



ORDER: GOD'S, MAN'S AND NATURE'S

Introduction by Eric Martin

Mechanistic explanations constitute a noteworthy departure from appeals to universal natural laws. While biological laws have long seemed difficult to come by, working out satisfactory alternatives has been an on-going project for scientists and philosophers alike. Of particular interest to the Order Project is the local character of mechanistic explanation. Mechanisms explain, where they do, in particular domains, in specific contexts, and in certain organisms. The extent of their reach must be determined empirically and depends largely on evolutionary conservation, but in any case they are far from earlier deductive-nomological accounts of explanation advanced by some logical positivists. Mechanistic explanations provide an image of life sciences that advance by piecemeal understanding of discrete parts of nature. If scientific practice provides our best warrant for belief in nature, then such discreteness is precisely the image of nature that the Order Project seeks to investigate further.

Localizing Cognitive Operations in Parts of Dynamically Active Brains

William Bechtel

Department of Philosophy, Center for Chronobiology, and
Interdisciplinary Programs in Cognitive Science and Science Studies
University of California, San Diego

The project of localizing cognitive operations in the brain has a long history and many impressive successes. It is a core element in the practice of giving mechanistic explanations of mental abilities. But it has also been challenged by prominent critics. One of the critics' claims is that brain regions are not specialized for specific cognitive operations. Most recently this claim has been advanced by theorists promoting a dynamical-systems perspective on cognition. There are, however, two ways to view the dynamical-systems perspective. The first is as a competitor to the mechanist perspective, rejecting altogether the conception of the brain as a mechanism or set of mechanisms underlying mental phenomena. The second is as a corrective to an overly simplistic conception of a mechanism and as complementary to a more adequate understanding of how mechanisms function. In this chapter I defend the later perspective. On this perspective, the traditional project of localizing mental operations in the brain is still important, but both the cognitive operations and brain regions in which they are localized must be conceived in the context of a dynamically active system.

In the first section, I describe the traditional project of localizing cognitive operations in the brain and situate it within the framework of mechanistic explanations of psychological phenomena. On this view, brain regions are construed as parts of a mechanism that are specialized for specific (information) processing tasks and perform these tasks when called upon by appropriate initiating conditions. I characterize the accompanying conception of brain regions in terms of Simon's conception of hierarchically organized, nearly decomposable components. I also briefly describe the opposition of some contemporary dynamists who reject the very project of parceling of the brain into regions performing distinct operations.

In the second section I develop a different perspective on brain regions, one that construes them as active components in a dynamically self-organizing system. On this account, the individual regions of the brain are endogenously active and as a result of this activity organize into specialized processing components. Yet even as the regions of the brain specialize, they remain integrated with other regions in a mode of organization known as a *small-world network*. In such organization local clustering gives rise to specialized regions but long-range connections link processing within these regions to activity elsewhere in the system, allowing activity elsewhere to modulate the behavior of local clusters. I present several sources of evidence that the brain exhibits small-world properties at multiple levels of organization.

In the final section I return to the question of localizing cognitive operations in brain regions and first examine evidence that small-world organization is also exhibited in the endogenous functioning of the brain and that functionally characterized small-world networks correlate with those characterized structurally. As a result, neither the decomposition of the mechanism into parts nor operations yields the sort of distinct parts or distinct operations commonly assumed in a mechanistic account, but rather parts and operations that are highly integrated with each other in

ways that undermine traditional approaches to link well-specified operations to clearly differentiated parts. Instead, the operations performed by a part must be recognized to vary substantially depending on their interactions with other parts. To understand this conception of the brain, a different conception of mechanistic explanation is required, one I refer to as *dynamic mechanistic explanation*. Dynamic mechanistic explanations are still mechanistic, and so make reference to operations localized within parts, but respect the dynamic processes that require characterizing both parts and operations relationally in terms of how they are situated in endogenously active dynamic networks.

I. The traditional approach to and criticism of localization of mental operations in brain mechanisms

Although claims that particular brain regions are responsible for specific mental phenomena were advanced early in the development of neuroscience (most notoriously by Gall around 1800), researchers only began to utilize naturally occurring or experimentally induced lesions or electrical stimulation to link mental processes experimentally to specific brain regions in the second half of the 19th century. Especially noteworthy accomplishments were Broca's (1861) localization of articulate speech in the region in left prefrontal cortex that came to bear his name (based on deficits exhibited in patients with lesions to this region), Fritsch and Hitzig's (1870) localization of motor control of muscle contractions to regions of the motor strip (based on responses after mild electrical stimulation to the cortex in dogs), and Munk's (1881) localization of basic visual processes to the occipital lobe (based on deficits following lesions to the occipital cortex in various species).

These successes seemed to answer the skeptical worries about localization advanced by Flourens (1824) in criticizing Gall. In his lesion experiments on pigeons, Flourens had succeeded in differentiating functions between the brain stem, cerebellum, and cerebrum, but failed to find specialized regions for perception, memory, or problem solving within the cerebrum. Rather, he claimed that lesions to the cerebrum generated deficits in proportion to how much was removed and embraced a holist perspective of the cerebrum. In the wake of the success of Broca, Fritsch and Hitzig, and Munk, it was no longer plausible to reject all localization of function in cortex, but many investigators adopted the perspective advanced by Lashley (1950)—primary sensory and motor processing was localized, but further processing and especially formation of memories depended on generalized *association* areas which operated by a principle of mass action, whereby performance decreased proportional to the amount of cortex lost, with no specific deficits corresponding to particular losses. Although lesion studies contributed to answering Lashley, further successes in localizing information processing in the brain relied heavily on the development of techniques for recording action potentials from individual neurons and correlating increased generation of action potentials in particular neurons with the presentation of specific stimuli. After Hubel and Wiesel (1962, 1968) showed that the cells in occipital cortex that Munk had viewed as the locus of visual processing seemed only to differentiate edges, they and other researchers looked more anteriorly and identified areas in temporal and parietal cortex responsive to other features of visual stimuli including motion, color, shape, and object identity.

I have analyzed the history of localizing different steps in processing visual inputs in different cortical regions elsewhere (Bechtel 2008). What is important for present purposes is the

conception of brain processing assumed in this research. The assumption is that the brain contains a mechanism for processing visual inputs and that the goal was to discover this mechanism. A mechanism for these purposes is a system that consists of distinguishable parts that perform specific operations which are organized to generate the phenomenon of interest (Bechtel & Richardson 1993; Bechtel & Abrahamsen 2005; Machamer, Darden, & Craver 2000). Differentiating parts and operations required the development of appropriate research techniques to decompose the brain structurally and functionally. Using staining techniques to differentiate the distribution of neurons in cortex, for example, Brodmann (1909/1994) and his contemporaries developed maps of different brain regions early in the 20th century. These maps were further differentiated using such tools for identifying patterns of connectivity between neurons and functional maps (Felleman & van Essen 1991). It was by recording from these areas as different stimuli were presented to an animal (typically a macaque monkey) that researchers localized various stages in the processing of visual information in cortex.

Implicitly, in differentiating brain areas and localizing functions to them, neuroscientists adopt the conception of a *hierarchically organized nearly decomposable* system articulated by Simon (1962). A fully decomposable system (which Wimsatt 1986, characterizes as an aggregative system) is one in which each component functions independently of others, and the whole is just the collection of separate components. Because the components do not cooperate to accomplish anything, such a system is not a mechanism. In a nearly decomposable system, the components perform independent functions but interact in ways that are “not negligible.” Simon characterizes such nearly decomposable systems in two propositions: “(a) in a nearly decomposable system, the short-run behavior of each of the component subsystems is approximately independent of the short-run behavior of the other components; (b) in the long run, the behavior of any one of the components depends in only an aggregate way on the behavior of the other components” (Simon 1962, p. 474). Simon illustrates such organization by describing a building that is thermally insulated from the environment and is divided into rooms that are well, but not perfectly insulated and that initially vary in temperature. In the short run each room goes to thermal equilibrium, little affected by the other rooms. More gradually the whole building reaches thermal equilibrium as heat dissipates from warmer to cooler rooms. The equations describing the initial changes need only employ variables for properties of the local room (as the other rooms will have negligible effects on it) and the equations for the second round of changes need only include variables for mean values of the rooms. Applied to the brain, in the short term processing within local brain regions will proceed, affected little by processing in other regions, whereas over the longer term the completed processing in one area will affect that in other areas to which it is connected.

Simon offers two sorts of arguments for the ubiquity of nearly decomposable systems, especially in living organisms. First, they are the most likely to evolve. He refers to the subsystems within nearly decomposable systems as *stable subassemblies* and illustrates the advantage of building larger systems from stable subassemblies with a tale of two watch makers, Tempus and Hora. Each makes equally fine watches consisting of approximately 1,000 parts that take nearly a day to assemble. Without stable subassemblies, Tempus has to position all 1,000 pieces before his watches are stable. Hora builds his watches by first building assemblies of about 10 pieces that remain stable, and then combining about ten of these into further stable assemblies, and finally combining these into the whole watch. With even a modest rate of interruptions (one every 100

steps of assembly), Tempus will only produce a stable product 44 times per million attempts, while Hora will complete an assembly before interruption 90% of the time. Simon extrapolates this lesson to biological evolution, arguing that if more complex biological systems evolved from stable simpler systems, natural selection would be able to generate complex life form far faster than supposed in various anti-evolutionist criticisms. Simon also argues for hierarchical organization of nearly decomposable systems by suggesting that only such systems would be intelligible to us:

If there are important systems in the world that are complex without being hierarchic, they may to a considerable extent escape our observation and our understanding. Analysis of their behavior would involve such detailed knowledge and calculation of the interactions of their elementary parts that it would be beyond our capacities of memory or computation (p. 477).

Although not strictly required by the basic conception of a mechanism or by Simon's conception of hierarchically organized nearly decomposable systems, a natural assumption is that a mechanism built in such a manner will function through a sequence of specialized operations, each of which makes its product available for further processing by components conceived as down stream of it. Thus, mammalian visual processing is assumed to begin with processing by the retina and continue through the LGN to V1, V2, etc. This conception was slightly complicated by the discovery of two different pathways from V1, one culminating in medial temporal cortex and involved in object recognition and a second culminating in parietal cortex and involved either in identifying where the stimulus is located (Ungerleider & Mishkin 1982) or the motor response one might make to it (Milner & Goodale 1995). As there are projections between the pathways at various stages of processing, van Essen and Gallant (1994) propose we think of them as streams which, while mostly flowing in isolation, sometimes sending tributaries to the other. They, however, retain the sequential perspective, which is indeed highly natural for humans trying to understand the operation of a mechanism. We think of a mechanism as doing something (e.g., locating an object in the space around us) and we investigate how it performs the sequence of operations needed to do so. Machamer, Darden, and Craver (2000) in fact enshrine the perspective in their definition of a mechanism as "productive of regular changes from start or set-up to finish or termination conditions."

As I noted at the outset, this classical conception of the brain as a mechanism has been challenged by theorists who advocate a dynamical systems approach to brain function. This approach characterizes brain function in terms of reciprocal causal interactions whereby "each and every component of a system contributes to every behavior of the whole system" (van Orden & Paap 1997, p. S92). While it might seem that such a holistic perspective would provide no tools for understanding specific activities, dynamical systems theory provides a variety of ways to identify different patterns of relations between variables characterizing the system (e.g., by describing trajectories in the state-space defined in terms of these variables and the patterns of these trajectories) that can be related to different cognitive activities. The resulting approach emphasizes patterns in the global behavior of the brain that correspond to particular cognitive activities while eschewing any attempt to localize these cognitive activities in regions of the brain. In arguing for this approach as preferable to the mechanistic approach of localizing operations within the brain through techniques such as neuroimaging (van Orden & Paap 1997; Uttal 2001) or analysis of lesions (van Orden, Pennington, & Stone 2001), dynamicist critics

emphasize the persistent variability in the data appealed to in localization studies; while this variability is usually treated as noise, the critics see it as a indicator of the flawed conception of the brain that underlies localization.

In the following sections I propose to reconcile the seemingly incompatible localizationist, mechanistic perspective and the dynamical systems approach, a reconciliation that requires reconceptualizing the project of localization. I begin by focusing on how recent research on dynamical activity in the brain requires reconceiving brain regions.

II. Reconceiving brain regions from a dynamical, interactive perspective

As plausible as Simon's account of hierarchically organized nearly decomposable systems and Machamer, Darden, and Craver's conception of the sequential operation are as ways of understanding mechanisms, they fit poorly with the modes of organization increasingly being identified in actual biological systems. Within the visual system, and indeed in the brain generally, there are as many projections backwards through the presumed processing stages as there are forward projections, and there are even more lateral projections within and between brain regions (Lorente de Nó 1938). While widely acknowledged, these projections are often neglected in functional accounts since it is not clear how they promote the presumed processing ends of the system. One of their effects, as we will see, is to undercut the view of brain regions as nearly decomposable parts while still allowing for localized regions that perform specialized operations.

To motivate reconceptualizing brain regions, consider again Simon's proposal that complex systems are constructed from stable sub-assemblies. Such assembly of independent components into larger systems appears to occur infrequently in evolution and even when it does, the result is an integrated and so less decomposable system. One of the best supported proposals for such an evolutionary process hypothesizes that mitochondria in animals and chloroplasts in plants arose when one prokaryotic organism engulfed another to form eukaryotic cells. Mitochondria on this account resulted from the incorporation of proteobacteria (probably purple non-sulfur bacteria) while chloroplasts resulted from the incorporation of cyanobacteria.¹ In the case of mitochondria, the host that previously had relied on glycolysis for energy acquired from the symbiont the ability to perform oxidative phosphorylation, thereby extracting significantly more energy (in the form of ATP) from its food sources. Although the symbiont retained its DNA and ribosomes and replicates through a process of division (which provides important evidence for its origin as an independent organism), it nonetheless has been highly integrated into the life of the host cell (e.g., some of its DNA has been transferred to the host's nucleus, and the symbiont has become dependent on proteins generated by the host and the conditions within it). Thus, even when separate components, originally capable of functioning independently, are brought together, over time they integrate their operation and become far less decomposable.

¹ This proposal, known as the endosymbiotic theory and commonly associated with Lynn Sagan (Margulis) (Sagan 1967), was initially suggested (in the case of chloroplasts in plants) by the German botanist Andreas Schimper in 1883, further developed by the Russian botanist Konstantin Mereschkowsky in 1905, and extended to mitochondria by Ivan Wallin in the 1920s. It was, however, largely ignored until electron microscopy provided further evidence for the similarities between chloroplasts and cyanobacteria and Stocking and Gifford demonstrated the occurrence of DNA in mitochondria. Margulis (1981) extended the endosymbiotic theory to explain the origins of flagella and cilia, but these proposals are harder to support since flagella and cilia do not have their own DNA.

More commonly, however, evolution proceeds by expanding, often by duplicating (e.g., through an extra round of cell division during development) existing parts of the mechanism and then allowing structural and function differentiation within the system. That is, component parts and operations arise from specialization of regions within an initially more homogenous part of the system. I will offer a sketch of such a specialization process in the context of introducing a mode of organization that has only been seriously explored in the past two decades but is increasingly recognized as extremely common in evolving systems, including living organisms—small-world networks. Small-world networks contain specialized components, but they are not nearly as independent in their operation as the sub-systems in a hierarchically organized nearly decomposable system.

Through much of the second half of the 20th century, mathematicians in the subfield of graph theory focused on three designs for networks: regular lattice networks, totally connected networks, and randomly connected networks. Two measures that are particularly useful in understanding network functioning are its clustering coefficient and its characteristic path length (Bullmore & Sporns 2009). The clustering coefficient (C_j) for a given node j measures the percentage of the possible connections between neighbors of j that are actually realized. A high clustering coefficient for a whole network (C) indicates that components are highly connected to those in their neighborhood, a necessary condition for them to cooperate in performing specialized processing. The characteristic path length (L) is the mean of the shortest path lengths between nodes, where path length is measured in the number of connections that must be traversed. A low L enables high integration of the activity of all units in the network. In a regular lattice, nearby units are connected, resulting in a high value for C , but a large number of connections typically must be traversed to go from a given unit to a specific other unit, resulting in a high value for L (Simon's nearly decomposable systems are similar in this respect to regular lattices). In a network in which connections between units are determined randomly, if there is a path between two units, there will likely be a relatively short path between them (hence a low value for L). However, there also will be little clustering (low C) as the connections are as likely to be between distant as nearby units.

Ideal conditions for information processing arise when C is high, enabling local clusters of units to work together, and L is low, enabling coordination between the clusters. This condition obtains in a fully connected network, but because all units are connected together there is little opportunity for specialization. It also obtains, though, in a form of organization known as *small-world*, a concept introduced by Milgram (1967) on the basis of an experiment in which individuals in the U.S. were requested to forward a letter they received to someone they knew personally with the goal of eventually getting it to a designated individual in Cambridge, MA. Surprisingly, on average those letters that made it to the destination did so in less than six steps. This gave rise to the popular notion that on average six degrees of separation exist between any two human beings. Watts and Strogatz (1998) provided a conceptual grounding for this phenomenon when they characterized a class of networks that result from randomly reconnecting a few of the connections in a regular lattice, providing rapid long-distance connection across a large network. After rewiring only a relative small percentage of the connections, they obtained networks that still exhibited high C but greatly reduced values for L .

In addition to characterizing such a network structure, Watts and Strogatz also demonstrated that a variety of real-world networks, including networks of movie actors linked by co-appearances, the electrical power-grid of the Western U. S., and the neural network of the nematode worm *Caenorhabditis elegans*, exhibited the characteristics of a small-world. Watts and Strogatz also examined the functional properties of small-world networks, showing how they allow for rapid spread of infectious diseases, enable efficient problem solving in cellular automata, and reduce the likelihood of cooperation in iterated prisoner dilemma games. Of particular interest, they examined coupled phase oscillators and demonstrated that synchronization occurred almost as fast in small-world networks as in fully-connected networks. They comment: “This result may be relevant to the observed synchronization of widely separated neurons in the visual cortex if, as seems plausible, the brain has a small-world architecture” (p. 442).

Subsequent research has indeed demonstrated a small-world architecture in the mammalian cortex, including the visual system. Sporns and Zwi (2004) developed connectivity matrices based on data from published studies of neuroanatomy—Felleman and van Essen’s (1991) study of the macaque’s visual cortex, Young’s (1993) study of the whole macaque cortex, and Scannell et al.’s (1999) study of the cat cortex. Sporns and Zwi show that these networks exhibit the properties of a small-world network: high C and low L .²

Watts and Strogatz introduced a procedure for creating small-world networks by rewiring regular lattice networks to replace a small percentage of local connections with more distant ones. Such a process is similar to the proposal above that even if a complex system developed from combining previously independent stable mechanisms, over time the components would develop increased connections with each other. But Rubinov, Sporns, van Leeuwen, and Breakspear (2009) offer an intriguing alternative account of how small-world organization might arise, one that fits with the process of specialization after replication described above. They start with the idea that neurons are not passive elements awaiting an input but endogenous oscillators. There is considerable evidence of spontaneous oscillations in ion concentrations across cell processes (as well as selective synchronized oscillations in response to oscillation of other neurons at a particular frequency), which can give rise to spontaneous spiking (Llinás 1988).³ To analyze how

² Other research has identified features characteristic of small-world organization, including reciprocal and clustered connections in local circuits in mammalian neocortex through multielectrode recording (Song, Sjöström, Reigl, Nelson, & Chklovskii 2005). Very recently tensor diffusion imaging has enabled MRI studies to identify fiber bundles in the human brain, and they too provide evidence of small-world organization (Hagmann et al. 2008).

³ The inquiries of Hodgkin, Huxley, and Katz in the 1940s and 1950s into the generation of action potentials revealed voltage dependent membrane conductances in potassium and sodium. Investigations conducted in invertebrates, especially mollusks (Kandel 1976), revealed numerous other voltage dependent conductances beyond those responsible for action potentials which eventually were found to figure in mammalian neurons. These discoveries, together with a developing appreciation of intracellular signaling processes and of the variety of neurotransmitters and neuromodulators operating across neurons, resulted in a much richer understanding of physiological processes in individual neurons. Among the critical discoveries was Llinás and Yarum’s (1981) finding of low threshold Ca^{++} conductances in neurons in the inferior olive in mammals, which is more active when the neuron is at its resting potential than when it is hyperpolarized. They showed that this conductance renders the neuron into a spontaneous oscillator or resonator (a cell that oscillates in response to endogenous oscillators at specific frequencies). The spontaneous spiking of inferior olive neurons generates synchronized activity in a wide population of Purkinje cells in cerebellar cortex, which then provide inhibitory synchronized inputs to cerebellar nuclear neurons. Their rebound responses activate motor neurons in a synchronized manner, generating the physiological tremor with a frequency of approximately 10 Hz. Llinás (1988) draws out broad consequences of the

such behavior could generate small-world organization, Rubinov et al. employ a neural network developed by Gong and van Leeuwen (2004) that uses a logistic map as the basis for the activation function so that the activation value of a unit at the next time step $f(x)$ depends upon its current activation (x):

$$f(x) = 1 - \alpha x^2$$

Under appropriate values of α (in their simulations they used 1.7), a unit can exhibit either regular oscillations or chaotic behavior. Gong and van Leeuwen then couple oscillators so that the activation of one depends upon not just its previous activation but that of each of the other units (j) to which it is connected:

$$x_{n+1}^i = (1 - \varepsilon) f(x_n^i) + \frac{\varepsilon}{M_i} \sum_{j \in B(i)} f(x_n^j)$$

Here ε determines how much a given unit is responsive to the previous activation of other units (in their simulations $\varepsilon=0.5$). Such a network will generate complex dynamic activity in which various units will spontaneously synchronize their oscillations and subsequently spontaneously desynchronize. Gong and van Leeuwen then rewire the network by, at each time step, pruning the connection between a unit and another with which it is least synchronized and adding a new connections to that unit with which it is not directly connected but with which it exhibits the greatest synchrony. From an initially random pattern of connectivity, such a network develops a small-world organization.

What is particularly intriguing about this model for the development of small-world organization is that it shows how structural organization can arise spontaneously relying on endogenous activity within a network. In the model, the network was not employed in any information processing task. Were the network employed in a task, though, the same processes of forming clusters integrated with other clusters in a small-world organization would be employed. The clusters formed could then subserve relevant information processing tasks. I turn to functional considerations in the next section.

The emerging conception of the structure of the brain that I have explored in this section departs significantly from Simon's conception of hierarchically organized decomposable system or from the sequential conception of a mechanism presented above. Rather than a brain region that constitutes a component in a mechanism being largely isolated from others except for receiving the products of processing by other components and directing its products to other components, a brain region will be interconnected with numerous others. As a result, its behavior at any given time will depend not just on its inputs and its internal constitution, but on the signals it receives from these other units. Such a perspective is suggestive of the holism that inspired skeptics of localization, but there is an important difference. The significance of clustering in small-world networks is that different parts of the network can still specialize for processing different kinds of information. Thus, the project of localization of function still makes sense, but it needs to be understood as involving components that are not operating in isolation but are highly sensitive to ongoing activity elsewhere in the mechanism.

spontaneous activity in the inferior olive: "The organization of the IO nucleus demonstrates that the oscillatory properties of single neurons, arising from a congruous set of electrical events, can activate a large number of neurons over a wide area. The ability to project oscillatory rhythms and to generate synchronous firing in large populations of cells may be one of the important properties of intrinsic electroresponsiveness" (p. 1661). Llinás also reviews a vast array of other ionic conductances found in many neurons in the brain.

III. Reconceiving localization from a dynamical mechanistic perspective

In introducing their review of evidence for a small-world architecture in cortex, Sporns and Zwi (2004) identify the dual role cortical connectivity plays in neural processing: “First, it is critical in generating functional specificity (i.e., information) of local cell populations and areas within cortex. Second, it allows the integration of different sources of information into coherent behavioral and cognitive states” (p. 146). I will briefly explore how both of these are manifest in brain activity before turning to recent research using new approaches to fMRI that suggests how large-scale networks in the brain coordinate their functional activity through synchronization. I will then focus on the implications of these perspectives for the project for understanding the type of mechanism involved in cognition and how it supports referring to operations localized within such a mechanism.

The first functional consequence of small-world organization, functional specificity, results from the clustering of units into local subsystems. Even though many other components modulate the behavior of units in a local cluster, it is still possible, in part by focusing on the task the overall mechanism is performing, to identify and refer to the specific processing occurring in the local cluster. In particular, the traditional strategies for localization remain informative in characterizing the processing in local areas. MT is appropriately characterized as playing a role in motion detection, as revealed in the range of techniques (lesion, recording, stimulation) that indicate its importance to that task (Britten, Shadlen, Newsome, & Movshon 1992). These strategies of referring to localized operations, however, should not be viewed as completing inquiry, but only as a probative beginning. Once MT is identified as contributing to motion perception, researchers need to understand how its contribution is modulated by activity elsewhere in the brain.

The second functional consequence of small-world organization, integration into coherent global states, is illustrated at a coarse grain by the role oscillations in thalamic neurons play in producing global states such as attentive awakesness, drowsiness, and sleep, which modulate processing in many local circuits, including MT. When thalamic neurons are slightly depolarized, they oscillate at 10 Hz, but when hyperpolarized they oscillate at 6 Hz. Given the loops between thalamic and cortical neurons, the thalamus is able to differentially modulate cortical activity depending on which of the two states it is in. Firing at 10 Hz, thalamic neurons act as relay elements communicating information to and from cortical neurons. Firing at 6 Hz, on the other hand, they serve to entrain cortical neurons, generating alpha rhythms or the pattern of spindling exhibited in early stages of sleep (Steriade & Llinás 1988). Motion detection in MT, as well as many other cortical operations, is modulated by the frequency of such oscillations.

Integration is achieved, not just at such a global level, but in more local circuits within the brain. One way these circuits have been revealed is through a relatively recent way of interpreting fMRI results. The BOLD (blood oxygen level dependent) signal employed in fMRI research registers the oxygen concentrations in the brain within areas that can be as small as 2 mm. In the most familiar applications of fMRI, changes in the BOLD signal are correlated with stimuli presented to or tasks performed by a person in the scanner (Posner & Raichle 1994). Well before this approach to fMRI studies was developed, however, researchers measured blood flow and

brain activity when no tasks were being performed. Using the xenon 133 inhalation technique to measure regional cerebral blood flow, Ingvar (1975) showed high levels of frontal activity when subjects were at rest and proposed that this activity reflected “undirected, spontaneous, conscious mentation, the ‘brain work,’ which we carry out when left alone undisturbed.” Studies on resting brains were temporarily superseded once it became possible to measure blood flow during performance of cognitive tasks. They were revived, though, as Raichle and his colleagues tried to understand why a range of brain areas including the precuneus and posterior cingulate cortex, inferior parietal cortex, left dorsal lateral prefrontal cortex, a medial frontal strip that continued through the inferior anterior cingulate cortex, left inferior frontal cortex, left inferior frontal gyrus, and the amygdala seemed to be regularly deactivated in task situations. Raichle and his collaborators proposed that these areas that are more active in the absence of task requirements and deactivated in task conditions constitute a *default network* (Raichle et al. 2001). Raichle and Mintun (2006) provide further clarification of what they intend by a default network, “We concluded these regional decreases, observed commonly during task performance, represented the presence of functionality that was ongoing in the resting state and attenuated only when resources were temporarily reallocated during goal-directed behaviors, hence our original designation of them as default functions.”

Characterizing these areas as a network requires more than establishing that they all exhibit less blood flow in task conditions—one must establish that reduced activations are somehow coordinated. The avenue to doing this involved applying an innovative fMRI analysis pioneered by Biswal, Yetkin, Haughton, and Hyde (1995), who showed how to use fMRI time series data to examine the micro-temporal dynamics of blood flow. These investigators obtained BOLD signal values every 250 msec. after a subject performed a hand movement and identified low frequency (less than 0.1 Hz) synchronized fluctuations in the left and right motor areas. Synchronizing oscillations requires the communication of a signal between the independent oscillators, indicative of a functionally integrated network of brain areas. Cordes et al. (2000) used a similar approach to identify networks exhibiting synchronized blood flow within regions activated in sensorimotor, visual, auditory, and expressive and receptive language tasks. The technique of identifying networks by findings areas in which fluctuations in the BOLD signal are correlated is known as functional connectivity MRI (fcMRI).

Using fcMRI, Greicius, Krasnow, Reiss, and Menon (2003) demonstrated, in the context of a working memory task, “a cohesive, tonically active, default mode network” (p. 256) (cohesiveness refers to synchronized activity). It is noteworthy that synchronized oscillations in the default network are maintained even in task conditions, which results in lower overall activity in this network that is nonetheless still synchronized. One can then compare how oscillations in the regions constituting the default network relate to those in other brain areas. Fox et al. (2005) conducted fcMRI analysis of imaging results during an attention demanding task and found synchronized oscillations in a set of areas distinct from the default network, including the intraparietal sulcus, frontal eye field, middle temporal region, supplementary motor areas, and the insula. Oscillations in these areas were synchronized with each other, but not with areas in the default network. Fox et al. refer to this as anticorrelation and comment:

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 (p. 9677).

Subsequently, researchers have distinguished six anticorrelated networks (Mantini, Perrucci, Del Gratta, Romani, & Corbetta 2007).⁴

Synchronization of processing in different brain regions has also been demonstrated in electrophysiological investigations measuring local field potentials in multi-cell recording or in EEG studies. The best known of these are the synchronized oscillations in the range of 40 Hz found in brain areas involved in representing different features of visual stimuli, which has been proposed to explain how the brain binds together different features of a visual stimulus as they are processed in different regions of the visual system (Gray & Singer 1989). The synchronized oscillations found with fMRI are at a much lower frequency (< 0.1 Hz) than those typically reported with local field potentials or EEG (usually in the 1-80 Hz range). However, oscillations at different frequencies may be related. Laufs et al. (2003) identified significant correlations between alpha (8-12 Hz) and beta oscillations (12-30 Hz) in EEG and BOLD fluctuations in specific brain networks. Mantini, Perrucci, Del Gratta, Romani, and Corbetta (2007) related the six anticorrelated networks they found in fMRI with power fluctuations in all EEG bands, and concluded: "Each brain network was associated with a specific combination of EEG rhythms, a neurophysiological signature that constitutes a baseline for evaluating changes in oscillatory signals during active behavior" (p. 13170). For example, the default network showed positive correlations with alpha and beta band oscillations while the attention network exhibited negative correlations.

Going beyond mere correlations, a number of researchers have explored the potential causal relations between oscillations at different frequencies. One relevant finding is that in the cortex of mammals, the amplitude (power density) of EEG oscillations is inversely proportional to their frequency ($1/f$). In addition, the phase of lower-frequency oscillations seems to modulate the amplitude of those at higher frequencies, which results in a nesting relation between the frequency bands. Lakatos et al. (2005) refer to this as the "oscillatory hierarchy hypothesis" (see also Canolty et al. 2006.) Since oscillations at lower frequencies tend to synchronize over more widely distributed brain regions than those at higher frequencies (Buzsáki & Draguhn 2004), this suggests that the relatively slow oscillations detected with fMRI might serve to modulate specific processing, reflected in the higher frequency oscillations, in local brain areas:

The power density of EEG or local field potential is inversely proportional to frequency (f) in the mammalian cortex. This $1/f$ power relationship implies that perturbations occurring at slow frequencies can cause a cascade of energy dissipation at higher frequencies and that widespread slow oscillations modulate faster local events. These properties of neuronal oscillators are the result of the physical architecture of neuronal networks and the limited speed of neuronal communication due to axon conduction and synaptic delays. Because most neuronal connections are local, the period of oscillation is constrained by the size of the neuronal pool engaged in a given cycle. Higher frequency

⁴ Achard, Salvador, Whitcher, Suckling, and Bullmore (2006) applied wavelet analysis to time series data derived from fMRI to estimate frequency-dependent correlation matrices for the whole human brain, and identified a small-world topology of sparse connections. They identified a small number of highly connected hubs located in multimodal association regions.

oscillations are confined to a small neuronal space, whereas very large networks are recruited during slow oscillations” (Buzsáki & Draguhn 2004, p. 1926)

In this section so far I have focused on functional connectivity without linking the activity in any processing area to specific cognitive or mental operations. To establish the latter connections, investigators rely on the techniques employed in more traditional neural imaging—identifying what types of cognitive tasks elicit changes in blood flow in particular brain regions. These findings, however, are now understood within the context of the networks in which the areas function. In one of the studies that first drew attention to the decrease in activity in the default network in tasks conditions, Andreasen et al. (1995) compared blood flow during an episodic memory task to that produced both during a resting state and a semantic memory task. They found that both the resting state and episodic memory tasks generated higher BOLD levels in regions that were later identified as constituting the default network than did the semantic memory task. In an attempt to figure out what was eliciting activity in these areas during the resting state, the researchers queried their subjects about what they were doing during while lying in the scanner and found that they reported being engaged in “a mixture of freely wandering past recollection, future plans, and other personal thoughts and experiences”—activities that plausibly draw upon episodic memory. This suggestion has been developed further by Buckner, Andrews-Hanna, & Schacter, (2008), who link mental activity during the resting state to what Antrobus, Singer, Goldstein, & Fortgang (1970) characterized as *mind-wandering* and “hypothesize that the fundamental function of the default network is to facilitate flexible self-relevant mental explorations—simulations—that provide a means to anticipate and evaluate upcoming events before they happen” (p. 2).

Although the approach to identifying what cognitive operations are performed in brain regions begins in a manner comparable to more traditional approaches to neuroimaging, the strategies for further research go considerably beyond the traditional localizationist project. First, activities such as mind-wandering and self-relevant mental explorations are not being localized in a single brain region, but in a network of regions, and additional work is required to determine the specific operations performed by the different components of the network. The synchronization within the network then must be considered in characterizing each of these operations. Second, the coordination of the dynamic activities in different networks whose activity is anticorrelated also becomes an important issue. One of the interests in mind-wandering is that it often leads to failures to notice or attend to other stimuli. This may well be manifest in competitive interactions between different networks that modulate cognitive processing within each network. The ability of network dynamics to alter processing has been demonstrated in a more local manner in an fMRI study which showed that the ongoing endogenous oscillation in motor areas could account for the variability in the BOLD signal recorded during motor behavior (button pressing) and moreover accounted for most of the variation in the strength of button presses (Fox, Snyder, Vincent, & Raichle 2007).

In the previous section I discussed evidence for structural small-world organization in the brain whereas in this section I have discussed comparable evidence for functional organized networks. It is, of course, possible that these are unrelated, but that does not seem to be the case. Rather, the functionally characterized default network maps closely onto a structural network identified through diffusion tensor imaging tractography (Greicius, Supekar, Menon, & Dougherty 2009). In

addition to reviewing how, on a variety of measures, the small-world architecture found through structural and functional oscillations at low-frequencies tend to correspond, Bullmore and Sporns (2009) note how, at short-temporal intervals there are dynamic functional changes, especially involving high-frequency oscillations, that are not reflected in changed structure. Nonetheless, the structural organization is likely to affect the faster dynamic processing and Bullmore and Sporns identify several questions to address in further research:

We need to understand more about the non-stationarity or metastability of brain functional networks. How does functional network topology change over time? Do functional networks exist in a dynamically critical state at some or all frequency intervals? What constraints on the itinerancy of network dynamics are imposed anatomically and how does the long-term history of functional activity in a network feed back on the development and remodelling of the anatomical connections between nodes?

Understanding processing in networks such as I have been describing requires adopting a different conception of a mechanism than one involving the sequential operation of nearly decomposable parts. Instead, it requires a perspective in which researchers characterize the components of mechanisms in terms of variables and represent the changes in values of these variables in terms of (differential) equations. One can then analyze the functioning of mechanisms in terms of the patterns of change over time in properties of their parts and operations, generating *dynamic mechanistic explanations* (Bechtel & Abrahamsen in press). In developing equations to characterize the relations between operations performed by various parts, researchers are not limited to sequential operations, but are able to accommodate any dependency relation between parts, including those whereby activity elsewhere in a mechanism modulates the operation of a given part. The behavior of a small-world network can often be extremely complex, and only via simulation and the invocation of tools such as those of dynamical systems theory (e.g., analyses of attractors in multidimensional phase spaces) is often required to anticipate how such mechanisms will behave. But unlike the dynamicists who oppose localization, researchers pursuing dynamical mechanistic accounts are still describing mechanisms. Even if the parts and especially the operations are not fixed but contextually modulated, they can still be characterized (e.g., in the equations developed to model their behavior).

The route to an adequate dynamic mechanistic explanation begins with strategies to localize operations in parts, where these are viewed as nearly-decomposable constituents and assumed to operate in roughly a sequential manner. However, the results of such efforts only serve to characterize the parts and operations of the mechanism *to a first approximation* (Bechtel & Richardson 1993). Especially in biology, researchers increasingly expect that the parts and operations will be affected in a variety of ways by activity elsewhere in the mechanism, and as these additional interactions are identified, they need to be characterized in equations which can be used to understand the resulting complex dynamics.

Conclusions

The endeavor of localizing cognitive operations in the brain has traditionally been approached from the perspective of mechanistic explanation in which the regions of the brain are independent parts performing distinct information-processing operations. But research on both

the structure and functional processes in the brain suggests a very different picture, in which the brain is organized as a small-world network. This still allows for local clusters to specialize in performing different operations, but the components are also regularly modulated by activity elsewhere in the brain. Such interactions do not entail a holism that defeats localization, but a perspective of a dynamically organized and modulated system in which the component operations are contextually modified by activity elsewhere. The initial assumption of a hierarchically organized nearly decomposable system must be modified, in the course of research, to take into account the sorts of modulation that occurs in a small-world network. These modulating effects are represented in equations whose solution reveals the temporal dynamics of the system. Within the context of such dynamic mechanistic explanations, localizing functions still plays an important role, but it is only an initial step and the resulting localization claims must be modified as researchers recognize how the whole mechanism functions in time by modifying the operation of its own constituents. The reference to parts and operations in the brain needs to be couched within a dynamic perspective in which both the parts and operations change through time in complex ways.

References

- Achard, Sophie, Salvador, Raymond, Whitcher, Brandon, Suckling, John, and Bullmore, Ed 2006. 'A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs', *Journal of Neuroscience* 26: 63-72
- Andreasen, Nancy C., O'Leary, Daniel S., Cizadlo, Ted, Arndt, Stephan, Rezai, Karim, Watkins, G. Leonard, Ponto, Laura L., and Hichwa, Richard D. 1995. 'Remembering the past: two facets of episodic memory explored with positron emission tomography', *American Journal of Psychiatry* 152: 1576-1585
- Antrobus, John S., Singer, Jerome L., Goldstein, S., and Fortgang, M. 1970. 'Mindwandering and cognitive structure', *Transactions of the New York Academy of Sciences* 32: 242-252
- Bechtel, William 2008. *Mental mechanisms*. London: Routledge
- Bechtel, William and Abrahamsen, Adele 2005. 'Explanation: A mechanist alternative', *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 421-441
- Bechtel, William and Abrahamsen, Adele in press. 'Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science', *Studies in History and Philosophy of Science Part A*
- Bechtel, William and Richardson, Robert C. 1993. *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton, NJ: Princeton University Press
- Biswal, Bharat, Yetkin, F. Zerrin, Haughton, Victor M., and Hyde, James S. 1995. 'Functional connectivity in the motor cortex of resting human brain using echo-planar MRI', *Magnetic Resonance in Medicine* 34: 537-541
- Britten, Kenneth H., Shadlen, Michael N., Newsome, William T., and Movshon, J. Anthony 1992. 'The analysis of visual motion: A comparison of neuronal and psychophysical performance', *The Journal of Neuroscience* 12: 4745-4765
- Broca, Paul 1861. 'Remarques sur le siège de la faculté du langage articulé, suivies d'une observation d'aphémie (perte de la parole)', *Bulletin de la Société Anatomique* 6: 343-357
- Brodmann, Korbinian 1909/1994. *Vergleichende Lokalisationslehre der Grosshirnrinde* (L. J. Garvey, Trans.). Leipzig: J. A. Barth

- Buckner, Randy L., Andrews-Hanna, Jessica R., and Schacter, Daniel L. 2008. 'The Brain's default network', *Annals of the New York Academy of Sciences* 1124: 1-38
- Bullmore, Ed and Sporns, Olaf 2009. 'Complex brain networks: graph theoretical analysis of structural and functional systems', *Nature Reviews Neuroscience* 10: 186-198
- Buzsáki, György and Draguhn, Andreas 2004. 'Neuronal oscillations in cortical networks', *Science* 304: 1926-1929
- Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., Berger, M. S., Barbaro, N. M., and Knight, R. T. 2006. 'High gamma power is phase-locked to theta oscillations in human neocortex', *Science* 313: 1626-1628
- Cordes, Dietmar, Haughton, Victor M., Arfanakis, Konstantinos, Wendt, Gary J., Turski, Patrick A., Moritz, Chad H., Quigley, Michelle A., and Meyerand, M. Elizabeth 2000. 'Mapping functionally related regions of brain with functional connectivity MR imaging', *American Journal of Neuroradiology* 21: 1636-1644
- Felleman, Daniel J. and van Essen, David C. 1991. 'Distributed hierarchical processing in the primate cerebral cortex', *Cerebral Cortex* 1: 1-47
- Flourens, J. P. M. 1824. *Recherches Expérimentales sur les Propriétés et les Fonctions du Système Nerveux dans les Animaux Vertébrés*. Paris: Crevot
- Fox, Michael D., Snyder, Abraham Z., Vincent, Justin L., Corbetta, Maurizio, Van Essen, David C., and Raichle, Marcus 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, Michael D., Snyder, Abraham Z., Vincent, Justin L., and Raichle, Marcus E. 2007. 'Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior', *Neuron* 56: 171-184
- Fritsch, Gustav Theodor and Hitzig, Eduard 1870. 'Über die electrische Erregbarkeit des Grosshirns', *Archiv für Anatomie und Physiologie*: 300-332
- Gong, Pulin and van Leeuwen, Cees 2004. 'Evolution to a small-world network with chaotic units', *Europhysics Letters* 67: 328-333
- Gray, Charles M. and Singer, Wolf 1989. 'Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex', *Proceedings of the National Academy of Sciences of the United States of America* 86: 1698-1702
- Greicius, Michael D., Krasnow, Ben, Reiss, Allan L., and Menon, Vinod 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, Michael D., Supekar, Kaustubh, Menon, Vinod, and Dougherty, Robert F. 2009. 'Resting-state functional connectivity reflects structural connectivity in the default mode network', *Cerebral Cortex* 19: 72-78
- Hagmann, Patric, Cammoun, Leila, Gigandet, Xavier, Meuli, Reto, Honey, Christopher J., Wedeen, Van J., and Sporns, Olaf 2008. 'Mapping the structural core of human cerebral cortex', *PLoS Biology* 6: e159
- Hubel, David H. and Wiesel, Torsten N. 1962. 'Receptive fields, binocular interaction and functional architecture in the cat's visual cortex', *Journal of Physiology* 160: 106-154
- Hubel, David H. and Wiesel, Torsten N. 1968. 'Receptive fields and functional architecture of monkey striate cortex', *Journal of Physiology* 195: 215-243

- Ingvar, David H. 1975. 'Patterns of brain activity revealed by measurements of regional cerebral blood flow', in David H. Ingvar & Niels A. Lassen (eds.), *Brain work : the coupling of function, metabolism, and blood flow in the brain: Proceedings of the Alfred Benzon Symposium VIII, Copenhagen, 26-30 May 1974, held at the premises of the Royal Danish Academy of Sciences and Letters, Copenhagen*. New York: Academic
- Kandel, Eric R. 1976. *Cellular basis of behavior: An introduction to behavioral neurobiology*. San Francisco: W. H. Freeman
- Lakatos, Peter, Shah, Ankoor S., Knuth, Kevin H., Ulbert, Istvan, Karmos, George, and Schroeder, Charles E. 2005. 'An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex', *Journal of Neurophysiology* 94: 1904-1911
- Lashley, Karl S. 1950. 'In search of the engram', *Symposia of the Society for Experimental Biology* IV. Physiological Mechanisms in Animal Behaviour: 454-482
- Laufs, Helmut, Krakow, K., Sterzer, P., Eger, E., Beyerle, A., Salek-Haddadi, A., and Kleinschmidt, A. 2003. 'Electroencephalographic signatures of attentional and cognitive default modes in spontaneous brain activity fluctuations at rest', *Proceedings of the National Academy of Sciences of the United States of America* 100: 11053-11058
- Llinás, Rodolfo R. 1988. 'The intrinsic electrophysiological properties of mammalian neurons: Insights into central nervous system function', *Science* 242: 1654-1664
- Llinás, Rodolfo R. and Yarom, Y. 1981. 'Properties and distribution of ionic conductances generating electroresponsiveness of mammalian inferior olivary neurones in vitro', *The Journal of Physiology* 315: 569-584
- Lorente de Nó, Rafael 1938. 'Analysis of the activity of the chains of internuncial neurons', *Journal of Neurophysiology* 1: 207-244
- Machamer, Peter, Darden, Lindley, and Craver, Carl F. 2000. 'Thinking about mechanisms', *Philosophy of Science* 67: 1-25
- Mantini, D., Perrucci, M. G., Del Gratta, C., Romani, G. L., and Corbetta, Maurizio 2007. 'Electrophysiological signatures of resting state networks in the human brain', *Proceedings of the National Academy of Sciences* 104: 13170-13175
- Margulis, Lynn 1981. *Symbiosis in cell evolution: Life and its environment on the early Earth*. San Francisco: W. H. Freeman
- Milgram, Stanley 1967. 'The small world problem', *Psychology Today* 2: 60-67
- Milner, A. David and Goodale, Melvyn G. 1995. *The visual brain in action*. Oxford: Oxford University Press
- Munk, Hermann 1881. *Über die Funktionen der Grosshirnrinde*. Berlin: A. Hirschwald
- Posner, M. I. and Raichle, M. E. 1994. *Images of Mind*. San Francisco: Freeman
- Raichle, Marcus E., MacLeod, Ann Mary, Snyder, Abraham Z., Powers, William J., Gusnard, Debra A., and Shulman, Gordon L. 2001. 'A default mode of brain function', *Proceedings of the National Academy of Sciences of the United States of America* 98: 676-682
- Raichle, Marcus E. and Mintun, Mark A. 2006. 'Brain work and brain imaging', *Annual Review of Neuroscience* 29: 449-476
- Rubinov, Mikail, Sporns, Olaf, van Leeuwen, Cees, and Breakspear, Michael 2009. 'Symbiotic relationship between brain structure and dynamics', *BMC Neuroscience* 10: 55
- Sagan, Lynn 1967. 'On the origin of mitosing cells', *Journal of Theoretical Biology* 14: 225-274, IN221-IN226

- Scannell, J. W., Burns, G. A. P. C., Hilgetag, C. C., O'Neil, M. A., and Young, Malcolm P. 1999. 'The connectional organization of the cortico-thalamic system of the cat', *Cerebral Cortex* 9: 277-299
- Simon, Herbert A. 1962. 'The architecture of complexity: hierarchic systems', *Proceedings of the American Philosophical Society* 106: 467-482
- Song, Sen, Sjöström, Per Jesper, Reigl, Markus, Nelson, Sacha, and Chklovskii, Dmitri B. 2005. 'Highly nonrandom features of synaptic connectivity in local cortical circuits', *PLoS Biology* 3: e68
- Sporns, Olaf and Zwi, Jonathan D. 2004. 'The small world of the cerebral cortex', *Neuroinformatics* 2: 145-162
- Steriade, Mircea and Llinás, Rodolfo R. 1988. 'The functional states of the thalamus and the associated neuronal interplay', *Physiological Reviews* 68: 649-742
- Ungerleider, Leslie G. and Mishkin, Mortimer 1982. 'Two cortical visual systems', in D. J. Ingle, M. A. Goodale & R. J. W. Mansfield (eds.), *Analysis of Visual Behavior* (pp. 549-586). Cambridge, MA: MIT Press
- Uttal, William R. 2001. *The new phrenology: The limits of localizing cognitive processes in the brain*. Cambridge, MA: MIT Press
- van Essen, David C. and Gallant, Jack L. 1994. 'Neural mechanisms of form and motion processing in the primate visual system', *Neuron* 13: 1-10
- van Orden, Guy C. and Paap, Kenneth R. 1997. 'Functional neural images fail to discover the pieces of the mind in the parts of the brain', *Philosophy of Science* 64: S85-S94
- van Orden, Guy C., Pennington, Bruce F., and Stone, Gregory O. 2001. 'What do double dissociations prove? Inductive methods and isolable systems', *Cognitive Science* 25: 111-172
- Watts, Duncan and Strogatz, Steven 1998. 'Collective dynamics of small worlds', *Nature* 393
- Wimsatt, William C. 1986. 'Forms of aggregativity', in A. Donagan, N. Perovich & M. Wedin (eds.), *Human nature and natural knowledge* (pp. 259-293). Dordrecht: Reidel
- Young, Malcolm P. 1993. 'The organization of neural systems in the primate cerebral cortex', *Proceedings of the Royal Society of London. Series B: Biological Sciences* 252: 13-18