

The Endogenously Active Brain: The Need for an Alternative Cognitive Architecture

William Bechtel

Department of Philosophy, Center for Chronobiology, and
Interdisciplinary Program in Cognitive Science,
University of California, San Diego (USA)

Résumé : La plupart des projets d'architectures cognitives dans les sciences cognitives et des explications des processus cérébraux dans les neurosciences interprètent l'esprit/cerveau comme réactif : le processus est déclenché par un stimulus et se termine par une réponse. Mais il y a de plus en plus de données laissant penser que le cerveau est endogéniquement actif : des oscillations de l'activité électrochimique à de multiples fréquences ont lieu dans le cerveau même en l'absence de stimuli et ceux-ci servent plutôt à moduler ces oscillations qu'à déclencher l'activité. De plus, des données montrent que cette activité endogène est utilisée dans diverses activités de traitement de l'information. Je m'appuie sur les données issues des enregistrements cellulaires unitaires, de l'EEG et de l'IRMf en situation passive, afin de soutenir la thèse d'un comportement oscillatoire continu dans le cerveau et d'identifier plusieurs manières dont ce comportement peut contribuer à la cognition. Si les sciences cognitives ont pour but de comprendre comment nous effectuons des tâches cognitives, elles ont besoin de développer des architectures cognitives qui intègrent le type d'activité endogène dynamique présentée par le cerveau.

Abstract: Most proposals of cognitive architectures in cognitive science and accounts of brain processes in neuroscience construe the mind/brain as reactive: processing is initiated by a stimulus and terminates in a response to it. But there is growing evidence that brains are endogenously active: oscillations in electrochemical activity at multiple frequencies are ongoing in the brain even in the absence of stimuli and stimuli serve to modulate these oscillations rather than initiate activity. Moreover, evidence is growing that this endogenous activity is used in various information processing activities. I appeal to evidence from single-cell recording, EEG, and resting state fMRI to support the claim of ongoing oscillatory behavior in the brain and identify several ways it may contribute to cognition. If cognitive science is to understand how we perform

cognitive tasks it needs to develop cognitive architectures that incorporate the sort of endogenous dynamic activity exhibited by the brain.

1 The search for cognitive architectures

The cognitive tradition is distinguished from its behaviorist predecessor by focusing on information processing mechanisms that are thought to exist within the mind-brain and hypothesized to explain behavior. Cognitive researchers have often assumed, either implicitly or explicitly, that the brain takes in information through the senses, represents it, performs operations on the representations, and responds by either changing its internal state or planning and executing actions. In the early decades of cognitive science it was difficult to identify the neural processes that served as representations for high-level cognitive processes such as memory, reasoning, and problem solving and the neural operations through which they were processed. As a result, most cognitive theories of information processing had to rely on indirect measures such as reaction times and error patterns to guide and evaluate hypotheses as to the representations employed and the operations performed on them. For example, Sternberg used the time required for subjects to determine whether a test item was on a list of items they had memorized to determine that humans perform exhaustive serial search [Sternberg 1966].

One way cognitive scientists, especially those focused on computational modeling of cognition, further constrained their inquiry was to develop proposals as to the nature of the cognitive architecture. A cognitive architecture specifies the primitive operations the mind/brain is thought to perform. Proposed architectures such as Newell's SOAR [Laird, Newell & Rosenbloom 1987] and Anderson's ACT-R [Anderson 1990, 2007] were defended on theoretical grounds that they possessed the appropriate primitive capacities for performing the cognitive activities performed by humans. For those skeptical of the proposal that the mind performs operations on symbol structures, the parallel distributed processing framework or connectionism provided an alternative class of architectures [Rumelhart & McClelland 1986], emphasizing in particular feedforward networks trained by backpropagation. (For an overview and taxonomy of cognitive architectures, see [Duch, Oentaryo & Pasquier 2008].) Researchers working with a particular architecture constrain themselves to account for various aspects of human cognitive activity using the primitive operations provided in the architecture. None of these architectures was originally developed primarily with an eye to characterizing the representations and processes used in the brain (although connectionists sometimes refer to their approach as *brain-style computing*), but as techniques have been developed to relate cognitive processing to neural activity, advocates of some cognitive architectures have tried to show that their architecture fits with our understanding of how the brain functions.

These cognitive architectures typically adopt a reactive perspective on the mind/brain. Cognitive activity is assumed to begin with the presentation of a task or stimulus, which is represented and the representation is then transformed via operations specified by the architecture. This reactive conception of cognition (it occurs in response to a stimulus) has also been shared as the neurosciences began to provide insight into the representations and operations performed. The first successes in identifying neural processes that represent information resulted from investigations of sensory and motor processing in which it was possible to link brain activity (typically spiking rates of neurons) with sensory stimuli or motor activities. With respect to visual processing, for example, researchers beginning with Kuffler [Kuffler 1953] and Hubel and Wiesel [Hubel & Wiesel 1962, 1968] employed such techniques as single- and multi-cell recording to determine what features of visual stimuli were correlated with specific neuronal activity. The activity of these neurons was then viewed as representing the correlated features of the visual stimulus, and researchers hypothesized operations through which these representations were successively transformed in a hierarchy of processing areas ([van Essen & Gallant 1994]; for an analysis of this history, see [Bechtel 2008]). In the 1970s researchers developed techniques for relating the electrical signal recorded at the cortical surface (electroencephalography or EEG) to stimulus presentation (measuring what are known as Evoked Response Potentials or ERPs). While ERP studies provided little information about the spatial locus of activity, they offered information about the temporal order of processing. For example, Kutas and Hillyard found that when the last word of a sentence was anomalous the EEG exhibited a negative deflection peaking about 400 milliseconds later (hence, an N400 response) [Kutas & Hillyard 1980]. The introduction of Positron Emission Tomography (PET) in the late 1980s and functional Magnetic Resonance Imaging (fMRI) in the 1990s allowed researchers to begin to locate cognitive operations in brain regions in humans. Petersen, Fox, Posner, Mintun, & Raichle, for example, adapted the subtractive method, developed by Donders in the mid-19th century for identifying the time required for a cognitive operation, to identify the brain regions responsible for such operations [Petersen, Fox, Posner *et al.* 1988]. Thus, in their pioneering PET study they subtracted the blood flow measured when subjects simply read a noun aloud from that measured when they first generated a related verb and read it aloud. They found increased activity in the left prefrontal cortex (also the cerebellum and the anterior cingulate), which they argued to be the locus of the semantic processing required for the task.

Researchers adopting the reactive perspective have provided a great deal of information about information processing in the brain, especially in areas of sensory and motor processing and increasingly with respect to memory, attention, and emotional responses. But there are reasons to be skeptical about the adequacy of the reactive perspective. Since the pioneering studies of Lorente de Nó [Lorente de Nó 1938], researchers have recognized that there are at least as many, and likely many more, backwards and collateral projections than for-

ward ones.¹ The reactive perspective, however, has been able to provide little insight into what contribution to information processing these operations perform since whatever contribution they make has already affected how brain areas respond to stimuli and typically cannot be separately identified. (One exception has been the investigation of how attention modifies processing of perceptual stimuli. See [Corbetta & Shulman 2002].)

Although it has received far less attention, there is an alternative tradition in brain research that has emphasized the endogenous activity of the brain, viewing stimuli as perturbing on-going brain activity, not initiating activity. This tradition traces back at least to Thomas Graham Brown, who studied neural mechanisms for motor behavior in decerebrate cats in the Liverpool laboratory of Charles Scott Sherrington, a chief proponent of adopting a reactive perspective on the brain [Brown 1911, 1914]. Both were interested in motor activity such as walking which appear to originate from inside the organism. Sherrington extended the reactive framework to explain such behavior in terms of a sequential reflex mechanism, by which peripheral input (e.g., to the cat's feet when placed on a moving treadmill) produced a sequence of neural signals (to the spine, within the spine, and out to flexor and then extensor muscles) [Sherrington 1923]. Each cycle of stepping resulted in renewed input (sensory feedback) and hence ongoing, rhythmic stepping movements. Brown discovered that he could obtain similar rhythmic stepping even after isolating the spinal cord from afferent (peripheral) input by cutting the dorsal root nerves. He accounted for its rhythmic outputs by proposing a coupled network of spinal neurons—one for flexion and one for extension—that each inhibited the other's behavior to generate oscillations. Brown's proposals anticipated research on what came to be known as the central pattern generator [Wilson & Wyman 1965]. More recently, Brown's perspective on endogenous activity has been advanced by researchers using the very tools employed by the reactive tradition—single-cell recording, EGG, and fMRI.

In this paper I will argue that the evidence for the endogenously active perspective on the mind-brain is extremely compelling and that in light of it cognitive researchers should fundamentally reconceive their conceptions of cognition and cognitive architectures to incorporate and recognize the significance of endogenous activity. I begin by considering what it means to claim that a biological mechanism exhibits endogenous activity and focus on how individual neurons are endogenously active, then explore how EEG and fMRI are revealing synchronous activity of populations of neurons and inspiring proposals as to how they figure in neural processing. I then conclude by emphasizing how this body of research points to the need for fundamental revisions in the cognitive architectures employed in cognitive science.

1. For the most part in this paper I will focus on connections, forward, recurrent, and collateral within the cerebral cortex and hippocampal formation. But these present only a portion of the story. Most sensory information reaches cortex via regions in the thalamus, and there are almost ten times as many recurrent projections from cortex back to these thalamic regions as there are forward projection.

2 Endogenously active mechanisms in the brain

Claiming that the brain is endogenously active may strike some as comparable to proposing that it is a perpetual motion machine. That is, however, far from what is being proposed. All living organisms, and accordingly those with a nervous system and a brain, are open in the thermodynamic sense to matter and energy from their environment. What is distinctive about biological organisms is that they are organized systems—hence not in equilibrium with their environment—and that they maintain themselves in this non-equilibrium state despite the tendency exhibited by closed systems towards equilibrium (high entropy). The equilibrium tendencies are manifest in the continual degradation of organic structures that require living systems regularly to repair themselves [Rosen 1991] or else cease to exist as organized systems. Moreover, they must perform the operations needed to construct themselves—incorporate matter into the organized structure with which they are identified. Maturana and Varela [Maturana & Varela 1980] have referred to this as *autopoiesis* and viewed it as the foundation of cognition (see [Lyon 2006], for a systematic discussion). Crucial to autopoiesis, but not emphasized by Maturana and Varela, is that organisms must capture and employ free energy from their environments in autopoiesis and repair. Ruiz-Mirazo and Moreno have made this a central feature in their characterization of living organisms as *autonomous* [Ruiz-Mirazo & Moreno 2004] and Barandiaran and Moreno have extended the view to cognition [Barandiaran & Moreno 2006].

Organisms must initiate these activities of autopoiesis and repair from within, and this entails that they be endogenously active, not reactive. As Gánti emphasized in his proposal of the chemoton as the minimal chemical system capable of exhibiting the characteristics of life, cyclic processes that can regularly return to the same condition are fundamental to autonomous systems—as a result of cyclic organization, an organism can regularly restore itself to the conditions where it can perform the operations necessary to build and repair itself [Gánti 2003]. As long as the organization of the cyclically organized system is adequate to recruit free energy from the environment, it can continue to iterate the stages in the cycle. If one tracks variables representing states of the system through time, they will oscillate. If there are time delays within the cyclic system, and especially if the operations within the cycle are non-linear, these oscillations can be sustained indefinitely provided sufficient matter and energy are available. Goodwin pioneered the analysis of sustained oscillations in biological systems [Goodwin 1963], showing in a computational model how a feedback loop inspired by Jacob and Monod’s operon model for gene regulation in bacteria [Jacob & Monod 1961] could, with appropriate parameters, generate a sustained oscillatory system. Subsequently, other theorists such as Goldbeter have pursued the approach and have drawn attention to a substantial number of endogenous oscillations

in biological systems [Goldbeter 1996]. (For further discussion see [Bechtel & Abrahamsen 2011].)

While mechanisms that are endogenously active clearly occur in nature and are fundamental to living systems, many of the accounts of mechanism and mechanistic explanation advanced by the new mechanistic philosophers of science are inadequate to characterize them. Following what has been the dominant emphasis in mechanistic research in biology over the past several centuries, these accounts emphasize the parts and operations of mechanisms and the processes by which researchers decompose mechanisms into their parts and operations ([Bechtel & Richardson 1993-2010]; [Bechtel & Abrahamsen 2005]; [Glennan 1996, 2002]; [Machamer, Darden & Craver 2000]; [Thagard 2003]; [Wimsatt 2007]). These accounts of mechanism also note the importance of how components are organized but, like many scientists, they tend to privilege sequential organization. Thus, Machamer, Darden, & Craver include in their definition of a mechanism that they are “productive of regular changes from start or set-up to finish or termination conditions” [Machamer, Darden & Craver 2000, 3]. Sequential organization, however, cannot produce endogenous activity—this requires negative feedback or other cyclic designs. Moreover, not all negative feedback systems sustain oscillations; only ones with non-linear components and appropriate parameter values can do so. Determining whether a particular mechanism will generate continued activity (e.g., oscillation) or settle into a stable state requires employing the tools of computational modeling and dynamical systems analysis. Bechtel and Abrahamsen ([Bechtel & Abrahamsen 2010]; see also [Bechtel 2011]) designate explanations that invoke computational modeling to understand patterns of change over time in the properties of the parts and operations of a mechanism as *dynamic mechanistic explanations*.

Sustained oscillators are the simplest endogenously active mechanisms—as long as they can recruit free energy from their environment they are continually active. Neurons are examples of such sustained oscillators. The endogenous nature of neuronal activity, however, is often not appreciated. Since early in the 20th century, when techniques were developed for recording the electrical activity of individual neurons, neuroscientists have often treated the generation of action potentials as the principal activity of neurons. The common picture, both in textbooks and in philosophical discussions, is that neurons reside at their resting potential until they receive sufficient inputs on their dendrites to raise the voltage above threshold, after which they generate an action potential on their axon and, over time, return to their resting potential. On this view, neurons are reactive components. Although most accounts of neural processing focus on action potentials, they represent only a small part of the electrical activity in the brain. Even before individual action potentials could be recorded, researchers such as du Bois-Reymond had recorded electrical potentials from muscle cells and neurons [Du Bois-Reymond 1848-1884]. Bernstein had identified a greater concentration of potassium (K^+) ions inside than outside the cell, and proposed that the electrical current result-

ing from diffusion out of the cell explained the negative potential of the cell at rest [Bernstein 1912]. He further proposed that, when excited, other ions would diffuse across the membrane, eliminating the diffusion potential of the K^+ ions. Hodgkin and Huxley determined that in addition to the potassium gradient, there is a sodium (N^+) gradient (with N^+ concentrations greater outside the cell during the resting potential phase) and that the conductance of each ion varies independently and depends on the voltage across the membrane [Hodgkin & Huxley 1952]. Except when the relative concentration of a given ion on each side of the membrane (its membrane potential) equals its reversal or equilibrium potential, ions will flow either into or out of the cell.

Hodgkin and Huxley's interest in these currents was principally to explain the action potential, which most researchers continued to view from a reactive perspective. But some researchers, working as Hodgkin and Huxley had with invertebrate neurons, found specialized *pacemaker* neurons that generated their own rhythmic action potentials [Alving 1968]. Others began to focus on the variety of voltage-gated and other currents that could be observed across neuronal membranes and explored the complex patterns of change in these currents not only in axons but also in the dendrites and cell bodies of neurons (reviewed by [Kandel 1976]).

Initially mammalian researchers, who followed in the tradition of Sherrington and treated neurons as integrating inputs and firing when these inputs pushed them above their threshold, viewed these studies skeptically. This began to change in the 1970s and 1980s when Llinás and his collaborators found a variety of functionally important ion currents in neurons of the inferior olive and cerebellum in mammals and birds. Most were spatially distributed and gated by voltage in a different manner than the sodium and potassium channels in the axon, equipping them for functions other than the direct generation of action potentials. Notably, the dendrites were endowed with channels providing high-threshold conductance to calcium (Ca^{2+}) ions, enabling dynamically complex dendritic excitation in contrast to earlier assumptions of passive transmission of signals from synapses.² Moreover, the cell bodies of some neurons in the inferior olive had a different kind of calcium channel with a seemingly paradoxical low-threshold conductance that, in interaction with sodium and high-threshold calcium conductances, enabled these neurons to function as single-cell oscillators “capable of self-sustained rhythmic firing independent of synaptic input” [Llinás 1988, 1659].³ They sent these rhythmic action potentials to target neurons in the cerebellum that were able to respond at the same frequency, qualifying them as *resonators* in

2. This linked nerve excitability with the Ca^{2+} -dependent second messenger system that is important for regulating general cellular functions.

3. For further exposition, see [Buzsáki 2006, 181–183], who comments: “These findings [...] illustrate that nature went to a lot of trouble bringing together these channels at the right densities and location just to serve one purpose: oscillation.” For evidence extending the findings to sensory neurons in various mammalian species, see [Huguenard 1996].

the dynamical lexicon championed by Llinás—reacting, but in ways shaped by their internal properties. Llinás also investigated spontaneous oscillations in electrical potentials elsewhere in the brain.

Llinás' findings revealed that the neurons are oscillators—some of which maintain oscillations on their own while others resonate to oscillations initiated by others. This has consequences not just for how we conceive of neurons but how we understand their interactions. As Huygens had observed in his 1665 letters to de Sluse (letters n° 1333 of 24 February 1665, n° 1335 of 26 February 1665, n° 1345 of 6 March 1665 published in [Huygens 1888]), as long as there is a means to convey a signal (even a very weak one) between two pendulum clocks (a pendulum clock is an oscillator) with nearly the same periods, they tend to synchronize their oscillations and exhibit a common phase relation. When signals can be passed between multiple oscillators, some with different natural frequencies, complex patterns of activity can develop. A given component may synchronize with some but oscillate out of phase with other components.

Emphasizing the oscillatory dynamics across neural membranes is not in conflict with focusing on them as generating action potentials in response to stimuli. Inputs reduce the electrical potential of the neuron and action potentials arise when they cross a threshold. If the electrical potential of neurons is oscillating, there will be periods when they are less polarized and hence closer to threshold and other periods when they are more polarized. An input that may suffice to push neurons across the threshold when they are hypopolarized may be insufficient to do so when they are hyperpolarized. Moreover, when two neurons are linked in a circuit and their oscillations are synchronized, the periods when the first neuron is most likely to generate an action potential will correspond to periods when the second is likely to generate an action potential in response. On the other hand, when they are desynchronized, the second neuron will be less likely to generate an action potential at the time when the first neuron is most likely to. I turn now to considering how such dynamical behavior of neurons is manifest in populations of neurons engaged in cognitive tasks.

3 Oscillations detected with EEG and LFP

A variety of techniques have been developed that can detect synchronized electrical potentials of populations of neurons. One of the first involved electrodes, initially inserted into the scalp and subsequently placed on the scalp. In his pioneering research in which he coined the name *Elektrenkephalogramm* (*electroencephalogram* or EEG in English) for the technique, Berger reported that he could record electrical oscillations from lead foil electrodes affixed to the heads of human beings [Berger 1929, 1930]. In particular, he detected large-amplitude oscillations of approximately 10 Hz when subjects were awake with their eyes closed that he called *alpha* rhythms and smaller amplitude,




FIGURE 1: In this eight-second extract of a recording made by Berger [Berger 1930], the upper line is a subcutaneous EEG. It shows three seconds of predominantly alpha waves that were blocked 0.27 seconds after he stroked the subject's hand with a glass rod (indicated by arrow *B* on the 10 Hz timing signal at the bottom). For at least the next two seconds the EEG shows lower-amplitude, higher-frequency beta waves, not alpha waves. The middle line is an electrocardiogram recorded simultaneously. Extracted from Figure 5 in Gloor's translation [Gloor 1969, 82] of [Berger 1930].

faster oscillations (20-30 Hz) that became apparent once alpha rhythms declined when subjects opened their eyes, received input from another sense, or performed an attention demanding task (see Figure 1).

Berger and others initially assumed that the EEG was generated from action potentials. Eventually researchers recognized that they resulted not from action potentials but from synchronized post-synaptic potentials in dendrites discussed in the previous section [Bremer 1958]. Moreover, research such as that of Jahnsen and Llinás on the thalamus and thalamocortical relay neurons helped link dynamic behavior of individual neurons to the large-scale dynamics seen in EEG [Jahnsen & Llinás 1984].

In the decades after Berger's pioneering work, researchers identified both faster oscillations (greater than 30 Hz) that they designated *gamma* rhythms as well as a variety of slower oscillations (*delta*, 0-4 Hz, and *theta*, 4-8 Hz, rhythms). Typically oscillations at different frequencies are all combined in the electrical activity researchers record and techniques such as fast Fourier analysis are required to differentiate components in the overall signal. Moreover, researchers have also found that the same activity can be detected from electrodes implanted into brain tissue as long as the electrodes are not too close to any given neuron. When detected in this way the currents are referred to as local field potentials.

The most common application of the various rhythms detected with EEG has been to differentiate overall cognitive states—states of active awareness (gamma), quiet resting (alpha), and various stages of sleep (delta, theta). For many years little progress was made in linking EEG oscillations to more specifically characterized cognitive functions. Instead, the primary use of EEG in cognitive research has been in ERP studies that reflect the reactive conception of neural processing. The difference in the overall electrical activity produced

due to the stimulus is quite small in relation to the electrical activity that originates either endogenously or in response to other stimuli, so ERP researchers must average over many trials in order to extract a signal. Ironically, they are thereby removing from consideration the endogenous background activity that Berger had first identified.

In recent decades, however, some researchers have begun to identify correlates of cognitive activity directly in the EEG signal itself. One of the first hypotheses resulted from the detection of gamma oscillations in visual processing areas in anesthetized cats presented with moving bars with a particular orientation [Gray, König, Engel *et al.* 1989]. Singer and his colleagues developed an interpretative framework in which gamma oscillations served to temporarily bind together the neural representations of different features of a particular object which are represented in the action potentials of neurons in different visual processing regions (Singer characterized this as the binding-by-synchronization hypothesis [Singer 1999]). Subsequent researchers have detected gamma oscillations in awake, alert cats, monkeys, humans, and other species, in response to the presentation of shapes and smoothly transforming shapes (reviewed by [Fries 2009]). Gamma oscillations have also been found in somatosensory and auditory processing areas. Recent studies have shown that gamma synchronization requires not just inputs but top-down activation due to attention.

These gamma oscillations appear to be the result of initial inputs to pyramidal cells that then elicit responses in basket cells that distribute an inhibitory signal that blocks further action potentials in the pyramidal cells on which they synapse until their effects wear off in a synchronized manner, allowing synchronized spiking behavior. The population of basket cells are then entrained to maintain a regular oscillation which leaves windows in which action potentials can be generated in the pyramidal cells [Hasenstaub, Shu, Haider *et al.* 2005]. The effect of synchronized action potentials allows the outputs of many cells to be received by a downstream neuron in a common temporal window so that they sum sufficiently to produce an action potential in it.

Binding-by-synchronization retains a reactive conception of brain processing since it represents the features elicited by a stimulus as coming to be bound. But other investigations of gamma oscillation suggest that they may reflect an endogenous component of brain processing. Whittington, Traub, & Jefferys found that pharmacologically-isolated inhibitory neurons, when provided with a tonic excitatory drive (activation of metabotropic glutamate receptors), generate gamma oscillations without specific inputs [Whittington, Traub & Jefferys 1995]. Researchers soon found other conditions that would generate gamma oscillations and Wang and Buzsáki pioneered the project of constructing computational models to examine conditions under which networks of inhibitory interneurons would generate gamma oscillations [Wang & Buzsáki 1996]. However, their model required conditions not found in the experimental preparations. Kawaguchi, Katsumaru, Kosaka, Heizmann, & Hama's discovery of fast-spiking inhibitory neurons that express the calcium

binding protein parvalbumin provided a clue as to how these oscillations arise [Kawaguchi, Katsumaru, Kosaka *et al.* 1987]. Parvalbumin expressing neurons form networks of mutually connected neurons in which gamma oscillations are created and then transmitted to pyramidal cells via their output connections to these cells [Bartos, Vida & Jonas 2007]. The parvalbumin-expressing cells are highly active during gamma oscillations during which they generate action potentials once per cycle in a phase-locked manner [Gloveli 2005], supporting the hypothesis that they play a fundamental role in generating gamma oscillations. Moreover, they affect the pyramidal cells in a distinctive way—they do not cause them to hyperpolarize, but use chloride channels to shunt electrical activity. The result is a mechanism that is robust against the heterogeneity of the tonic drive—with high excitation an early conductance-dominated phase causes subsequent action potentials to occur later whereas with low excitation the depolarization-dominated phase causes subsequent action potentials to be advanced. The result is to create regular, clock-like epochs in which action potentials can arise; information processing is thus structured into discrete temporal windows.

Fries proposes that these gamma oscillations may be located within slower theta oscillations in such a fashion that theta oscillations provide larger temporal windows in which gamma synchrony may develop but is then interrupted so as to process new input [Fries 2009]. As evidence for this proposal he cites Rollenhagen and Olson's recordings from individual inferotemporal neurons after first presenting a preferred stimulus and then 600 ms later adding a non-preferred to the preferred stimulus [Rollenhagen & Olson 2005]. Presentation of the second stimulus resulted in a reduction in firing rate, followed by an approximately 5 Hz (theta range) oscillation between periods of enhanced and reduced firing, suggesting a theta frequency oscillation between the response to the preferred and non-preferred stimulus. Since such oscillation is also observed in recording from individual neurons when only the preferred stimulus is presented, Fries proposed the theta oscillations are an endogenous process that allows a period for developing a response to a stimulus followed by a period that opens the processing system up to seeking new input rather than simply continuing to respond to the first stimulus. It thus insures that any response to a stimulus is only metastable and that processing does not stop with it.

Both gamma and theta oscillations are also exhibited in the hippocampus where researchers have advanced hypotheses as to their contribution to cognitive processing of spatial information. The hippocampus is organized into a loop structure (Figure 2)—it receives input from the entorhinal cortex (EC) that is transmitted along a pathway to the dentate gyrus, the CA3 fields, the CA1 fields, and subiculum before sending inputs back to the entorhinal cortex. There is also a pathway directly from the EC to the CA1 fields.

Following reports that rats with hippocampal lesions exhibit deficits in the sorts of navigation tasks that had led Tolman to propose that they possess a cognitive map of their environment [Tolman 1948], O'Keefe began detailed




FIGURE 2: The loop architecture of the hippocampus in which input from entorhinal cortex is either processed in the dentate gyrus and CA3 before reaching CA1 or is directly received by CA1.

functional analysis of the response characteristics of hippocampal neurons. Recording from individual cells in the CA1 and CA4 regions (CA4 is a small region not commonly discussed) of the hippocampus, O'Keefe and Dostrovsky identified a number of cells that generated several bursts of action potentials whenever a rat was in a particular location in its enclosure [O'Keefe & Dostrovsky 1971]. O'Keefe later termed these *place cells* and argued that they constituted the rat's cognitive map of its environment [O'Keefe & Nadel 1978]. During the same period Ranck identified cells that generated bursts of action potentials in the theta range when the rat was moving [Ranck 1973]. Over the following decade research focused primarily on place cells, especially on conditions involving changes in the enclosure that would cause different cells to produce bursts of spikes despite the rat being in the same location. O'Keefe and Recce initiated a new phase of research when they explored the relation between place cell activity and theta oscillations and determined that the bursts of action potentials from place cells exhibited a somewhat higher frequency than the ongoing theta oscillation [O'Keefe & Recce 1993]. As shown in Figure 3, the initial burst would typically occur during the peak of the theta oscillations and subsequent bursts would precess earlier in the theta oscillations so that by the time the rat left the place field (the region that elicited activity from a particular cell), the bursts would have advanced nearly a full cycle. Considering only the activity of a single place cell, the amount of precession specifies how long the rat has been in the place field and thus how far it has advanced through it. From a population perspective, however, the information is even richer. Among cells whose place fields partially overlap, those place cells whose action potentials have precessed the most represent places the rat entered earlier while the ones in which action potentials have precessed the least represent those the rat has just entered. Thus, the activity across the population of place cells in relation to theta specifies the rat's trajectory.

In addition to theta activity, the hippocampus also exhibits gamma oscillations at two different frequencies, one originating in medial EC, which exhibits fast gamma oscillations (> 60 Hz), and the other in the CA3 region, which




FIGURE 3: Illustration of theta precession. As rat runs along the maze, it crosses the place field of a cell. The place cell spikes, shown in red, precess against the underlying theta oscillation, firing first just after the peak and moving progressively early on subsequent theta cycles. Reprinted by permission from The American Physiological Society: *Physiological Reviews*, Wang, X.-J., 90, ©2010, Figure 1A, [Wang 2010].

generates slow gamma oscillations (< 60 Hz). As shown in Figure 2, these represent the two regions from which CA1 receives input. I noted above that neurons are most likely to fire when they are synchronized with those from which they receive input. From these considerations and the fact that CA1 at different times exhibits fast and slow gamma oscillations, Colgin, Denninger, Fyhn, Hafting, Bonnevie, Jensen, Moser, & Moser have hypothesized that these different gamma oscillations regulate the processing of information—at different times the CA1 fields synchronize with the fast gamma of the entorhinal cortex and respond to input from it and at other times synchronize with the slow gamma of CA3 and hence respond to inputs from it [Colgin, Denninger, Fyhn *et al.* 2009]. The processing in DG and CA3 is thought to facilitate recognizing a location as a previously experienced one; when that fails, CA1 responds directly to the assumed to be novel place represented by the EC activity. An important process in the hippocampus is long-term potentiation (LTP), which strengthens synaptic connections on which an input is received when the recipient neuron actually generates an action potential. LTP occurs only when a neuron generates a spike during the through of the theta cycle, which happens when CA1 is synchronized with entorhinal cortex. This limits learning a new representation to occasions when the DG-CA3 loop fails to recognize the location, a situation in which it is appropriate to acquire a new representation. When an old location is recognized, CA1 synchronizes with CA3 and spikes occur at the peak of the theta cycle when LTP is blocked.

After a period during which the only way researchers were able to relate EEG activity to cognition was through ERP, which adopted a reactive perspective, researchers are beginning to formulate hypotheses as to how oscillations, especially those in the theta and gamma bands, that arise independently of specific stimuli figure in information processing. These oscillations have been proposed to create time windows in which neural processes may develop responses to inputs but then break out of the response to sample other inputs as well as ways of gating the flow of information between brain regions and controlling processes such as LTP. These proposals suggest that the brain is a dynamically active processing system in which the ongoing dynamics structures and regulates processing in time.

4 Oscillations detected with PET and fMRI

One of the limitations of EEG for studying cognitive processing in the brain is that while it provides high-resolution information about the timing of electrical activity, it is very difficult to localize spatially the source of the electrical signal. Thus although a small cadre of researchers employed EEG and ERP to investigate cognitive processing in the brain in the 1970s and 1980s, their investigations failed to spark the development of cognitive neuroscience as a prominent field of inquiry. Rather, this happened only after the development of techniques such as positron emission tomography (PET) in the late 1980s and functional magnetic resonance imaging (fMRI) in the 1990s. These tools used tomographic methods to localize the source of either a radioactive or magnetic signal that was assumed to be related to specific cognitive activities. As described above, in the initial research with PET and fMRI studies researchers sought to identify regions in which blood flow was correlated with cognitive activities. Thus like ERP researchers, PET and fMRI researchers initially viewed the cognitive system as a reactive mechanism and focused on detecting increased activity that could be attributed to the task a subject was performing. Accordingly, they paid little attention to the varying background signal except insofar as it posed challenges in detecting the signal attributed to the stimulus or task.

In the mid-1990s some researchers (e.g., [Ghatan, Hsieh, Wirsén-Meurling *et al.* 1995]; [Baker, Rogers, Owen *et al.* 1996]) drew attention to the fact that some brain regions seemed regularly to exhibit reduced activity (were *deactivated*) in task conditions than in a resting state (a condition in which the subject lay quietly in the scanner with no assigned task). To determine how widespread this effect was, Shulman, Corbetta, Buckner, Fiez, Miezin, Raichle, & Petersen conducted a meta-analysis of PET studies in which a task condition was compared to a non-task condition in which the same stimulus was presented (this produced results very similar to the resting condition with no stimulus) [Shulman, Corbetta, Buckner *et al.* 1997]. They identified a specific set of brain areas that were reliably less active in task situations: the junction of precuneus and posterior cingulate cortex (PCC), the inferior parietal cortex (IPC), the left dorsolateral prefrontal cortex (left DLPFC), a medial frontal strip that continued through the inferior anterior cingulate cortex (inferior ACC), the left inferior frontal cortex, the left inferior frontal gyrus, and the amygdala.

Shortly thereafter Raichle and his collaborators shifted their perspective from viewing these areas as deactivated in task conditions to viewing them as being more active in non-task conditions and began to characterize them as constituting a *default mode network*—one which performs actual functions best carried out when there are no external task demands. One clue as to their function was provided by Andreasen, O’Leary, Cizadlo, Arndt, Rezai, Watkins, Ponto, & Hichwa who found that episodic memory tasks were one type of task that did not result in reduced activity in these regions [Andreasen, O’Leary,

Cizadlo *et al.* 1995]. Some researchers drew upon this finding to suggest that the default mode network supports undirected thinking or what had been labeled *mind-wandering*. This coheres with the reports by Andreasen *et al.*'s subjects that, when required to lay in the scanner with no task requirements, they thought about their own past or planned future activities. These are both activities that draw upon episodic memory.

Adopting the mind-wandering hypothesis, Buckner, Andrews-Hanna, and Schacter link mind-wandering to the ability to carry out “flexible self-relevant mental explorations—simulations—that provide a means to anticipate and evaluate upcoming events before they happen” [Buckner, Andrews-Hanna & Schacter 2008, 2]. In defending this view they cite not only Andreasen *et al.*'s results but also correlations found by Mason *et al.* between stimulus independent thoughts and activity in the default network [Mason, Norton, Van Horn *et al.* 2007]. Gilbert, Dumontheil, Simons, Frith, & Burgess offer an alternative view that activity in the default network generates low-level generalized awareness or watchfulness [Gilbert, Dumontheil, Simons *et al.* 2007]. This alternative gains support from Hahn, Ross, & Stein's findings of increased activity in the default network in a target-detecting task when the target could appear anywhere, but not when it was expected in a specific location [Hahn, Ross & Stein 2007]. A potential problem for both of these treatments of the function of the default mode network is that it has been found to be active not just when individuals are awake but also during sleep ([Fukunaga, Horovitz, van Gelderen *et al.* 2006]; [Larson-Prior, Zempel, Nolan *et al.* 2009]) and under anesthesia ([Vincent, Patel, Fox *et al.* 2007])—circumstances in which neither spontaneous conscious thoughts (such as Andreasen *et al.*'s subjects report) nor generalized awareness and watchfulness occur. Even if activity in the default mode network is found in the absence of conscious thought, the default mode network may nonetheless provide a neurophysiological foundation for mind-wandering or watchfulness when the subject is engaged in conscious thought. Mind-wandering and watchfulness are features of our mental life that have been neglected by the reactive framework but as Buckner *et al.* suggest may be important cognitive activities that enable us to anticipate and cope with our environments.

The initial studies of the default mode network did not focus on the dynamics of the activity in these areas, but a study by Biswal, Yetkin, Haughton, & Hyde showed how fMRI could be used to identify dynamic processes in task activated networks that were soon extended to the default mode network [Biswal, Yetkin, Haughton *et al.* 1995]. These researchers performed a time series analysis of fMRI recordings taken every 250 ms and observed very low-frequency oscillations ($< .1$ Hz or one cycle every 10-15 seconds) as subjects performed a simple motor task (moving their hand). Moreover, they found that these oscillations synchronized across sensory and motor regions of the brain in both hemispheres, which they interpreted as indicating functional connections between these regions. Cordes, Haughton, Arfanakis, Wendt, Turski, Moritz, Quigley, & Meyerand found similar oscillations in resting state BOLD

signals in networks of areas previously identified as jointly exhibiting increased activation in sensorimotor, visual, receptive language, or expressive language tasks [Cordes, Haughton, Arfanakis *et al.* 2000]. Moreover, their *functional connectivity MRI (fcMRI)* analysis—applying correlational statistics to resting state BOLD time series data to determine patterns of synchronization—yielded functional networks very similar to those identified from activity during tasks. That is, areas within the same network had correlated patterns of activity across time (rising and falling in synchrony) regardless of whether overall level of activity was relatively high (e.g., the sensorimotor network while moving a hand) or relatively low (e.g., the same network in a resting state condition).

To determine whether the default mode network also exhibited synchronized activity, Greicius, Krasnow, Reiss, & Menon employed fcMRI with two seed areas identified with the default mode network, the PCC and inferior ACC [Greicius, Krasnow, Reiss *et al.* 2003]. Using the PCC as the seed, they found that its resting state oscillations were correlated with those in much of medial prefrontal cortex (including inferior ACC and orbitofrontal cortex), left DLPFC, IPC bilaterally, left inferolateral temporal cortex, and left parahippocampal gyrus, virtually the same range of areas as Shulman *et al.* had identified in their meta-analysis. They regarded their results as providing “the most compelling evidence to date for the existence of a cohesive, tonically active, default mode network” [Shulman, Corbetta, Buckner *et al.* 1997, 256]. When instead the ventral ACC was employed as the seed area, they demonstrated correlated activity not just in the PCC but also in the medial prefrontal cortex/orbitofrontal cortex, the nucleus accumbens, and the hypothalamus/midbrain and they argued that these primarily paralimbic and subcortical areas comprised a separate network important for calibrating affective and autonomic operations. They argued further that the strong connection between inferior ACC and PCC provided a crucial link between this and the default mode network. Subsequently Greicius and Menon found that the default network included the hippocampus [Greicius & Menon 2004] and Vincent *et al.* determined that by seeding an analysis with a hippocampal region they could find correlated activity in the rest of the default network [Vincent, Snyder, Fox *et al.* 2006]. Buckner, Andrews-Hanna, & Schacter summed up the perspective provided by this research:

The default network is a brain system much like the motor system or the visual system. It contains a set of interacting brain areas that are tightly functionally connected and distinct from other systems within the brain. [Buckner, Andrews-Hanna & Schacter 2008, 4–5]

The Cordes *et al.* results noted above showed that even in the resting state one could demonstrate correlated activity across networks that would then be jointly activated in task conditions. Fox *et al.* found correlated activity in the resting state in a network that was especially active during attention-demanding tasks (intraparietal sulcus, frontal eye field, middle temporal re-

gion, supplementary motor areas, and the insula) and showed that activity in that network and in the default network were anticorrelated [Fox, Snyder, Vincent *et al.* 2005]. That is, the areas that were positively correlated within each network were negatively correlated with areas in the other network—an outcome more interesting than a zero correlation. Fox *et al.* emphasized that:

anticorrelations may be as important as correlations in brain organization. Little has been said previously in the neuronal synchrony literature regarding the role of anticorrelations. While correlations may serve an integrative role in combining neuronal activity subserving similar goals or representations, anticorrelations may serve a differentiating role segregating neuronal processes subserving opposite goals or competing representations. [Fox, Snyder, Vincent *et al.* 2005, 9677]

The oscillations found using fMRI are one to two orders of magnitude slower than those observed with EEG. Do they, like theta and gamma oscillations, play a role in cognition? So far not a lot is known about their functional significance, but a very intriguing finding by the Raichle group suggests that it may have behavioral consequences. As I have noted, in reactive studies researchers constantly confront the fact that there is great variability in the recorded signal, which is often treated as noise. Fox, Snyder, Zacks, & Raichle devised an innovative strategy for showing that a large part of this variability may be the endogenous oscillation [Fox, Snyder, Zacks *et al.* 2006]. As subjects performed a button-press task with their right hand whenever they identified a new event in a video, a motor activity that is directed by the left side of the brain, the researchers identified the ongoing oscillation in sensory-motor areas on both sides of the brain. They found that the oscillations in the right side of the brain accounted for much of the variation found in the fMRI signal on the left side of the brain. Moreover, in a subsequent study the same group found that the spontaneous fluctuation exhibited in the contralateral hemisphere to that controlling the button push accounted for the force with which a subject pressed the button [Fox, Snyder, Vincent *et al.* 2007]. This finding is suggestive that endogenous activity may directly affected mental processes and behavioral activities.

5 Dynamically organized networks in the brain

The pattern of activity within correlated networks and between anticorrelated networks, whether these networks are task characterized or the default mode network, during conditions where subjects are not required to perform any tasks reveals that there is highly organized endogenous activity in the brain. During the same period as when fMRI was revealing these functional networks, other techniques (e.g., diffusion tensor MRI) were being developed to

analyze structural connectivity in the brain. The structurally and functionally characterized networks appear to be highly correlated. Greicius, Supekar, Menon, & Dougherty showed that the regions in the default mode network are anatomically connected [Greicius, Supekar, Menon *et al.* 2009]; van den Heuvel, Mandl, Kahn, & Pol found that eight of the nine networks they identified in the resting state correspond to ones that can be characterized anatomically as connected by fiber tracts [van den Heuvel, Mandl, Kahn *et al.* 2009].

Moreover, these networks have a similar type of structure. Whether analyzed structurally or functionally, they have been found to exhibit a small-world structure [Watts & Strogatz 1998]. Small-world organization is a form of organization that lies between regular lattices and randomly organized networks—as in lattices most connections are between nearby units but a few connect between distant locations. For information processing purposes, lattice structures have the virtue of creating modules of units that function together to perform a particular task whereas random networks exhibit a short path of communication across the network. Small-world networks possess both properties and many real-world networks have been found to exhibit small world organization, including the default mode network and task directed networks [Sporns 2010].

In thinking about cognitive architectures, it is common to construe them as set down before cognitive activity and simply setting the possibilities for such activity. But activity in the brain is now understood as capable of modifying the structural organization of the brain. The brain is plastic and processes such as LTP, discussed above, can change the connectivity of the brain. An intriguing possibility that Raichle has advanced is that the synchronized slow oscillations within networks may play a role in sculpting the brain. Gong and van Leeuwen showed in a computational model that when changing connection strengths between units occurred when units exhibit synchronized activity, small-world organization would develop in the network [Gong & van Leeuwen 2004]. Even when these networks begin with random connectivity, they self-organize into clusters linked to each other through hubs. Rubinov, Sporns, van Leeuwen, & Breakspear appeal to Gong and van Leeuwen’s model to suggest that as patterns of synchronized activity occur in the brain, the brain might self-organize into a small-world network with hubs [Rubinov, Sporns, van Leeuwen *et al.* 2009]. Some of the synchrony between neurons might be the result of specific stimuli or tasks the brain is required to perform. But if the synchronization between brain regions is ongoing, occurring during the resting state as well as under tasks conditions, then some of the synchronized activity that shapes the wiring of the brain will be a result of endogenous functioning, not just stimulus induced activity.

6 Developing an endogenously active cognitive architecture

In the preceding sections I have presented evidence from a variety of sources that brains not only respond to tasks or stimuli but also are endogenously active and that this activity may be important for understanding cognition. To begin with, neurons exhibit endogenous oscillations in the electrical activity on their membranes and they are more likely to generate action potentials when they are in a hypopolarized phase than when they are hyperpolarized. In addition, some neurons are capable of generating action potentials independently of any input. Further, at the level of large populations of neurons both EEG and fMRI reveal ongoing oscillatory activity that has traditionally been treated as noise. EEG detects oscillations in a variety of frequency bands above 1 Hz. Berger identified activity dominantly in the alpha band when subjects are sitting quietly without stimulation, but even during this period there are oscillations in other frequency ranges. Both theta and gamma oscillations in cortex appear to be important in creating temporal windows with gamma windows appearing to enable organized processing of inputs and theta oscillations serving to periodically reset the processing system to respond to alternative inputs. In the hippocampus different regions oscillate at different frequencies, with some regions able to entrain to and then respond to the oscillatory activity of different regions. Theta oscillations provide a timing window in which the relative timing of action potentials carries information and regulates whether activity such as LTP is to occur. The recent finding of oscillations at less than .1 Hz with fMRI provide further perspective on the role of endogenous activity in the brain, especially in revealing ongoing coordination of activity in multiple brain networks even during the resting state. There is suggestive evidence that this activity is manifest in behavior and underlies ongoing thought not tied to tasks (mind-wandering). The possible role of this activity in sculpting the brain suggests that it is important in establishing the processing capacities of the brain—i.e., the cognitive architecture.

The evidence of endogenous activity in the brain and suggestions that it may be important for cognition poses a challenge to cognitive science that relies on proposed cognitive architectures that are purely reactive. Moreover, it raises the question of how cognitive researchers might develop accounts of the cognitive architecture that incorporate endogenous activity. With individual neurons, endogenous activity appears in the ongoing oscillations in the electrical potentials over the membranes. One way of creating an endogenously active system is to construct it out of oscillators. Populations of oscillators have the potential to synchronize, but depending on the pattern of connections between them, they may only form temporary coalitions of synchronized components and so exhibit meta-stability. Accordingly, oscillators with processes for synchronizing them constitute systems capable of the sorts of endogenous activity found in the brain.

Cognitive architectures are instantiated in computational models. It is relatively easy to develop computational models that exhibit ongoing oscillations—in fields of biology dealing with oscillatory phenomena, such as circadian biology, researchers employ models of oscillatory components. Some of these represent the components of the mechanism (involving in the circadian case, the transcription and translation of genes) while others abstract from these details and employ, for example, equations describing a van der Pol oscillator [Bechtel & Abrahamsen 2010]. Populations are modeled by adding terms to the basic equations for oscillators to characterize the transmission of signals between them. In the context of neuroscience, computational modelers have increasingly directed their attention to modeling membrane potentials and using these model to understand the role of membrane potentials in generating population-level rhythms [Destexhe & Sejnowski 2003]. These models provide resources for creating computational models of cognitive architectures that exhibit endogenous dynamical activity. Developing and using such alternative cognitive architectures would provide cognitive researchers a basis for investigating the possible contributions of endogenous activity to cognition. On the other hand, by relying only on cognitive architectures that do not incorporate endogenous activity, cognitive scientists may lack the resources for modeling cognition that occurs in brains that do make endogenous activity available as a cognitive resource.

Bibliography

ALVING, B. O.

1968 Spontaneous activity in isolated somata of *Aplysia* pacemaker neurons, *Journal of General Physiology*, 51, 29–45.

ANDERSON, J. R.

1990 *The Adaptive Character of Thought*, Hillsdale, NJ: Lawrence Erlbaum Associates.

2007 *How can the Human Mind Occur in the Physical Universe?*, Oxford: Oxford University Press.

ANDREASEN, N. C., O'LEARY, D. S., CIZADLO, T., ARNDT, S., REZAI, K. & WATKINS, G. L. ET AL.

1995 Remembering the past: Two facets of episodic memory explored with positron emission tomography, *American Journal of Psychiatry*, 152, 1576–1585.

ANTROBUS, J. S., SINGER, J. L., GOLDSTEIN, S. & FORTGANG, M.

1970 Mindwandering and cognitive structure, *Transactions of the New York Academy of Sciences*, 32, 242–252.

- BAKER, S. C., ROGERS, R. D., OWEN, A. M., FRITH, C. D., DOLAN, R. J. & FRACKOWIAK, R. S. J. ET AL.
1996 Neural systems engaged by planning: A PET study of the Tower of London task, *Neuropsychologia*, 34, 515–526.
- BARANDIARAN, X. & MORENO, A.
2006 On what makes certain dynamical systems cognitive: A minimally cognitive organization program, *Adaptive Behavior*, 14, 171–185.
- BARTOS, M., VIDA, I. & JONAS, P.
2007 Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks, *Nature Reviews Neuroscience*, 8, 45–56.
- BECHTEL, W.
2008 *Mental Mechanisms*, London: Routledge.
2011 Mechanism and biological explanation, *Philosophy of Science*, 78, 533–557.
- BECHTEL, W. & ABRAHAMSEN, A.
2005 Explanation: A mechanist alternative, *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36, 421–441.
2010 Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science, *Studies in History and Philosophy of Science Part A*, 41, 321–333.
2011 Complex biological mechanisms: Cyclic, oscillatory, and autonomous, in *Philosophy of complex systems. Handbook of the philosophy of science*, edited by HOOKER, C. A., New York: Elsevier, vol. 10, 257–285.
- BECHTEL, W. & RICHARDSON, R. C.
1993-2010 *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*, Cambridge, MA: MIT Press, 2nd ed., first published in 1993, Princeton University Press.
- BERGER, H.
1929 Über das Elektrenkephalogramm des Menschen, *Archiv für Psychiatrie und Nervenkrankheiten*, 87, 527–570.
1930 Über das Elektrenkephalogramm des Menschen. Zweite Mitteilung, *Journal für Psychologie und Neurologie*, 40, 160–179.
- BERNSTEIN, J.
1912 *Elektrobiologie. Die Lehre von den elektrischen Vorgängen im Organismus auf moderner Grundlage dargestellt*, Braunschweig: Vieweg & Sohn.

- BISWAL, B., YETKIN, F. Z., HAUGHTON, V. M. & HYDE, J. S.
 1995 Functional connectivity in the motor cortex of resting human brain using echo-planar MRI, *Magnetic Resonance in Medicine*, 34, 537–541.
- BREMER, F.
 1958 Cerebral and cerebellar potentials, *Physiological Reviews*, 38, 357–388.
- BROWN, T. G.
 1911 The intrinsic factors in the act of progression in the mammal, *Proceedings of the Royal Society of London, Series B, Containing Papers of a Biological Character*, 84, 308–319.
 1914 On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a theory of the evolution of function in the nervous system, *The Journal of Physiology*, 48, 18–46.
- BUCKNER, R. L., ANDREWS-HANNA, J. R. & SCHACTER, D. L.
 2008 The brain's default network, *Annals of the New York Academy of Sciences*, 1124, 1–38.
- BUZSÁKI, G.
 2006 *Rhythms of the Brain*, Oxford: Oxford University Press.
- COLGIN, L. L., DENNINGER, T., FYHN, M., HAFTING, T., BONNEVIE, T. & JENSEN, O. ET AL.
 2009 Frequency of gamma oscillations routes flow of information in the hippocampus, *Nature*, 462, 353–357.
- CORBETTA, M. & SHULMAN, G. L.
 2002 Control of goal-directed and stimulus-driven attention in the brain, *Nature Reviews Neuroscience*, 3, 201–215.
- CORDES, D., HAUGHTON, V. M., ARFANAKIS, K., WENDT, G. J., TURSKI, P. A. & MORITZ, C. H. ET AL.
 2000 Mapping functionally related regions of brain with functional connectivity MR imaging, *American Journal of Neuroradiology*, 21, 1636–1644.
- DESTEXHE, A. D. & SEJNOWSKI, T. J.
 2003 Interactions between membrane conductances underlying thalamo-cortical slow-wave oscillations, *Physiological Reviews*, 83, 1401–1453.
- DU BOIS-REYMOND, E.
 1848–1884 *Untersuchungen über thierische Elektrizität*, Berlin: Reimer.

- DUCH, W., OENTARYO, R. J. & PASQUIER, M.
2008 Cognitive architectures: Where do we go from here?, in *Proceedings of the 2008 conference on Artificial General Intelligence*, IOS Press, 122–136.
- FOX, M. D., SNYDER, A. Z., VINCENT, J. L., CORBETTA, M., VAN ESSEN, D. C. & RAICHLER, M. E.
2005 The human brain is intrinsically organized into dynamic, anticorrelated functional networks, *Proceedings of the National Academy of Sciences of the United States of America*, 102, 9673–9678.
- FOX, M. D., SNYDER, A. Z., VINCENT, J. L. & RAICHLER, M. E.
2007 Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior, *Neuron*, 56, 171–184.
- FOX, M. D., SNYDER, A. Z., ZACKS, J. M. & RAICHLER, M. E.
2006 Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses, *Nature Neuroscience*, 9, 23–25.
- FRIES, P.
2009 Neuronal gamma-band synchronization as a fundamental process in cortical computation, *Annual Review of Neuroscience*, 32, 209–224.
- FUKUNAGA, M., HOROVITZ, S. G., VAN GELDEREN, P., DE ZWART, J. A., JANSMA, J. M. & IKONOMIDOU, V. N. ET AL.
2006 Large-amplitude, spatially correlated fluctuations in BOLD fMRI signals during extended rest and early sleep stages, *Magnetic Resonance Imaging*, 24, 979–992.
- GÁNTI, T.
2003 *The Principles of Life*, New York: Oxford University Press.
- GHATAN, P. H., HSIEH, J.-C., WIRSÉN-MEURLING, A., WREDLING, R., ERIKSSON, L. & STONE-ELANDER, S. ET AL.
1995 Brain activation induced by the perceptual maze test: A PET study of cognitive performance, *Neuroimage*, 2, 112–124.
- GILBERT, S. J., DUMONTEIL, I., SIMONS, J. S., FRITH, C. D. & BURGESS, P. W.
2007 Comment on “Wandering minds: The default network and stimulus-independent thought”, *Science*, 317, 43b.
- GLENNAN, S.
1996 Mechanisms and the nature of causation, *Erkenntnis*, 44, 50–71.
2002 Rethinking mechanistic explanation, *Philosophy of Science*, 69, S342–S353.
- GLOOR, P.
1969 *Hans Berger on the Electroencephalogram of Man*, Amsterdam: Elsevier.

- GLOVELI, T.
2005 Differential involvement of oriens/pyramidale interneurons in hippocampal network oscillations in vitro, *Journal of Physiology*, 562, 131–147.
- GOLDBETER, A.
1996 *Biochemical Oscillations and Cellular Rhythms: The molecular bases of periodic and chaotic behaviour*, Cambridge: Cambridge University Press.
- GONG, P. & VAN LEEUWEN, C.
2004 Evolution to a small-world network with chaotic units, *Europhysics Letters*, 67, 328–333.
- GOODWIN, B. C.
1963 *Temporal Organization in Cells; A dynamic theory of cellular control processes*, London: Academic.
- GRAY, C. M., KONIG, P., ENGEL, A. K. & SINGER, W.
1989 Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties, *Nature*, 338, 334–337.
- GREICIUS, M. D., KRASNOW, B., REISS, A. L. & MENON, V.
2003 Functional connectivity in the resting brain: A network analysis of the default mode hypothesis, *Proceedings of the National Academy of Sciences of the United States of America*, 100, 253–258.
- GREICIUS, M. D. & MENON, V.
2004 Default-mode activity during a passive sensory task: Uncoupled from deactivation but impacting activation, *Journal of Cognitive Neuroscience*, 16, 1484–1492.
- GREICIUS, M. D., SUPERKAR, K., MENON, V. & DOUGHERTY, R. F.
2009 Resting-state functional connectivity reflects structural connectivity in the default mode network, *Cerebral Cortex*, 19, 72–78.
- HAHN, B., ROSS, T. J. & STEIN, E. A.
2007 Cingulate activation increases dynamically with response speed under stimulus unpredictability, *Cerebral Cortex*, 17, 1664–1671.
- HASENSTAUB, A., SHU, Y., HAIDER, B., KRAUSHAAR, U., DUQUE, A. & MCCORMICK, D. A.
2005 Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks, *Neuron*, 47, 423–435.
- HODGKIN, A. L. & HUXLEY, A. F.
1952 A quantitative description of membrane current and its application to the conduction and excitation of nerve, *Journal of Physiology*, 117, 500–544.

- HUBEL, D. H. & WIESEL, T. N.
1962 Receptive fields, binocular interaction and functional architecture in the cat's visual cortex, *Journal of Physiology*, 160, 106–154.
1968 Receptive fields and functional architecture of monkey striate cortex, *Journal of Physiology*, 195, 215–243.
- HUGUENARD, J. R.
1996 Low-threshold calcium currents in central nervous system neurons, *Annual Review of Physiology*, 58, 329–348.
- HUYGENS, C.
1888 *Œuvres complètes*, La Haye: M. Nijhoff.
- JACOB, F. & MONOD, J.
1961 Genetic regulatory systems in the synthesis of proteins, *Journal of Molecular Biology*, 3, 318–356.
- JAHNSEN, H. & LLINÁS, R. R.
1984 Electrophysiological properties of guinea-pig thalamic neurones: an in vitro study, *The Journal of Physiology*, 349, 205–226.
- KANDEL, E. R.
1976 *Cellular Basis of Behavior: An introduction to behavioral neurobiology*, San Francisco: W. H. Freeman.
- KAWAGUCHI, Y., KATSUMARU, H., KOSAKA, T., HEIZMANN, C. W. & HAMA, K.
1987 Fast spiking cells in rat hippocampus (CA1 region) contain the calcium-binding protein parvalbumin, *Brain Research*, 416, 369–374.
- KUFFLER, S. W.
1953 Discharge patterns and functional organization of mammalian retina, *Journal of Neurophysiology*, 16, 37–68.
- KUTAS, M. & HILLYARD, S. A.
1980 Reading senseless sentences: brain potentials reflect semantic incongruity, *Science*, 207, 203–205.
- LAIRD, J. E., NEWELL, A. & ROSENBLOOM, P. S.
1987 Soar: An architecture for general intelligence, *Artificial Intelligence*, 33, 1–64.
- LARSON-PRIOR, L. J., ZEMPEL, J. M., NOLAN, T. S., PRIOR, F. W., SNYDER, A. Z. & RAICHLE, M. E.
2009 Cortical network functional connectivity in the descent to sleep, *Proceedings of the National Academy of Sciences*, 106, 4489–4494.
- LLINÁS, R.
1988 The intrinsic electrophysiological properties of mammalian neurons: Insights into central nervous system function, *Science*, 242, 1654–1664.

- LORENTE DE NÓ, R.
1938 Analysis of the activity of the chains of internuncial neurons, *Journal of Neurophysiology*, 1, 207–244.
- LYON, P.
2006 The biogenic approach to cognition, *Cognitive Processing*, 7, 11–29.
- MACHAMER, P., DARDEN, L. & CRAVER, C. F.
2000 Thinking about mechanisms, *Philosophy of Science*, 67, 1–25.
- MASON, M. F., NORTON, M. I., VAN HORN, J. D., WEGNER, D. M., GRAFTON, S. T. & MACRAE, C. N.
2007 Wandering minds: The default network and stimulus-independent thought, *Science*, 315, 393–395.
- MATURANA, H. R. & VARELA, F. J.
1980 Autopoiesis: The organization of the living, in *Autopoiesis and cognition: The realization of the living*, edited by MATURANA, H. R. & VARELA, F. J., Dordrecht: D. Reidel, 59–138.
- O'KEEFE, J. A. & DOSTROVSKY, J.
1971 The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat, *Brain Research*, 34, 171–175.
- O'KEEFE, J. A. & NADEL, L.
1978 *The Hippocampus as a Cognitive Map*, Oxford: Oxford University Press.
- O'KEEFE, J. A. & RECCE, M. L.
1993 Phase relationship between hippocampal place units and the EEG theta rhythm, *Hippocampus*, 3, 317–330.
- PETERSEN, S. E., FOX, P. T., POSNER, M. I., MINTUN, M. & RAICHEL, M. E.
1988 Positron emission tomographic studies of the cortical anatomy of single-word processing, *Nature*, 331, 585–588.
- RANCK, J. B.
1973 Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats: Part I. Behavioral correlates and firing repertoires, *Experimental Neurology*, 41, 462–531.
- ROLLENHAGEN, J. E. & OLSON, C. R.
2005 Low-frequency oscillations arising from competitive interactions between visual stimuli in macaque inferotemporal cortex, *Journal of Neurophysiology*, 94, 3368–3387.
- ROSEN, R.
1991 *Life itself: A comprehensive inquiry into the nature, origin, and fabrication of life*, New York: Columbia.

- RUBINOV, M., SPORNS, O., VAN LEEUWEN, C. & BREAKSPEAR, M.
2009 Symbiotic relationship between brain structure and dynamics, *BMC Neuroscience*, 10, 55.
- RUIZ-MIRAZO, K. & MORENO, A.
2004 Basic autonomy as a fundamental step in the synthesis of life, *Artificial Life*, 10, 235–259.
- RUMELHART, D. E. & MCCLELLAND, J. L.
1986 *Explorations in the microstructure of cognition. Volume 1. Foundations*, Cambridge, MA: Bradford Books, MIT Press.
- SHERRINGTON, C. S.
1923 *The Integrative Action of the Nervous System*, New Haven: Yale University Press.
- SHULMAN, G. L., CORBETTA, M., BUCKNER, R. L., FIEZ, J. A., MIEZIN, F. M. & RAICHLE, M. E. ET AL.
1997 Common blood flow changes across visual tasks: I. increases in subcortical structures and cerebellum but not in nonvisual cortex, *Journal of Cognitive Neuroscience*, 9, 624–647.
- SINGER, W.
1999 Neuronal synchrony: A versatile code for the definition of relations?, *Neuron*, 24, 49–65.
- SPORNS, O.
2010 *Networks of the Brain*, Cambridge, MA: MIT Press.
- STERNBERG, S.
1966 High-speed scanning in human memory, *Science*, 153, 652–654.
- THAGARD, P.
2003 Pathways to biomedical discovery, *Philosophy of Science*, 70, 235–254.
- TOLMAN, E. C.
1948 Cognitive maps in rats and men, *Psychological Review*, 55, 189–208.
- VAN DEN HEUVEL, M. P., MANDL, R. C. W., KAHN, R. S. & POL, H. E. H.
2009 Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain, *Human Brain Mapping*, 30, 3127–3141.
- VAN ESSEN, D. C. & GALLANT, J. L.
1994 Neural mechanisms of form and motion processing in the primate visual system, *Neuron*, 13, 1–10.
- VINCENT, J. L., PATEL, G. H., FOX, M. D., SNYDER, A. Z., BAKER, J. T. & VAN ESSEN, D. C. ET AL.
2007 Intrinsic functional architecture in the anaesthetized monkey brain, *Nature*, 447, 83–86.

- VINCENT, J. L., SNYDER, A. Z., FOX, M. D., SHANNON, B. J., ANDREWS, J. R. & RAICHLER, M. E. ET AL.
2006 Coherent spontaneous activity identifies a hippocampal-parietal memory network, *Journal of Neurophysiology*, 96, 3517–3531.
- WANG, X.-J.
2010 Neurophysiological and computational principles of cortical rhythms in cognition, *Physiological Review*, 90, 1195–1268.
- WANG, X.-J. & BUZSÁKI, G.
1996 Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model, *Journal of Neuroscience*, 16, 6402–6413.
- WATTS, D. & STROGRATZ, S.
1998 Collective dynamics of small worlds, *Nature*, 393, 440–442.
- WHITTINGTON, M. A., TRAUB, R. D. & JEFFERYS, J. G. R.
1995 Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation, *Nature*, 373, 612–615.
- WILSON, D. M. & WYMAN, R. J.
1965 Motor output patterns during random and rhythmic stimulation of locust thoracic ganglia, *Biophysical Journal*, 5, 121–143.
- WIMSATT, W. C.
2007 *Re-engineering Philosophy for Limited Beings: Piecewise approximations to reality*, Cambridge, MA: Harvard University Press.