. Due to the mechanisms depicted above, the cell assembly can change rapidly with stimuli and internal variables. In other words, membership in a cell assembly is highly flexible. In a Hopfield network, for example, each gamma cycle may contain a cell assembly that primes the formation of another assembly in the next cycle of iteration toward convergence. In a recurrent network with short-term synaptic dynamics, this could lead to transitions between attractors and the generation of temporally sequential activity of ensembles of neurons.

Experiments show that the assembly of coactive hippocampal neurons, defined within a period of about 20 ms, can change rapidly. This is partly related to phase precession. As the rat walks through the environment, place cells fire a series of spikes at different phases of the theta cycle. The group of coactive cells within any 20 ms period depends not only on the phase of the theta rhythm and position of the animal but also on other variables (e.g., running speed, head direction). Thus, a multimodal, dynamic, and rapidly evolving representation emerges.

Two additional mechanisms by which cell assemblies can become dynamic are asynchronous background activity and top-down influence. These factors can raise the level of depolarization of the cell, thereby altering its responsiveness to short, efficient bursts of synchronous inputs. This is particularly effective when these inputs target the fast-spiking interneurons, which can then entrain a subset of pyramidal cells, and could explain the influence of top-down inputs on the rapid reorganization of neural responses.

In such scenarios, it is conceivable that the relevant parameter for describing the network dynamics is not the group of synchronously active cells, or cell assembly, but the group of synchronously active synapses or a synapse assembly through which the information flows. Neuromodulators and their receptors, located extra-synaptically, can directly modulate the activity pattern of the synapse assembly, which is not restricted to a single cell and which may or may not result in the modulation of the cell assembly. Theoretical studies are needed to determine how such a dynamic synaptic assembly can also be stable and noise tolerant. In addition, experimental studies are needed to determine the structure of synaptic assemblies.

These cellular and synaptic assemblies in the oscillating and balanced E-I networks are examples of dynamic equilibrium, in which a network is kept maximally responsive to changing patterns of inputs, while keeping energy expenditures low. Such self-organized systems are often characterized by a power-law spectrum of event amplitudes. Supporting evidence may be found in the neural systems in terms of the power-law-shaped spectra of the activity of ensembles, such as the LFP or EEG. However, these systems occasionally

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

produce very large events, which would be catastrophic for neural systems. Perhaps the strong, fast, and reliable feedback within an E-I circuit can help prevent such runaway events while keeping the system more responsive, near a critical point, through the generation of neural oscillations.

How Does the Brain Dynamically Select Cell Assemblies to Make a Decision?

Decisions are likely to occur by coordinated activity patterns in neural ensembles across different regions. These cell assemblies across regions can coordinate or compete in a number of ways to generate a winner assembly that drives the decision. First, the long-range excitatory-excitatory connections between the E-I assemblies in different regions can synchronize their activities and increase their firing rates, thereby making a selected group more active than others. Second, the excitatory-inhibitory connections between two different E-I assemblies could raise the level of inhibition and suppress the activity of the inhibited population. Third, the long-range inhibitory-inhibitory connections could either decrease or increase the firing rates of an E-I circuit. It would seem that this inhibitory transmission between cell assemblies may result in reduced activity of the second assembly. However, it has been shown that in recurrent E-I networks, under some parameter regimes, increased inhibitory inputs result in increased overall activity due to suppression of inhibition via recurrent inhibitory synapses. Thus, inhibitory connection between two cell assemblies may serve a dual purpose of (a) synchronizing their gamma rhythmic activity and (b) increasing their overall firing rates, thereby allowing this group of cell assemblies to drive the downstream group of neurons toward decision.

In addition to these mean firing rate-based effects and mechanisms, precisely timed inhibitory inputs could synchronously release distant cell assemblies from inhibition, thereby making them coactive in brief windows of time. Here, oscillations could facilitate this rapid synchrony and competition by synchronously activating and inactivating large neural ensembles across multiple brain areas. Thus, the decision-making process may be a phase-dependent rather than a rate-dependent process.

Neuromodulators could play a key role in these processes by altering the E-I balance and rhythms, thereby generating state-dependent synchrony of cell assemblies that determine the winner ensemble. The above mechanisms address direct competition between assemblies. An advantage is energy efficiency. However, it is possible that competition between assemblies occurs at the level of their efficacy in driving a downstream structure, such as the prefrontal cortex, resulting in competition between synaptic assemblies, which in turn can bias information processing in the upstream network. Such a recurrent process of decision making across networks could be slow, but can be speeded up through the use of a phase code, where the top-down and bottom-up influences

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.

can arrive at different phases of oscillations, thereby determining a winner ensemble within an oscillation cycle.

Conclusions

Our discussion emphasizes a few key mechanisms of neural information processing. At the heart of this discussion is the ubiquitous neural circuit of recurrently connected groups of excitatory and inhibitory neurons. The E-I circuit can remain at a dynamic equilibrium, allowing it to respond rapidly to inputs. The E-I module can be easily replicated to generate larger circuits, perhaps during evolution, with each component using a similar language. Further, the dynamic equilibrium would allow a small number of inputs to alter the state of the network and make neurons respond. Thus, the network could have sparse activity, thereby making it energy efficient. In such sparsely active E-I networks, synchronous activity would be transmitted efficiently and rapidly.

Under many conditions, this E-I circuit oscillates. Interaction between these oscillations and inputs from external stimuli would generate a phase-coded representation of input that is rapid, efficient, and malleable; one that can facilitate learning via mechanisms of spike time-dependent synaptic plasticity. Such phase codes could bind multiple neural representations for brief periods in a flexible fashion and determine the computational interactions between different signals. The duration of each syllable in an E-I network phase code, called a cell assembly, would be about 20 ms, corresponding to a gamma cycle. Cell assemblies across multiple gamma cycles can either converge to a Hopfield-like attractor, in the presence of stationary stimuli, or generate history-dependent responses with dynamic stimuli; alternatively, in the absence of stimuli, it can spontaneously transition across a sequence of cell assemblies due to short-term dynamics. The oscillations of E-I circuit can be synchronized across different regions, allowing dynamic coordination of phase codes across brain regions. Other processes, such as lower frequency oscillations and neuromodulators, can influence coordination and competition between the E-I assemblies, generating a state-dependent winning ensemble or decision. While some tantalizing support is available for these mechanisms of the emergence of mind from neurons, much remains to be theoretically understood and experimentally tested.

Chapter from "Dynamic Coordination in the Brain: From Neurons to Mind," edited by C. von der Malsburg, W. A. Phillips, and W. Singer. Strüngmann Forum Report, vol. 5. ISBN 978-0-262-01471-7. Cambridge, MA: The MIT Press.