of the limb, simplifying the process of control. However, there is now convincing evidence that in programming motor commands to the muscles, the brain does take into account the dynamics of the task (see SENSORIMOTOR LEARNING). The motor commands result in an attractor trajectory that leads the hand in simple reaching movements. It would be expected that both the trajectory of the attractor and the shape of the restoring field about the attractor would change as the dynamics of the task change. Although it is quite possible that biological motor commands can be described in terms of changes in the equilibrium position of the limb, the hypothesis that control of movements by the brain is explicitly performed through this manipulation because it somehow simplifies the process of control appears to be inconsistent with the current data.

Road Maps: Mammalian Motor Control; Dynamic Systems

Related Reading: Arm and Hand Movement Control; Cerebellum and Motor Control; Geometrical Principles in Motor Control; Limb Geometry, Neural Control; Muscle Models; Optimization Principles in Motor Control

References

- Bennett, D. J., Hollerbach, J. M., Xu, Y., and Hunter, I. W., 1992, Time varying stiffness of the human elbow joint during cyclic voluntary movement, *Exp. Brain Res.*, 88:433–442.
- Burdet, E., Osu, R., Franklin, D. W., Milner, T. E., and Kawato, M., 2001, The central nervous system stabilizes unstable dynamics by learning optimal impedance, *Nature*, 414:446–449.
- Feldman, A. G., 1966, Functional tuning of the nervous system with control of movement or maintenance of a steady posture: II. Controllable parameters of the muscles, *Biophysics*, 11:565–578. ◆

- Feldman, A. G., and Orlovsky, G. N., 1972, The influence of different descending system on the tonic stretch reflex in the cat, *Exp. Neurol.*, 37:481–494.
- Gomi, H., and Kawato, M., 1996, Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement, *Science*, 272:117–120. ◆
- Gribble, P. L., Ostry, D. J., Sanguineti, V., and LaBoissiere, R., 1998, Are complex control signals required for human arm movement? J. Neurophysiol., 79:1409–1424.
- Hodgson, A. J., and Hogan, N., 2000, A model-independent definition of attractor behavior applicable to interactive tasks, *IEEE Trans. Syst. Man Cybern.*, 30:105–118. ◆
- Hoffer, J. A., and Andreassen, S., 1981, Regulation of soleus muscle stiffness in premammillary cats: Intrinsic and reflex components, J. Neurophysiol., 45:267–285.
- Hogan, N., 1985, The mechanics of multi-joint posture and movement control, *Biol. Cybern.*, 52:315–331.
- Katayama, M., and Kawato, M., 1993, Virtual trajectory and stiffness ellipse during multijoint arm movements predicted by neural inverse models, *Biol. Cybern.*, 68:353–362.
- Milner, T. E., 1993, Dependence of elbow viscoelastic behavior on speed and loading in voluntary movements, *Exp. Brain Res.*, 93:177–180.
- Mussa-Ivaldi, F. A., Hogan, N., and Bizzi, E., 1985, Neural, mechanical, and geometric factors subserving arm posture, *J. Neurosci.*, 5:2732–2743. ◆
- Shadmehr, R., and Arbib, M. A., 1992, A mathematical analysis of the force-stiffness characteristics of muscles in control of a single joint system, *Biol. Cybern.*, 66:463–477.
- Shadmehr, R., Mussa-Ivaldi, F. A., and Bizzi, E., 1993, Postural force fields of the human arm and their role in generating multi-joint movements, *J. Neurosci.*, 13:45–62.
- Wang, T., Dordevic, G. S., and Shadmehr, R., 2001, Learning dynamics of reaching movements results in the modification of arm impedance and long-latency perturbation responses, *Biol. Cybern.*, 85:437–448.
- Won, J., and Hogan, N., 1995, Stability properties of human reaching movements, *Exp. Brain Res.*, 107:125–136.

Event-Related Potentials

Steven L. Bressler

Introduction

It is commonly believed that cognition intimately depends on the functioning of the cerebral cortex. Understanding the neural basis of cognition therefore will likely require knowledge of cortical operations at all organizational levels, which may usefully be grouped as microscopic, mesoscopic, and macroscopic. The cellular mechanisms of cortical neurons operate at the microscopic scale and are measured by a host of techniques targeted at that level. Individual cortical neurons contribute to cognitive function, however, by joining in the cooperative actions of neural networks, which operate at the mesoscopic and macroscopic scales. At the microscopic scale, the cooperative fraction of any single neuron's total activity may be exceedingly small, but the cooperative activity of the network exerts effects that are relevant for cognition. The mesoscopic level concerns the cooperative activity of neurons locally in ensembles and area networks, and the macroscopic level concerns the cooperative activity of neurons globally in large-scale networks and entire systems. Thus, many important cortical functions reside in the operations of neural networks and are measured by specialized techniques targeted at the mesoscopic and macroscopic levels.

The event-related potential (ERP) is a neural signal that reflects coordinated neural network activity. The cortical ERP provides a window onto the dynamics of network activity in relation to a variety of different cognitive processes at both mesoscopic and macroscopic levels on a time scale comparable to that of single-neuron activity. Cortical ERPs arise from synchronous interactions among large numbers of participating neurons. These include dense local interactions involving excitatory pyramidal neurons and inhibitory interneurons, as well as long-range interactions mediated by axonal pathways in the white matter. (See NEUROANATOMY IN A COM-PUTATIONAL PERSPECTIVE.) Multiple feedback loops involving both excitatory and inhibitory interactions typically cause ERPs to be oscillatory, meaning that they fluctuate within bounds around a central value. Depending on the types of interaction that occur in a specific behavioral condition, cortical networks may display different states of synchrony, causing their ERPs to oscillate in different frequency bands, designated delta (0–4 Hz), theta (5–8 Hz), alpha (9–12 Hz), beta (13–30 Hz), and gamma (31–100 Hz).

The physiological basis of the cortical ERP lies in fields of potential generated by interacting neurons (Lopes da Silva, 1991). Field potentials are largely dendritic in origin, resulting from the summed extracellular currents generated by electromotive forces (EMFs) in the dendrites of synchronously active cortical neurons, primarily pyramidal cells. The EMFs, arising from synaptic activation of postsynaptic ion channels, circulate current in closed loops across the cell membrane and through the intracellular and extracellular spaces. Summed closed-loop currents generated by an ensemble of neighboring neurons flow across the external resistance to form the local ensemble mean field potential (Freeman, 2000).

Depending on the location and size of the recording and reference electrodes, recorded cortical field potentials integrate neural activity over a range of spatial scales: from the intracortical local field potential (LFP) to the intracranial electrocorticogram (ECoG) to the extracranial electroencephalogram (EEG). The LFP (Figure 1) is the most spatially localized signal, integrating the field potential on a submillimeter scale; the ECoG integrates on a submillimeter to millimeter scale; and the EEG integrates over centimeters. The term "field potential" will be used here in reference to the general class of signal subsuming the LFP, ECoG, and EEG. (The intracellular components of the same closed-loop currents that give rise to field potentials are responsible for the closely related magnetic fields, recorded extracranially as the magnetoencephalogram, or MEG.)

A general problem in the investigation of ERPs is that field potential recordings most often contain a combination of potentials, in unknown proportions, from multiple sources. Thus, in addition to the ERP, which is derived from specific networks associated with a behavioral event, the field potential typically also contains potentials derived from the more general field activity of large neural populations. Owing to their fortuitous geometric arrangements and synchronous behavior, these later potentials are mixed with the ERP waveform. Thus, a primary task of all ERP studies is to extract the event-related portion of the recorded field potential. The next section deals with some basic methodology by which this is accomplished for different kinds of ERP.

ERP Varieties and Their Analysis

Whether reflecting mesoscopic or macroscopic activity, the cortical ERP is an electrical signal generated by neuronal networks in relation to a behaviorally significant event. (The corresponding event-related magnetic field has many of the same dynamic and functional properties as the ERP.) Two general classes of ERP are distinguished by whether the relevant event is discrete or continuous. In the case of discrete events, the associated transient ERP is analyzed in short epochs that are time-locked to the event. In the case of continuous events, which usually are periodically modulated sensory stimuli such as a visual flicker, the concurrent steadystate ERP is analyzed in a relatively long time segment.

The traditional approach to the analysis of transient ERPs is to consider the ERP as a characteristic waveform that occurs in relation to the behaviorally significant discrete event. As a simplifying assumption, the ERP waveform is usually treated as if it possesses the same amplitude and phase each time that the event is repeated on multiple trials, although recent analysis shows that this assumption may not always be valid (Truccolo et al., 2002). Nonetheless, as was discussed above, the recorded single-trial field potential contains contributions from network activity that are both associated (ERP signal) and not associated (noise) with the event. Therefore, averaging of the single-trial field potential time series, time-locked to the event, is commonly employed to extract the ERP from the non-event-related noise. When the relevant event is a sensory stimulus, such phase-locked ERPs are called "evoked." Averaged evoked potentials (Figure 2) are most commonly described in terms of the succession of waveform components that follow stimulus presentation. These components are typically identified according to their polarity (positive or negative) and their time latency following stimulus onset. (Note that the time latency is equivalent to phase in this context.)

Transient ERP waveform components having variable phase may also reliably occur in relation to the repeated event. In this case, time series averaging does not reveal the ERP but instead is destructive, since components of opposite polarity on successive trials tend to be canceled. Non-phase-locked ERPs are referred to as "induced" when they occur following a stimulus and "spontaneous" in the period prior to a stimulus or motor response. This type of ERP may be effectively analyzed by averaging the frequency content of single-trial time series rather than the time series themselves.

Non-phase-locked transient event-related phenomena are detected as frequency-specific changes in the ERP time series. These phenomena may consist of either an event-related increase or decrease of power in one or more of the aforementioned frequency bands. Since the level of ERP power is typically considered to reflect the degree of synchrony within local neuronal populations. a power increase is called event-related synchronization, and a power decrease is called event-related desynchronization (Pfurtscheller and Lopez da Silva, 1999). Frequency analysis has the further advantage of allowing measurement of event-related phase synchronization of ERPs from different cortical sites (Varela et al., 2001). ERP phase synchronization in different frequency ranges has been identified as a fundamental neural correlate of basic sensory and motor processes, as well as higher cognitive processes such as perception and recall of semantic entities. (See SYNCHRO-NIZATION, BINDING AND EXPECTANCY.)

The study of steady-state ERPs also depends on a variant of frequency analysis. Field potentials recorded during periodically modulated sensory stimulation are narrow-bandpass filtered around the frequency of the driving periodicity to derive the steady-state (periodic) ERPs. Variations in the amplitude and phase of the steady-state ERP are interpreted in terms of driving frequency, spatial location, and behavioral state.



Figure 1. A local field potential (LFP) recorded from the posterior parietal cortex of a macaque monkey in relation to a visual stimulus presented on a display screen for 100 ms, starting at time 0. The LFP was recorded from a chronically implanted bipolar transcortical electrode consisting of 51-µm-diameter Teflon-coated platinum wires with 2.5-mm tip separation.



Figure 2. The averaged event-related potential from the same posterior parietal cortex site as in Figure 1, computed from an ensemble of 888 trials. Note the flat prestimulus baseline as compared to the single trial in Figure 1. This illustrates the fact that rhythmic prestimulus activity that is not phase-locked to the stimulus is canceled out by the averaging process.

The Theory of Large-Scale Cortical Networks

Evidence from a variety of sources indicates that neural networks in the cerebral cortex are organized both locally in anatomically segregated areas and on a large scale encompassing multiple distributed areas (Bressler, 2002). Although research on cortical network properties is still in its infancy, a rough depiction of some basic operational features is now possible. Local-area networks process and store information related to specialized sensory, motor. and executive functions, and local synaptic interactions lead to the manifestation of coherent spatial ERP patterns in these specialized informational domains. These interactions also modify the local synaptic matrix with learning. The modified synaptic matrix exerts an essential control on pattern formation in the local-area network by attracting its dynamics to learned (attractor) patterns. In this regard, artificial neural networks that operate according to attractor dynamics bear a resemblance to cortical networks at the local level. (See COMPUTING WITH ATTRACTORS.)

An essential element of overall cortical network function, however, is missing from most artificial network models. Following training on pattern recognition problems, traditional artificial neural networks converge to fixed solutions for a given class of input patterns. Although this behavior has well-known advantages for pattern recognition, it represents an excessive processing rigidity, since these networks lack the ability to adapt to changing external constraints such as are found in real-world situations. Adaptability, in this sense, is a distinguishing feature of normal cortical function.

Theoretical considerations suggest that processing adaptability in the cerebral cortex derives from an essential property of largescale network dynamics called *metastability*. Cortical metastability refers to a state of dynamic balance among multiply interacting local networks in which the tendency for independent local expression is offset by the tendency for large-scale entrainment (Bressler and Kelso, 2001). The property of metastability permits local networks that are interconnected within the large-scale network architecture of the cortex to coordinate their activities without becoming locked in a fixed pattern of coordination from which they cannot escape.

The ability of local-area networks to form transient coordination relations may represent a basic cortical mechanism for the rapid and flexible association of information from different informational domains. (See ASSOCIATIVE NETWORKS.) It is to be expected that the concurrent coordination of multiple local-area networks imposes conjoint constraints on the spatiotemporal patterning of activity in each local network. The imposition of such constraints may have the important effect of creating associations between activity patterns in different informational domains during the learning process, through the modification of synapses of axons that project from one local network to another. These learned associations would then act during recall on the attractor dynamics of multiple interacting local area networks, causing them to reach a conjunction of consensual patterns that represents an integration of their information. (See COMPUTING WITH ATTRACTORS.)

ERP Evidence for Large-Scale Cortical Network Organization

The theoretical considerations presented in the previous section lead to predictions about the large-scale cortical network organization underlying cognition. One straightforward prediction is that cognitive states should be characterized by unique configurations of interdependent cortical areas in large-scale networks. A confirmation of this prediction is found in the spatial patterning of coactivated cortical areas seen with functional brain imaging techniques such as PET and fMRI. (See COVARIANCE STRUCTURAL EQUATION MODELING.) Like these neuroimaging procedures, ERPs can provide information about the spatial distribution of large-scale network activity underlying a cognitive function. Moreover, because ERPs reflect neurodynamics on a fast time scale (that is inaccessible to current brain imaging technologies), ERPs can also reveal elementary neural subprocesses that subserve that cognitive function. This section uses working memory to illustrate how ERP results can relate large-scale network activity to different subprocesses of a cognitive function.

Working memory consists of several subprocesses for which prominent averaged ERP waveform components have revealed distinct underlying large-scale networks (McEvoy et al., 1998). The mismatch negativity is an early poststimulus ERP component that reflects the maintenance of sensory working memory in the auditory modality. It is elicited by auditory stimuli having physical acoustic properties that deviate from prior (standard) stimuli registered in auditory memory. Occurring between 80 and 200 ms after presentation of deviant auditory stimuli, thus overlapping the N1 and P2 components, the mismatch negativity is isolated by computing the difference wave between averaged ERPs evoked by deviant and standard stimuli. The mismatch negativity is subserved by a large-scale network that includes, in addition to auditory cortical areas, dorsolateral prefrontal cortex, which may serve to control the maintenance of sensory memory in the auditory cortex following one stimulus for comparison with subsequent stimuli (Alain et al., 1998).

A second ERP component, the P3b, occurring roughly 300 ms poststimulus, also results from the comparison of target stimuli with the content of working memory. However, rather than being tuned to the physical characteristics of stimuli, the widely distributed cortical network underlying the P3b is involved in the categorization of stimuli as significant events. Network strength has been found to reflect the degree of consonance resulting from comparison of stimulus attributes with a maintained "expectation" (Kok, 2001).

A third ERP component, related to semantic memory, is the negative-going N400. It occurs between 200 and 500 ms after presentation of a potentially meaningful information-bearing stimulus and varies systematically according to the preexisting context that is established by semantic and long-term memory influences. Specifically, N400 amplitude is reduced as a function of associative, semantic, and repetition priming within or across sensory modalities (Kutas and Federmeier, 2000). Variation of its scalp-recorded topographic distribution with task and stimulus type suggests that the N400 reflects the construction of meaning by cross-modal interactions in a widely distributed neural network. This view is supported by intracranial evidence that the N400 arises from similar waves of activity in multiple brain areas, particularly in the temporal and prefrontal cortices, during the retrieval of information from semantic memory.

Deeper insight into the dynamic organization of large-scale networks underlying working memory comes from studies of the phase synchronization between ERPs from distributed cortical areas. For example, long-range ERP phase synchronization has been reported in the theta frequency range between prefrontal and posterior association areas when subjects retain verbal and spatial items for short periods of time (Sarnthein et al., 1998) and in the beta frequency range between extrastriate areas when they retain visual object representations (Tallon-Baudry et al., 2001). These studies suggest that large-scale cortical network function is based not just on the co-activation of distributed neuronal ensembles, but also on the active coordination of ensemble activity, observable as ERP phase synchronization.

Finally, other ERP types have been used to examine the neural correlates of working memory load. In one investigation, the

steady-state visual ERP elicited by a diffuse 13-Hz visual flicker was used to study memory load during the retention period of an object working memory task (Silberstein et al., 2001). The steadystate visual ERP exhibited a load-dependent increase in amplitude at frontal and occipitoparietal sites. By comparison, in a study of event-related synchronization and desynchronization, significant effects of memory load were found in the frontal lobe during a visual sequential letter task (Krause et al., 2000). Event-related synchronization was found at theta frequencies during the initial stages of stimulus processing, whereas event-related desynchronization was observed at alpha frequencies.

Discussion

The cortical ERP reflects the coordinated behavior of large numbers of neurons in relation to a meaningful externally or internally generated event. Single neurons are actively coordinated in the operations of ensembles, local-area networks, and large-scale networks. ERP studies provide a unique avenue of approach to the dynamics of coordination in the cortex at the mesoscopic and macroscopic levels of organization. ERP analysis is an indispensable complement to single-cell neurophysiology and whole-head neuroimaging techniques and can supply a rich source of criteria for neural network modeling efforts.

ERP studies have shown that local cortical area networks are able to synchronize and desynchronize their activity rapidly with changes in cognitive state. These synchronization changes occur between neurons located both within individual local networks and in different local networks. The ability of local area networks to repeatedly reconfigure their activity patterns under constraint of large-scale coordinating influences may allow them to increase the degree of consensus of those local patterns in a short period of time, thereby causing the cortical system as a whole to evolve toward the solution of computational problems. Since it normally operates in a metastable dynamic regime, the cortex is able to balance the coordinated and independent behavior of local networks to maintain the flexibility of this process. When incorporated into artificial neural network designs, a similar computational process could prove useful in avoiding the processing rigidity of many current network models. A metastable large-scale neural network design that recruits and excludes subnetworks according to their ability to reach consensual local patterns has the potential to implement behavioral schema and adapt to changing environmental conditions. Such a system would represent an important advance in machine cognition.

Road Map: Cognitive Neuroscience

Related Reading: Covariance Structural Equation Modeling; EEG and MEG Analysis; Hippocampal Rhythm Generation; Schema Theory; Synchronization, Binding and Expectancy

References

- Alain, C., Woods, D. L., and Knight, R. T., 1998, A distributed cortical network for auditory sensory memory in humans, *Brain Res.*, 812:23– 37.
- Bressler, S. L., 2002, Understanding cognition through large-scale cortical networks, *Curr. Dir. Psychol. Sci.*, 11:58–61.
- Bressler, S. L., and Kelso, J. A., 2001, Cortical coordination dynamics and cognition, *Trends Cogn. Sci.*, 5:26–36.
- Freeman, W. J., 2000, Mesoscopic neurodynamics: From neuron to brain, J. Physiol. Paris, 94:303–322. ◆
- Kok, A., 2001, On the utility of P3 amplitude as a measure of processing capacity, *Psychophys.*, 38:557–577.
- Krause, C. M., Sillanmaki, L., Koivisto, M., Saarela, C., Haggqvist, A., Laine, M., and Hamalainen, H., 2000, The effects of memory load on event-related EEG desynchronization and synchronization, *Clin. Neurophysiol.*, 111:2071–2078.
- Kutas, M., and Federmeier, K. D., 2000, Electrophysiology reveals semantic memory use in language comprehension, *Trends Cogn. Sci.*, 4:463– 470.
- Lopes da Silva, F., 1991, Neural mechanisms underlying brain waves: From neural membranes to networks, *Electroenceph. Clin. Neurophysiol.*, 79:81–93. ◆
- McEvoy, L. K., Smith, M. E., and Gevins, A., 1998, Dynamic cortical networks of verbal and spatial working memory: Effects of memory load and task practice, *Cereb. Cortex*, 8:563–574.
- Pfurtscheller, G., and Lopez da Silva, F. H., 1999, Event-related EEG/MEG synchronization and desynchronization: Basic principles, *Clin. Neurophysiol.*, 110:1842–1857.
- Sarnthein, J., Petsche, H., Rappelsberger, P., Shaw, G. L., and von Stein, A., 1998, Synchronization between prefrontal and posterior association cortex during human working memory, *Proc. Natl. Acad. Sci. USA*, 95:7092–7096.
- Silberstein, R. B., Nunez, Pipingas, A., Harris, P., and Danieli, F., 2001, Steady state visually evoked potential (SSVEP) topography in a graded working memory task, *Int. J. Psychophysiol.*, 42:219–232.
- Tallon-Baudry, C., Bertrand, O., and Fischer, C., 2001, Oscillatory synchrony between human extrastriate areas during visual short-term memory maintenance, J. Neurosci., 21:RC177.
- Truccolo, W. A., Ding, M., Knuth, K. H., Nakamura, R., and Bressler, A., 2002, Trial-to-trial variability of cortical evoked responses: Implications for the analysis of functional connectivity, *Clin. Neurophysiol.*, 113:206– 226.
- Varela, F., Lachaux, J. P., Rodriguez, E., and Martinerie, J., 2001, The brainweb: Phase synchronization and large-scale integration, *Nat. Rev. Neurosci.*, 2:229–239. ◆

Evolution and Learning in Neural Networks

Stefano Nolfi

Introduction

Evolution and learning are two forms of adaptation that operate on different time scales. Evolution is capable of capturing relatively slow environmental changes that might encompass several generations. Learning allows an individual to adapt to environmental changes that are unpredictable at the generational level. Moreover, evolution operates on the genotype, but learning affects the phenotype and phenotypic changes cannot directly modify the genotype. Recently, the study of artificial neural networks subjected both to an evolutionary (see EVOLUTION OF ARTIFICIAL NEURAL NETWORKS) and a lifetime learning process has received increasing attention. These studies (see also Nolfi and Floreano, 1999) have been conducted with two different purposes: (1) looking at the advantages, in terms of performance, of combining two different adaptation techniques; (2) understanding the role of the interaction between learning and evolution in natural organisms. The general picture emerging from this body of research suggests that, within an evolutionary perspective, learning has several different adaptive functions: