

Network Organization in Health and Disease: On Being a Reductionist and a Systems Biologist Too

Author

W. Bechtel

Affiliation

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

Key words

- mechanistic explanation
- dynamical systems theory
- dynamic mechanistic explanation
- graph theory
- small worlds
- subgraphs
- mental disorders

Abstract

Whereas the challenge for traditional mechanistic science was to identify parts and operations, the current challenge in many fields of biology is to understand how the many parts of mechanisms are organized in networks and their operations coordinated across these networks. This paper explores how tools from graph theory are enabling analysis of organization at both macro-

and micro-levels. In applying these approaches to brain regions, systems neuroscientists are identifying both small-world organization with hubs at the macro-scale and frequently occurring subgraphs that link specific brain regions at a more micro-scale. This has led to the discovery of networks in which activity in multiple brain regions exhibits coherent oscillations and demonstrations that these networks are disrupted in various mental disorders.

Introduction

In multiple fields of biology researchers are recognizing that the phenomena of interest exhibit complex dynamics. Rather than functioning at a steady state, organisms exhibit oscillations at a wide range of frequencies from milliseconds to years. Coordination of behavior often involves synchronizing these oscillations and entraining them to external oscillations. For example, organisms from cyanobacteria to humans exhibit circadian (approximately 24-h) rhythms that are synchronized between cells in multi-celled organisms and entrained to the day-night cycle on our planet. Moreover, disruptions in these oscillations can produce a wide range of disease conditions. Understanding oscillatory behavior has required a fundamental shift in the explanatory strategies of biologists from one that focuses primarily on the parts and operations of biological mechanisms to one that emphasizes system organization and invokes tools such as graph theory to understand how these mechanisms are organized and linked to one another and dynamical systems theory to understand normal and pathological functioning of these mechanisms. I explore the transformations that are occurring in biology, drawing upon examples from recent neuroscience and identifying some of the insights that have been applied to mental disorders.

Biological research for the past several centuries has searched for mechanisms to explain various phenomena. Mechanistic research has often followed the script of identifying the putative mechanism for a given phenomenon and then decomposing that mechanism into its component parts and operations [1,2]. The behavior of mechanisms depends not just on their parts and operations, but also on how they are organized. With few exceptions, biologists tended to downplay the importance of organization, assuming the decomposed operations acted sequentially. The recent rise of systems biology provides an important corrective to this downplaying of organization since a systems focus leads biologists to address the more complicated and complex ways mechanisms are organized. Understanding complex organization, in which operations are not just organized sequentially, however, exceeds the ability of scientists to simulate the functioning of the mechanism in their heads. Instead, they must turn to computational models and the resources of dynamical systems theory. To mark the contrast with more traditional mechanistic accounts, I refer to the resulting explanations as dynamic mechanistic explanations [3]. Sometimes those advocating a systems approach present it as supplanting the reductionistic approach of traditional mechanistic science.

Bibliography

DOI <http://dx.doi.org/10.1055/s-0033-1337922>
 Pharmacopsychiatry 2013; 46 (Suppl. 1): S10–S21
 © Georg Thieme Verlag KG
 Stuttgart · New York
 ISSN 0176-3679

Correspondence

Prof. W. Bechtel

Department of
 Philosophy-0119
 University of California
 San Diego
 9500 Gilman Drive
 La Jolla
 CA 92093-0119
 USA
 Tel.: +1/858/822 4461
 Fax: +1/858/534 8566
 bill@mechanism.ucsd.edu

However, systems models that are advanced by theorists that divorce systems explanations from reductionistic discoveries of the parts and operations of mechanisms turn out to offer only speculative, how-possibly explanations.¹ Dynamic mechanistic explanations, in contrast, draw upon the resources of reductionistic research identifying parts and operations. When the variables and parameters in the computational models are grounded in identified parts and operations, the computational models are more than possible models – they are models of the mechanism that is thought to exist in the world. They show how the mechanism, as characterized in the explanatory account, would behave under a variety of conditions. Dynamic mechanistic explanations integrate reductionistic and systems approaches.

The distinctive contribution of systems biology is to recompose mechanisms (reversing the decomposition employed in reductionistic inquiry) and then, in turn, to situate mechanisms in specific contexts in which they function. Recomposing requires locating parts and operations within an organization. One aspiration of many systems biologists is to discover general principles of organization, which they sometimes refer to as *laws*, that characterize how differently constituted mechanisms that implement the same mode of organization will behave. Indeed, research in graph theory as it has been applied to biological mechanisms has identified a number of organizational principles. While these organizational principles will not function in the same way as laws, classically understood,² the prospect of identifying and understanding the implications of different patterns of organization offers promise of a more systematic inquiry than would occur if the effects of organization in each mechanism had to be analyzed *de novo*.³ In this paper I focus on how representing the organization of mechanisms in graphs and applying to them graph-theoretic analyses is providing a valuable resource for recomposing mechanisms, explaining their behavior, and understanding how breakdowns in organization result in pathology. To set this up, I will briefly introduce recent accounts of mechanistic explanation in philosophy of science and how they apply to major exemplars of explanation in neuroscience. I then in the section “Representing System Organization in Graphs” examine recent contributions to understanding organization through graph theory analyses of networks. In the section “Applying Graph Representations to Neural Mechanisms” I explore how these ideas have been applied to the brain and in the section “Altered Graphs, Disrupted Oscillations, and Mental Disorders” show how these have led to new perspectives on mental disorders.

¹The phrase how possibly explanations was introduced by Craver [4], who argues against the explanatory status of mathematical models. Even such speculative models can, nonetheless, be extremely useful if they serve as guides to identifying the components of the actual mechanisms functioning in the world.

²Laws were classically characterized as stating fundamental and necessary principles that were assumed to provide the explanation for the behavior of components. By filling in initial values for the variables in the equation representing the law one can derive how a component will behave. While mathematical representations are used to assess the mode of organization, what is most frequently “derived” from the analysis is a qualitative statement of how a certain type of mechanism will behave.

³Simon argued that if nature were not nearly decomposable into components whose contributions could be separately understood, it would not be intelligible to humans. A similar comment might be made about organization – if there were not an identifiable set of organizational arrangements that had regular effects on the mechanisms in which they were employed, the behavior of mechanisms might be unintelligible to humans.

A Brief Introduction to Mechanistic Explanation and Neuroscience



Ironically, while the search for mechanisms has been a dominant theme in biology, including neuroscience, over the past two centuries, it was been largely ignored in philosophy of science until the past 2 decades. Philosophical accounts of explanation emphasized laws and the derivation of explanations from laws by supplying initial conditions [5]. Though this approach worked well for some fields of physics, it did not fit actual research in biology in which there are few laws to invoke and yet many proposed explanations.⁴ Following the lead of biologists who regularly use the term mechanism, several philosophers have recently undertaken the challenge of characterizing what biologists mean by the term and showing how accounts of mechanism are explanatory [1,6,7]. Although vocabulary differs, these accounts concur in viewing mechanisms as involving parts performing operations organized so that together they generate the phenomenon for which an explanation is sought.

Research on fermentation in the late 19th and early 20th centuries illustrates the process of investigating mechanisms. Although Pasteur had viewed fermentation as a capacity of a whole living cell, and thus not subject to further explanation, Buchner's [8] discovery the fermentation occurred in pressed yeast juice when sugar was added led researchers to look within cells to identify potential intermediates and the enzymes that catalyzed each reaction. This initially proved challenging since the most plausible candidate intermediates, 3-carbon sugars such as methylglyoxal, did not themselves undergo fermentation when added to yeast-juice preparations as they should if they were intermediates. Once Embden, Deuticke, and Kraft [9] determined that the intermediates were phosphorylated forms of these compounds, though, numerous intermediates were quickly identified and within the decade they were organized into a sequential pathway. Glycolysis quickly became an exemplar of successful biochemical explanation of biological phenomena (for details, see [10]). Similar examples can be identified in many areas of neuroscience. Lesions studied in animals and autopsies of patients with deficits in the late 19th century led researchers concerned with vision to focus on the occipital lobe. Hubel and Wiesel's [11,12] research revealed that neurons in part of this region responded to edges but also made clear that many other areas of the brain must be involved in vision. By the time of Felleman and van Essen's [13] review over 30 brain regions were identified as involved in processing visual stimuli and the distinctive contributions of several had been determined (e.g., different types of motion in MT and MSTd) (for an analysis of this history, see [14]).

The process of decomposing mechanisms into their parts and operations can be iterated and in many cases further decomposing the parts initially identified into their components fills out the explanations. Research on learning and memory provides a useful example. Those researchers seeking an explanation of how memories are acquired sought to identify a brain region that was involved. In rodent research, a combination of single-cell recording studies and lesion studies focused attention on the hippocampus and the medial temporal lobe more generally. For

⁴These explanations may invoke laws from physics or chemistry, such as Olm's law, but these laws alone do not suffice as an explanation. Rather, what explains the biological phenomenon is the mechanism composed of the components characterized by laws.

example, focusing on rat navigation O'Keefe and Dostrovsky [15] identified neurons in the hippocampus that responded whenever a rat traversed a particular part of an enclosure. Drawing upon Tolman's idea that rats solve navigation problems by employing a cognitive map, O'Keefe and Nadel [16] concluded that the hippocampus generated cognitive maps. During the same period, Bliss and Lømo [17] discovered the phenomenon of long-term potentiation in hippocampal neurons, a process by which these neurons exhibited increased generation of action potentials after receiving a tetanus (for a detailed analysis of this research, see [18]). This finding was subsequently integrated with the research identifying cognitive maps in the hippocampus with the suggestion that long-term potentiation was responsible for the construction of maps in the hippocampus. Long-term potentiation is a process that occurs at synapses and subsequent research has revealed many of the key molecular processes by which the responsiveness of a post-synaptic cell to neurotransmitters is altered. In this case, research that started with behavioral studies has given rise to identification of the hippocampus as the organ involved, a process occurring at synapses in that organ, and ultimately to chemical reactions within the post-synaptic neuron that result in additional receptors being incorporated within the membrane at the synapse [19]. Although in principle researchers could choose to decompose the system further, as Machamer, Darden, and Craver emphasize, mechanistic research typically bottoms out at a level of decomposition at which the operations of parts can be seen as accounting for the key behavior with which research began.

The study of memory mechanisms in the hippocampus is just one of numerous examples of research into neural mechanisms underlying mental activity. The cognitive revolution that began in the 1950s was grounded on the idea that mental capacities could be explained by identifying the information processing mechanisms responsible for them. Lacking tools for specifically identifying neural substrates that performed different information processing operations, the components of mechanisms for different cognitive phenomena were primarily identified functionally in terms of the operations they performed. For example, drawing upon the different sorts of deficits found in patient populations, memory researchers differentiated memory systems for episodic memory, semantic memory, and various types of implicit memory, albeit with little success in identifying the component operations in each system. Lesion studies also provided one of the first tools for relating these systems to the brain; subsequently these have been complemented with EEG, PET, fMRI, and MEG. Initially these tools were applied to localize whole mechanisms, but more recently they have been used to discover components of responsible networks that contribute differentially to the phenomenon.

Insofar as research on mechanisms focuses on taking mechanisms apart to identify their component parts and operations, it exemplifies a reductionistic approach. However, all accounts of mechanistic explanation recognize that the parts and operations must be organized for a mechanism to generate any phenomenon beyond what individual parts perform. Locating components in an organized system is what I am referring to as recomposition. Accounts such as Machamer, Darden, and Craver's (2000), however, focus principally on sequential organization where, as in an assembly line, each component carries out an operation on the product of the previous operation. Although sequential organization is easy for humans to understand and invoke in design, biological mechanisms were not designed to

generate a phenomenon, but evolved through descent with modification. Modifications often involved connecting up operations in a manner that did not respect sequence – an operation that might be thought of as occurring later in a sequence might be connected to one that occurs earlier (thereby affecting subsequent iterations of the sequence). As such modifications accrued over evolutionary time, modes of organization became much more complex. Understanding the consequences for the behavior of mechanisms requires going beyond accounts of basic mechanistic explanation. Historically, researchers did this on a case-by-case basis. For example, after Krebs recognized that the intermediates in oxidative metabolism formed a cycle (the tricarboxylic or Krebs cycle), he puzzled over the organization and speculated as to its significance [20]. More recently, however, systems biologists have made use of tools from graph theory to represent modes of mechanism organization and their consequences for the functioning of mechanisms.

Representing System Organization in Graphs

▼ Graph theory provides a powerful tool for representing the organization of a mechanism while abstracting from its specific components and then analyzing the consequences for any mechanism instantiating that organization [21]. In a graph, each part of the mechanism is represented as a node, often drawn as a circle, and the operations through which one node affects others is represented by an edge or line between them, sometimes with an arrow to indicate the direction of effect. To understand the behavior of a mechanism that satisfies a graph one must follow out the connections between nodes. With some relatively simple graphs, theorists can do this in their heads, but as pathways multiply this becomes much more difficult.

As researchers began to develop graphs of particular networks, it was possible that each would have its own design and that there would be no common principles that could be elicited and applied more broadly. Each network might be analyzed in a computational model (with the graph providing a productive guide in developing the model; see [22]), but there might not have been any general principles that could be elicited. In fact, however, some powerful organizing principles have been identified and consequences elicited both for the normal behavior of mechanisms that employ them and the pathologies that result when organization is disrupted. These analyses have proceeded both with respect to large-scale graphs of whole mechanisms (or systems of multiple mechanisms) and small subgraphs that are frequent constituents of larger scale networks. I focus first on the organizational principles at the large-scale, then turn to the analyses of subgraphs.

Large-scale organization of graphs

Mathematicians in the mid-20th century performed some of the pioneering studies on the properties of particular types of graphs. Erdős and Rényi [23] explored graphs that began with a set of nodes and randomly added edges between them. They discovered that when the number of connections was much smaller than the number of nodes, only small, disconnected clusters of connected nodes would develop. However, when the number of connections was approximately equal to half the number of nodes a phase transition would occur in which a single giant cluster emerged. Within clusters there is usually a short connection path between any 2 nodes; as a result if the nodes exhibit

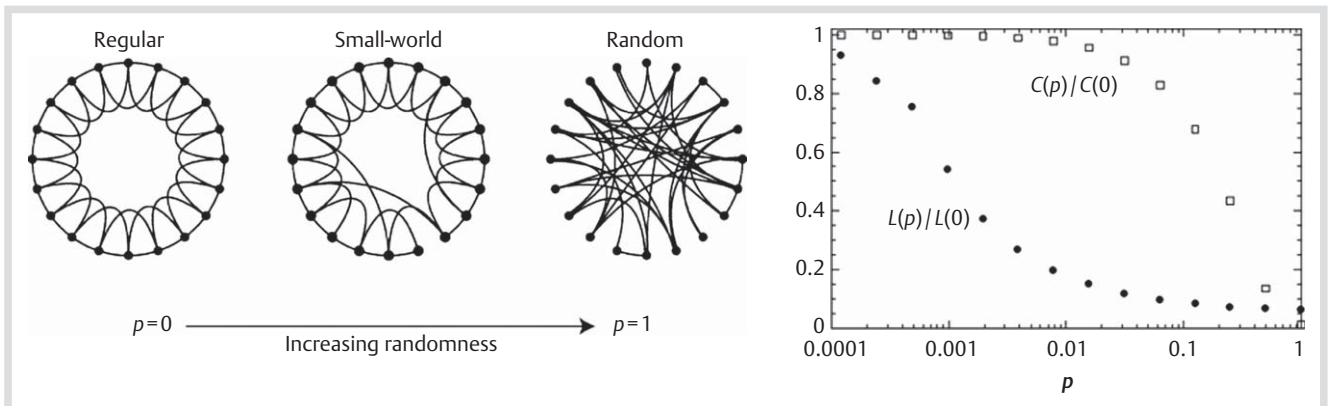


Fig. 1 Watts and Strogatz's (1998) representation of small-networks as arising as an intermediate between regular lattices and random networks as one begins randomly replacing connections in a regular lattice with longer-range ones. The graph shows how path length and clustering change as the probability of rewiring increases. Regular lattices exhibit both high clustering and long mean path length whereas random networks exhibit low clustering and short mean path length. The broad region in between in which clustering remains high while path length drops to near that of a random network is where small worlds reside. Reprinted by permission from Macmillan Publishers Ltd: Nature, 393, Figures 1 and 2, copyright 1998.

oscillatory behavior, nodes across the cluster rapidly synchronize their activity. These analyses continue to be useful in biology. Yook, Oltvai, and Barabási [24] have employed the notion of a giant cluster in which all elements would synchronize in their analysis of protein interactions in yeast. An alternative type of network that was explored involved regular lattices in which each node is connected to neighbors within a neighborhood of a specified size. In such structured networks the path between distant nodes is quite long. Rather than rapidly synchronizing activity across such a network, Ermentrout and Kopell [25] showed that such networks would create waves of activity that propagated across the network, fitting the pattern found in the central pattern generators that regulate motor activity in various animals.

In the previous paragraph I implicitly introduced one of the measures used to distinguish random networks and regular lattices: random networks have a short mean path length (fewest edges that must be traversed to move from one node to another) whereas regular lattices have a much longer mean path length. Another important measure is how nodes form clusters: the clustering coefficient measures the percentage of possible connections among units in a local neighborhood that are actually realized. A lattice scores high on this measure whereas a random network has a much lower score. Identifying short mean path length and high clustering as desirable properties for information processing, Watts and Strogatz [26] identified a type of network, which they termed small-world networks, that exhibit both properties.⁵ They started with a regular lattice and began substituting long-range connections for some of the local connections. They showed that relatively few such substitutions would produce a large drop in the mean path length while the clustering coefficient remained high (► Fig. 1). The result is a design of a network in which locally connected components can constitute specialized modules while remaining closely linked to other components and thus able to rapidly synchronize with them.

Watts and Strogatz identified numerous examples of small-world networks in both biological and social systems. One of

their examples was the neuronal network in the nematode worm *Caenorhabditis elegans*. Through reconstruction from electron micrographs of serial sections, White et al. [29] mapped out all the synaptic contacts between the 302 neurons found in *C. elegans* (the pattern is invariant across individuals).⁶ Among the claims Watts and Strogatz made for small-world networks is that actual mechanisms implementing small-world design would be extremely effective in processing information. Nodes that are highly clustered can be organized into appropriate networks for a given information-processing task but as a result of short path length, each local region can be modulated by activity occurring elsewhere.

Shortly after Watts and Strogatz focused attention on small-world networks, Barabási and his collaborators [32] focused attention on another way in which many real world networks differ from either lattices or random networks. They focused on degree, a measure of the number of edges connected to a given node. In random networks, degree is distributed in a Gaussian manner over a fairly narrow range, providing a scale of network connectivity. What Barabási discovered is that many real world networks degree is distributed according to a power-law in which most nodes are connected to only a small number of other nodes but a few have a very large number of connections. The extremely long tail on a power-law distribution means there is no characteristic scale over which degree is distributed and so such networks are referred to as scale-free. In his discussions of scale-free networks, Barabási emphasized their robustness in the face of random removal of nodes – since the vast majority of nodes have few connections with other nodes, removing them has little effect on how the network behaves. On the other hand, disrupting nodes with unusually high number of edges connected to them often has serious consequences since many paths through the network pass through them. In networks with high clustering into modules, these nodes are often referred to as *hubs*, emphasizing their central place in a module or in linking multiple modules. Although some biological networks such as

⁵The notion of a small-world has a much longer history. For example, it was analyzed mathematically by Pool and Kochen [27] in a paper that was circulated for 2 decades before it was published and was established empirically in a classical social psychology experiment by Milgram [28].

⁶Varshney, et al. [30] updated the original map, identifying or modifying 3000 chemical synapses, gap junction connections, and neuromuscular junctions. Sohn, et al. [31] employed a variety of analytic techniques to identify modules in the *C. elegans* nervous system and identified 5 clusters of neurons that could each be linked with functional circuits experimentally defined and that themselves organized into 2 super-clusters.



Fig. 2 A negative feedback subgraph.




Fig. 3 The triangular sub-circuit White (1985) identified as occurring frequently in the *C. elegans* nervous system.

protein networks are scale-free, in others the spatial limitations on numbers of connections and the cost of adding connections limits the maximum number of edges that can connect to a node; such networks are not fully scale-free. Nonetheless, their degree distribution may be broader than a normal distribution (they may exhibit an exponential distribution or an exponentially truncated power-law distribution). In these networks as well nodes with the greatest number of edges may serve as hubs.

Organization of subgraphs

Subgraphs are organizations involving a small number of nodes that may then be embedded in a larger network by adding additional connections from or to nodes outside the subgraph. These subgraphs can be analyzed for the characteristic behavior they support when embedded in larger networks. Prior to the beginning of the 21st century, a few exemplars of subgraphs attracted attention, but there were no systematic attempts to analyze different patterns of organization as components or building blocks of larger networks.

One of the first subgraph organizations to attract serious theoretical interest was negative feedback whereby an operation later in a sequence feeds back to inhibit an operation occurring earlier in the sequence (see Fig. 2). Although first proposed in his design for a water clock by Ktesibios in the 3rd century BCE as a means of maintaining a constant quantity of water in the supply vessel so as to generate a constant stream into the recording vessel, negative feedback did not become recognized as a general principle of organization for two thousand years. Rather, it was reinvented in different engineering contexts, such as to control furnaces and windmills, as needed until Watt designed the centrifugal governor for the steam engine and Maxwell [33] developed an abstract mathematical analysis of governors. In the 20th century instances of negative feedback were identified in physiological systems and the cyberneticists [34] celebrated it as a general principle of control in biological and engineered systems. Even so, an important feature of negative feedback, its propensity to generate oscillations when delays and non-linearities are introduced into the feedback loop, was largely ignored in biology except by investigators such as Goodwin [35] who were seeking to understand endogenous oscillations in biological systems.

A notable example of identifying a potentially functionally significant subgraph arose when, in mapping the complete neuronal network of *C. elegans*, White [36] noted “the preponderance of triangular sub-circuits” such as shown in Fig. 3 and speculated as to their functional significance. From the electron micrographs, White could not ascertain whether connections were excitatory or inhibitory, but he considered what would happen if the pathways from A to both B and C were excitatory but that from B to C was inhibitory: A signal from A would initially elicit a response from C, but this would soon be suppressed as a result of the negative connection from B to C. He suggests: “The whole system would therefore act as a differentiator, the output from [C] being proportional to the rate of change of stimulus. As the animal is constantly moving, this reformation is probably of more value to it than an absolute measure of the stimulus.”

The analysis of subgraphs advanced from considering individual examples to a wide-scale project in the research of Uri Alon as he was developing graphical analyses of gene transcription and metabolic networks in bacteria and yeast. He began to notice “recurring, significant patterns of interconnections” in subgraphs of 1 to 4 nodes that appeared much more frequently than would be expected by chance. The chance rate was assessed by the frequency of the subgraphs in randomly constructed networks with the same degree of node connectivity. Alon and his collaborators developed an algorithm for searching databases specifying network connectivity for unusually frequently occurring subgraphs, which he termed motifs.⁷

Alon found what he called a feedforward loop (see Fig. 4, left) as occurring in “hundreds of non-homologous gene systems” [38] in the transcription network of *E. coli*. This motif consists of three units in which an operon (X) responds to an input signal (S) by producing a transcription factor that both regulates an operon for an output protein (Z) and an intermediate operon (Y) which also produces a transcription factor affecting the output. (The specific substances identified in Fig. 4 are the ones involved in an instance of the motif in controlling synthesis of L-arabinose.) One operon could regulate another in an excitatory or in an inhibitory manner and Alon termed a feedforward loop *coherent* if both the direct and indirect pathways affected the output in the same manner (the loop in Fig. 4 is coherent). To determine what function such a motif might perform, Alon turned to mathematical analysis. He found it sufficient to use a step function to approximate the effect of one factor on another and an AND- or OR-gate to model the combined effect of X and Y on Z⁸. When it functioned as an AND-gate, as in the case illustrated, the motif functioned as a *persistence detector* in that an output was generated only if the input to X persisted. As shown on the right in Fig. 4, when the input X is transient, Y begins to respond but before it can reach a full response, the input ceases. There is no effect on Z. But when X persists for several seconds, Y builds up and slightly afterwards Z begins to be expressed. In other domains, Alon and his collaborators identified different subgraphs as meeting the criterion for motifs. In food webs, only a sequential arrangement of 3 nodes qualified as a motif,

⁷Shen-Orr et al. [37] introduced the term: “We generalize the notion of motifs, widely used for sequence analysis, to the level of networks. We define ‘network motifs’ as patterns of interconnections that recur in many different parts of a network at frequencies much higher than those found in randomized networks.”

⁸More typically the effect of an excitatory or an inhibitory stimulus is modeled with a Hill function, which produces an S-shaped activation curve; the step function approximates a very steep S-curve. For discussion, see Alon [39].




Fig. 4 One of the feed-forward loop motifs examined by Alon and his collaborators. In this case, Z behaves as an AND-gate, initiating production of araBAD only when it receives inputs from both X and Y. The graph on the right is from a mathematical simulation of the motif and shows that when X experiences a short-lasting increase, it has minimal effect on Y and none on Z. When X experiences a longer increase, sufficient amounts of Y accumulate and shortly thereafter Z begins to increase in concentrations. (Figure on right reprinted by permission from MacMillan Publishers Ltd: Nature Genetics, 31, Figure 2a, Copyright 2002).

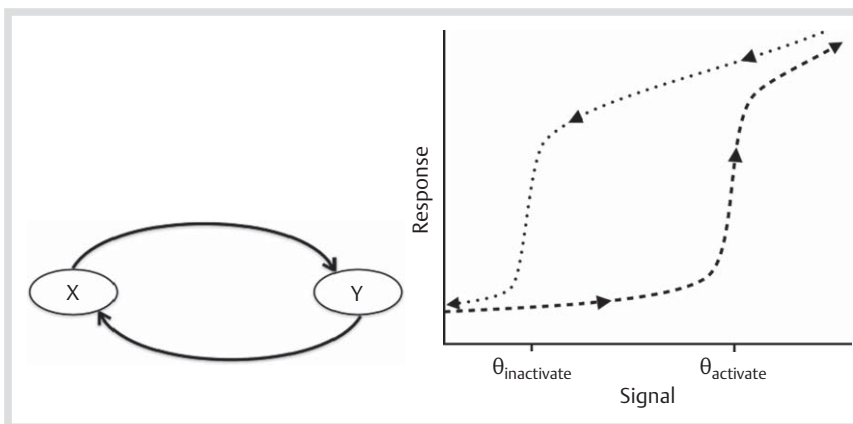


Fig. 5 Positive feedback loop subgraph (left) and graph showing how as signal increases, response sharply increases at θ_{activate} but does not return to lower levels until the signal drops below $\theta_{\text{inactivate}}$.

whereas in electronic circuits the only 3-node subgraph that qualified was a feedback loop. He suggests that the occurrence of subgraph design is related to how the larger network behaves and that the subgraph may have been selected for its contribution. Other theorists have picked up on Alon's language of motifs, but have applied the term to subgraphs without respect to their frequency in a given network. They then speak of motifs occurring more or less frequently than expected by chance. Tyson & Novák [40], for example, analyze (using ordinary differential equations rather than Boolean networks) a wide range of subgraphs and use these analyses to explain their functioning in any context in which they appear in a graph representation of a biological system. 2 simple subgraphs that Tyson and Novak analyzed are the positive and double negative feedback loops (► Fig. 5). Given appropriate non-linear relations and parameter values, they show both subgraphs can act as bistable toggle switches—depending on the value of an input signal (delivered, for example, to X), the values of the nodes switch from low to high but they switch on at a higher value than they switch off. Thus, once the input crosses the threshold that turns the switch on, merely dropping below that threshold will not turn the switch off. Rather, it stays on until the input drops below a significantly lower value (► Fig. 5, right). Examples of both the double negative and the positive feedback loops are found in the biochemical system that ensure progress through the stages of the eukaryotic cell cycle from G_1 (Gap 1 or growth phase) to S (synthesis or DNA replication phase), G_2 (Gap2 or continued growth phase), and M (mitosis). Progress through the transitions from G_1 to S, from G_2 to M, and returning from M to G_1 is regulated to insure passage in

only one direction. At the core of each step is a dimer of a cyclin and a cyclin dependent kinase (CDK:Cyclin or mitosis promoting factor [MPF]; different cyclins and CDKs are involved at each transition). The first involves a double negative feedback loop between the CDK:Cyclin and a complex of CDK:Cyclin with CKI, a cyclin dependent kinase inhibitor, while the second employs a positive feedback loop between CDK:Cyclin and its phosphorylated form (phosphorylation is catalyzed by a member of the Wee1 family and dephosphorylation by a member of the Cdc25 family). Both of these behave as toggle switches and insure that the cell does not return to the previous stage. The final transition from M to G_1 is regulated by a negative feedback oscillator that results in a large rise then fall of CDK:Cyclin activity [41].

Applying Graph Representations to Neural Mechanisms

▼ In this section I turn to the application of graph theory representations of organization to neural mechanisms in organisms with brains (in contrast to the nematode nervous system discussed in the previous section). I proceed as in the previous section by considering first organization of the large-scale networks relating brain regions and then turning to subgraphs found in these networks.

Large-scale organization of brain networks

In the section “A Brief Introduction to Mechanistic Explanation and Neuroscience” above I briefly referred to research on the mammalian visual system; over the second half of the 20th cen-




Fig. 6 Felleman and van Essen's (1991) matrix indicating connections between cortical visual areas is shown on the left. Each row shows whether a connection had been identified between the area shown on the left and the areas indicated at the top of each column. A plus indicates a connection has been found. A period indicates that a connection has been sought but not found. A blank square indicates that the pathway has not been tested for and a question mark indicates conflicting evidence. On the right is their graphical representation of the hierarchical organization among these regions as well as the sub-cortical areas and a few non-visual areas. From Felleman, D. J., & van Essen, D. C., Distributed hierarchical processing in the primate cerebral cortex, *Cerebral Cortex*, 1991, 1, Table 3 and Figure 4, by permission of Oxford University Press. (Color figure available online only.)

tury this research resulted in the identification of numerous brain regions distinguished by criteria such as cytoarchitecture, connectivity, and the topographical maps identified as researchers charted how neurons responded to stimuli located in different parts of the organism's visual field. Drawing upon a large number of studies, Felleman and van Essen [13] differentiated 32 cortical regions involved in visual processing and showed that 305 of the 992 possible connections were realized between them. They presented their results both in a matrix in which local clusters are apparent and in a graphical analysis that reflects the hierarchical pattern they identified by distinguishing feedforward, feedback, and lateral connections (► Fig. 6). Sporns and Zwi [42] calculated path length and clustering for the graph described by Felleman and van Essen as well as ones for the complete macaque and cat cortices characterized by other researchers and found that all 3 graphs exhibited high clustering with short path lengths characteristic of small-worlds. They also examined regions within the network and found that in each network areas differed significantly. For example, area V4 in the macaque exhibits both a low path length and a low clustering coefficient (characteristics of random networks) whereas areas V1, V4t, and STPa have high path lengths and high clustering (characteristic of regular lattices). In the case of V4, it is a highly connected area (21 incoming and outgoing connections) but the areas to which it is connected do not themselves form a common cluster. Other areas such as area A3a in the somatosensory cortex shows the opposite pattern – long path lengths and high clustering. It is connected only to areas A1 and A2. Given the relatively small numbers of nodes and connec-

tions in these databases, the analyses were not able to show a scale-free distribution, although all 3 databases did reflect significantly higher variance than random or lattice networks, suggestive of the occurrence of hubs.

The invasive techniques available for most of the 20th century for mapping neural connections limited researchers to non-human species such as the macaque and cat, but in recent decades a variety of ways of employing magnetic resonance imaging has enabled comparable research on the human brain. One approach has used detected correlations in thickness of grey matter between cortical areas (the cause of these correlations is currently unknown) in multiple subjects as predictive of connections [43]. Another is diffusion MRI that provides evidence of myelinated fiber tracks in cerebral white matter. Sporns, Tononi, and Kötter [44] introduced the term connectome⁹ for the comprehensive graph of brain connections at different levels of organization and a number of researchers are now combining their efforts to develop a detailed account of the human connectome. (Efforts are also being directed at developing the connectome of other species; for research on the fruit fly connectome, see [46].) Although still in its early phase, connectome research is already providing insights into the organization of the human brain. Applying measures such as mean path-length and clustering to graphs constructed with diffusion MRI has generated evidence that the human brain, like the cat and macaque, exhibits a small-world architecture [47,48]. In addition to measures of path

⁹Independently Hagmann [45] also introduced the term in his dissertation that employed diffusion MRI to identify structural networks.




Fig. 7 On the left is the dual-dyad motif found frequently in networks relating brain regions in mammals. The dual-dyad does not occur with elevated frequency in *C. elegans*, but the 2 sub-graphs on the right do.

length and clustering, connectome researchers have focused on identification of modules as brain areas with extensive interconnections and hubs that link them. I will focus more specifically on hubs when I turn to the identification of motifs in brain networks, but for now note that Hagmann et al. [49] identified numerous hubs located along the anterior-posterior medial axis of the brain which included the rostral and caudal anterior cingulate cortex, the paracentral lobule, and the precuneus. These hubs are highly connected to each other and between them connect to regions in virtually all other areas in both hemispheres. This suggests a central network that is important for directing communication through the brain.

A major reason for interest in the structural organization of the connectome is that connections between brain areas are likely to serve functional ends such as information exchange. Accordingly, connectome researchers have explored graphing information secured in various ways about functional connectivity, characterized in terms of statistical dependence between measures of brain activity in different brain regions [50]. Such dependence is often identified between recordings showing oscillatory activity such as EEG, which detects oscillations in electrical potentials in the 1–100 Hz range. It is difficult, however, to localize the source of activity recorded with EEG. Biswal et al.'s [51] discovery, through time-series analysis of ultraslow oscillations (<0.1 Hz) in fMRI recordings, provided a means of studying coherence between oscillations that could be localized to specific brain regions. During the same period Raichle and his collaborators began to analyze fMRI recordings made in the resting state in which subjects lay quietly in the scanner not directed to perform any task [52]. Their initial interest was in brain regions that showed greater activity in the resting state than in task conditions; they identified these regions as constituting the default mode network. Cordes et al. [53] developed functional connectivity MRI (fcMRI) analysis that applied correlational statistics to resting state BOLD time series data to determine patterns of synchronization and identified networks of regions with corrected activity. The approach has been applied in particular to the default node network [54].

As was hoped, evidence soon developed that the functionally characterized networks largely¹⁰ correspond to those identified structurally. For example, Greicius et al. [56] showed that the regions in the default mode network are anatomically connected while van den Heuvel et al. [57] found that 8 of the nine networks they identified in the resting state correspond to ones that can be characterized anatomically as connected by fiber tracts. Moreover, when analyzed graph-theoretically, these networks were found to exhibit modular small-world architecture [58]. Of particular importance, like the structural analyses discussed

above, these functional analyses identified medial areas in the default mode network, such as the precuneus and the posterior cingulate cortex, as extremely well-connected hubs [59]. These results in particular have elicited new interest in the precuneus and the posterior cingulate cortex and their role in various cognitive activities [50].

Organization of subgraphs linking brain regions

As with other biological networks, analysis of brain networks is revealing subgraphs that help explain brain function. Sporns and Kötter [60] counted the frequency of subgraphs of 2–5 nodes across the macaque visual cortex, macaque cortex, and the cat cortex and identified several frequently occurring subgraphs whose z-score was greater than 5 (i.e., their frequency was 5 standard deviations above the mean) across a variety of random and lattice networks. The one 3-node subgraph that met this condition is shown on the left of **Fig. 7**; Sporns [61] named this the dual dyad motif as it consisted of 2 sets of reciprocal connections (dyads) joined at a common node. As a comparison, when Sporns and Kötter examined the neuronal network of *C. elegans* this subgraph was not significantly increased in frequency but instead the two shown on the right in **Fig. 7** were. They took this as an indication of different processing needs in mammals and worms; in particular, Sporns has long emphasized that mammals must both segregate and integrate information processing and his analysis points to how the dual dyad may serve to integrate information processing performed in separate clusters. When they examined nodes that participated in the dual dyad motif, they found increased participation only in areas that constituted hubs – nodes that are characterized by relatively low clustering, short path lengths to the rest of the network, and high centrality (fraction of shortest paths that go through the node). The 2 dyads constituting the dual dyad make sense at such hubs as means of linking nodes from different clusters. Vicente et al. [62] demonstrated that dual dyads would promote zero phase-lag synchrony across long distances, suggesting that they can promote communication among brain areas (when regions are synchronized, action potentials received from one area are more likely to elicit response in the receiving area). Sporns, Honey and Kötter [63] expanded on this analysis. They used a somewhat different criterion for the increased frequency of a subgraph, treating it as significantly increased if it had a z-score greater than 3 when compared both to random and to lattice controls. By this criterion, a number of regions exhibit significantly increased participation in the dual dyad: VP, MSTd, V4, DP, FST, 46, 7a, 7b, Ig, STPp, and TH. 3 other areas, LIP, VIP, and FEF, exhibit increased participation in another motif (which adds a one way connection between the 2 nodes not connected in the dual dyad). Of the brain areas that participate in the dual dyad, V4 is the most frequent; moreover, when it appears in the dual dyad, V4 is typically the apex node. In this position, it serves to link 2 nodes that are constituents of distinguishable clusters or modules.

¹⁰There are numerous functionally related areas that are not directly connected structurally. However, relatively short indirect structural connections can typically be identified [55].

Using a variety of criteria such as node degree, motif participation, and centrality, Sporns, Honey, and Kötter identified V4, FEF, 46, 7a, TF, 5, and 7b as the most likely hubs in the macaque cortex. All but V4 counted as connector hubs due to the diversity and distance of the regions to which they connect whereas most of V4 connections are to other visual areas, including areas in both visual streams. It thus counts as a provincial hub. The connector hubs turned out to be highly connected among themselves, forming what Sporns, Honey, and Kötter call “hub complexes.” They analyzed the effects of removing the 2 types of hubs on the graph theoretic measures of the overall network structure – deleting a connector hub increased the small world character of the network as the increased clustering more than compensated for the increase in path length. In contrast, deleting a provincial hub reduced the small-world character of the network as a result of decreasing clustering.

Altered Graphs, Disrupted Oscillations, and Mental Disorders

Graph theory analyses of the networks relating brain areas are useful both for understanding how these networks support cognitive functions (in [64], I argue that they require developing a new conception of the cognitive architecture subserving cognitive performance) and how disruptions in graph structure characterize mental disorders. In this section I turn briefly to the latter use and show how identifying the altered graphs found in brains of patients with mental disorders are providing new insight into these disorders. Although the linkage between structural and functional graphs remains important for understanding these disorders, it is the functional graphs, based on synchronized oscillatory behavior at the lower frequencies observed in fMRI, that are proving especially insightful. Alterations in the pattern of synchronization between brain regions corresponds to altered exchange of information and this offers promise in explaining the altered cognition exhibited by patients suffering these disorders.

In the previous section I introduced the default mode network, a network of brain regions that was initially identified as being less active in task conditions than in the resting state. Episodic memory tasks were an exception: regions constituting the default mode network remained highly active in these tasks [65]. Drawing upon the literature on undirected thinking or mind-wandering [66], several researchers inferred that the default mode network was involved in ruminations about the events of one’s life and planning future activities that participants would pursue while resting quietly in the scanner. Thus, Buckner, Andrews-Hanna, and Schacter [67] 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” (p. 2). Although other researchers have advanced alternative interpretations of the primary function of activity in the default mode network, the case that it is employed in reflection and planning, and may figure importantly in how subjects conceptualize themselves, is compelling.

In the decade since the characterization of the default mode network, researchers have found altered activity in the default mode network in patients with a wide range of mental disorders including dementia, Alzheimer’s disease, autism, schizophrenia, anxiety and depression, obsessional disorders, attention-deficit/

hyperactivity disorder, and post-traumatic stress disorder [67–69]. This, however, only provides part of the picture. Other networks that can be identified by their coherent patterns of oscillation in the resting state become active in various task conditions, and a crucial part of normal brain function is the process of switching between networks. Altered behavior in the default mode network in a given mental disorder may be an effect or a cause of disrupted engagement between networks [69]. In this paper, however, I focus only on disruptions found in the default mode network.

Research on Alzheimer’s disease has revealed some of the strongest evidence of atypical activity in the default mode network. A clue to such atypical activity was the finding that Alzheimer’s patients exhibit reduced metabolism in brain regions corresponding to the posterior portions of the default mode network – the posterior cingulate cortex/retrosplenial cortex, the inferior parietal lobule, and lateral temporal cortex [70]. These same regions also exhibit atrophy in Alzheimer’s patients. When researchers turned to analyzing default mode network activity in fMRI studies, they found that these regions do not exhibit the reduction in activation in task conditions that is found both in younger subjects or in normally aging adults [71]. The intrinsic activity in these regions is also not correlated [72]. Drawing on the fact that the plaques found in autopsy of Alzheimer’s patients form first in regions of the default mode network has led Buckner and his collaborators to advance the “metabolism hypothesis” that the activity of the network throughout rest results in increased metabolism that generated increase in amyloid β protein that initiate the formation of plaques and tangles [59].

Corresponding to these functional findings, researchers applying graph-theoretical analyses have identified altered network structure in Alzheimer’s patients. Although path length is normal, they exhibit lower clustering than is found in normal participants [73]. The clustering coefficient is particularly reduced in the hippocampus, a part of the default mode network, but connectivity is increased in the frontal lobe outside the default network. The hub structure is especially altered, with the greatest loss of hubs within the default mode network – especially the posterior cingulate cortex and temporal lobe hubs – with minimal effect on frontal lobe hubs [74]. These researchers showed that using these network measures one can distinguish Alzheimer’s patients from others with mild impairments.

A quite different pattern of alterations in network structure is manifest in schizophrenic patients. Instead of a decrease, Garrity et al. [75] found an increase in default mode network activity, especially in medial prefrontal cortex and the posterior cingulate cortex/retrosplenial cortex during hallucinations, delusions, and thought confusions. In terms of network structure, whereas Alzheimer’s patients inhibited normal path lengths, schizophrenics exhibit both reduced clustering and increased path lengths [76]. Moreover, they exhibit disrupted hubs in frontal as well as parietal and temporal lobes [77]. Schizophrenics also manifest changes in a second network – salience network, that is usually anticorrelated with the default mode network. Thus, White et al. [78], demonstrated greater activity in 2 areas of the salience network – the anterior insula and the frontal operculum – which they argue is consistent with abnormally active monitoring for auditory inputs that might explain hallucinations.

While the study of differences in network organization, structural and function, in patients with mental disorders, is in its infancy, research on default mode network activity in patients

with other disorders is also yielding suggestive finding. For example, Kennedy, Redcay, and Courchesne [79], see also [80] found that autism patients failed to show normal deactivation of the default mode network in task conditions. Given the linkage between default mode activity and self-directed rumination and planning, the finding suggests that continued engagement of the default mode network may figure in the social deficits of individuals with autism. Particularly striking is the correlation they observed between medial prefrontal cortex (a region in the default network) activity and the degree of social impairment as measured by the Autism Diagnostic Interview-Revised. With depressed patients Greicius et al. [81] identified a different pattern of alteration in the default mode network, with enhanced prefrontal processing and increased recruitment of the subgenual cingulate into the network. Abnormal activity in the subgenual cingulate had been identified in several studies of major depression and this demonstration of its abnormal recruitment into the default mode provides perspective on how the pathology operates. These studies point to the promise that analyzing the altered dynamics in neural processing in networks such as the default mode network can provide valuable new insights into mental disorders.

Conclusions

Traditional approaches to understanding mechanisms emphasized strategies for decomposing mechanisms into their parts and operations. Although initially this led to identification of only a few parts and operation in given mechanisms, continued research, especially performed with more powerful search techniques such as genetic screens and neural imaging, identified many times more parts, although understanding the operations in which they figured often lagged. When only a few parts and operations were identified, researchers were often able to re-compose the mechanism in their heads (often supported by diagrams), tracing out the effects of individual components and accumulating them. But as the number of components and the pathways by which they were connected increased, researchers required new tools for understanding organization. Although still at an early stage of development (in part, because the graphs of brain networks are still in early stages of development), graph theoretic analyses are already bearing fruit in characterizing the large-scale organization of the brain and the local connectivity between regions, providing the basis for dynamic mechanistic explanations of mental activity. The investigation of how these are altered in patients with mental disorders is beginning to provide insights into these disorders.

At the large-scale, the human brain, as well as that of the macaque and the cat, has a small-world organization with short path lengths enabling rapid coordination across brain regions and high clustering, allowing for specialized processing modules. This alone is not terribly surprising since most natural occurring networks exhibit these characteristics. But researchers are also identifying more micro-organization such as a hierarchy of modules linked by hubs. Potentially of great significance is the network of hubs along the midline that may play a crucial role in coordinating processing across the brain. Not surprisingly, disruption to this hub-structure is a major feature in a variety of mental disorders.

Even if the prominence of small-world organization is not unique to the brain, focusing on it leads to a different perspec-

tive on brain organization than has been prevalent in neuroscience. The idea that the brain is divided into modules that perform different information processing tasks has played a central role in attempts to understand brain function, but in the context of small-world networks with hubs this notion is importantly recast. The short path length in such networks ensures ongoing interaction between differentiated modules. In particular, different areas are able to synchronize their activity, facilitating communication of information between them, and when this fails, mental disorders ensue.

At the small-scale researchers have both begun to identify sub-graphs and to analyze the contribution they make to larger mechanisms. Identifying those sub-graphs that appeared far more likely than expected by chance as motifs, Alon and his collaborators addressed the functional contribution of these particular motifs. Sporns extended the approach in the case of brain networks, both identifying the dual dyad as occurring frequently and then focusing on the particular parts of the network (hubs) in which it occurred. Others such as Tyson have examined sub-graphs more generally and developed computational models of their behavior. Such an approach could further advance the understanding of how local organization contributes to the behavior of component mechanisms in the brain and provide better understanding of how disruptions to such organization result in mental disorder.

The application of graph theoretic analyses brings a systems perspective to neuroscience, which has long been dominated by the reductionistic emphasis on decomposing parts and operations and localizing operations in specific parts of the brain. But it does not supplant the need to identify parts and identify their operations but complements it by providing new tools for re-composing the mechanisms as it identifies modes of organization and their effects on the operation of mechanisms. The quest to understand brain mechanisms and the pathologies that afflict them requires both reductionistic and holistic perspectives to generate dynamic mechanistic explanations.

Conflict of Interest

There are no conflicts of interest for the author.

References

- 1 Bechtel W, Richardson RC. Discovering complexity: Decomposition and localization as strategies in scientific research. Cambridge, MA: MIT Press; 2010
- 2 Bechtel W, Abrahamsen A. Explanation: A mechanist alternative. *Stud Hist Phil Biol Biomed Sci* 2005; 36: 421–441
- 3 Bechtel W. Mechanism and biological explanation. *Philosophy of Science* 2011
- 4 Craver CF. Explaining the brain: Mechanisms and the mosaic unity of neuroscience. New York: Oxford University Press; 2007
- 5 Hempel CG. Aspects of scientific explanation, in *Aspects of scientific explanation and other essays in the philosophy of science*. Hempel CG, ed. New York: Macmillan; 1965; 331–496
- 6 Glennan S. Mechanisms and the nature of causation. *Erkenntnis* 1996; 44: 50–71
- 7 Machamer P, Darden L, Craver CF. Thinking about mechanisms. *Phil Sci* 2000; 67: 1–25
- 8 Buchner E. Alkoholische Gärung ohne Hefezellen (Vorläufige Mitteilung). *Ber dtsh chem Ges* 1897; 30: 117–124
- 9 Embden G, Deuticke HJ, Kraft G. Über die intermediären Vorgänge bei der Glykolyse in der Muskulatur. *Klin Wschr* 1933; 12: 213–215
- 10 Bechtel W. Discovering cell mechanisms: The creation of modern cell biology. Cambridge: Cambridge University Press; 2006

- 11 Hubel DH, Wiesel TN. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 1962; 160: 106–154
- 12 Hubel DH, Wiesel TN. Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 1968; 195: 215–243
- 13 Felleman DJ, van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1991; 1: 1–47
- 14 Bechtel W. *Mental mechanisms*. London: Routledge; 2008
- 15 O'Keefe JA, Dostrovsky J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Res* 1971; 34: 171–175
- 16 O'Keefe JA, Nadel L. *The hippocampus as a cognitive map*. Oxford: Oxford University Press; 1978
- 17 Bliss TVP, Lomo T. Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. *J Physiol* 1973; 232: 331–356
- 18 Craver CF. The making of a memory mechanism. *J History Biol* 2003; 36: 153–195
- 19 Craver CF. Interlevel experiments and multilevel mechanisms in the neuroscience of memory. *Phil Sci* 2002; 69: S83–S97
- 20 Krebs HA. Cyclic processes in living matter. *Enzymologia* 1946; 8 12: 88–100
- 21 Levy A, Bechtel W. Abstraction and the organization of mechanisms. *Phil Sci* (in press)
- 22 Jones N, Wolkenhauer O. Diagrams as locality aids for explanation and model construction in cell biology. *Biol Philosophy* 2012; 27: 705–721
- 23 Erdős P, Rényi A. On the evolution of random graphs. *Proc Math Inst Hungarian Acad Sci* 1960; 5: 17–61
- 24 Yook S-H, Oltvai ZN, Barabási A-L. Functional and topological characterization of protein interaction networks. *PROTEOMICS* 2004; 4: 928–942
- 25 Ermentrout GB, Kopell N. Frequency plateaus in a chain of weakly coupled oscillators. 1. *Siam J Math Anal* 1984; 15: 215–237
- 26 Watts D, Strogatz S. Collective dynamics of small worlds. *Nature* 1998; 393: 440–442
- 27 Pool IdS, Kochen M. Contacts and influence. *Social Networks* 1978; 1: 5–51
- 28 Milgram S. The small world problem. *Psychol Today* 1967; 2: 60–67
- 29 White JG, Southgate E, Thomson JN *et al*. The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Phil Trans Sec B* 1986; 314: 1–340
- 30 Varshney LR, Chen BL, Paniagua E *et al*. Structural properties of the *Caenorhabditis elegans* neuronal network. *PLoS Comput Biol* 2011; 7: e1001066
- 31 Sohn Y, Choi MK, Ahn YY *et al*. Topological cluster analysis reveals the systemic organization of the *Caenorhabditis elegans* connectome. *PLoS Comput Biol* 2011; 7: e1001139
- 32 Barabási A-L, Albert R. Emergence of scaling in random networks. *Science* 1999; 286: 509–512
- 33 Maxwell JC. On governors. *Proc Roy Soc Lond* 1868; 16: 270–283
- 34 Wiener N. *Cybernetics: Or, control and communication in the animal and the machine*. New York: Wiley; 1948
- 35 Goodwin BC. *Temporal organization in cells; a dynamic theory of cellular control processes*. London: Academic; 1963
- 36 White JG. Neuronal connectivity in *Caenorhabditis elegans*. *Trends Neurosci* 1985; 8: 277–283
- 37 Shen-Orr SS, Milo R, Mangan S *et al*. Network motifs in the transcriptional regulation network of *Escherichia coli*. *Nature Genets* 2002; 31: 64–68
- 38 Mangan S, Zaslaver A, Alon U. The coherent feedforward loop serves as a sign-sensitive delay element in transcription networks. *J Mol Biol* 2003; 334: 197–204
- 39 Alon U. *An introduction to systems biology: Design principles of biological circuits*. Boca Raton, FL: Chapman & Hall/CRC; 2007
- 40 Tyson JJ, Novák B. Functional Motifs in Biochemical Reaction Networks. *Ann Rev Phys Chem* 2010; 61: 219–240
- 41 Tyson JJ, Chen KC, Novák B. Sniffers, buzzers, toggles and blinkers: dynamics of regulatory and signaling pathways in the cell. *Curr Opin Cell Biol* 2003; 15: 221–231
- 42 Sporns O, Zwi JD. The small world of the cerebral cortex. *Neuroinformatics* 2004; 2: 145–162
- 43 Chen ZJ, He Y, Rosa-Neto P *et al*. Revealing modular architecture of human brain structural networks by using cortical thickness from MRI. *Cereb Cortex* 2008; 18: 2374–2381
- 44 Sporns O, Tononi G, Kötter R. The human connectome: A structural description of the human brain. *PLoS Comput Biol* 2005; 1: e42
- 45 Hagmann P. From diffusion MRI to brain connectomics. 2005; EPFL
- 46 Chiang A-S, Lin CY, Chuang CC *et al*. Three-dimensional reconstruction of brain-wide wiring networks in drosophila at single-cell resolution. *Curr Biol* 2011; 21: 1–11
- 47 Gong G, He Y, Concha L *et al*. Mapping anatomical connectivity patterns of human cerebral cortex using in vivo diffusion tensor imaging tractography. *Cereb Cortex* 2009; 19: 524–536
- 48 Hagmann P, Kurrant M, Gigandet X *et al*. Mapping human whole-brain structural networks with diffusion MRI. *PLoS ONE* 2007; 2: e597
- 49 Hagmann P, Cammoun L, Gigandet X *et al*. Mapping the structural core of human cerebral cortex. *PLoS Biol* 2008; 6: e159
- 50 Sporns O. The human connectome: a complex network. *Ann N Y Acad Sci* 2011; 1224: 109–125
- 51 Biswal B, Yetkin FZ, Haughton VM *et al*. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* 1995; 34: 537–541
- 52 Shulman GL, Corbetta M, Buckner RL *et al*. Common blood flow changes across visual tasks: I. increases in subcortical structures and cerebellum but not in nonvisual cortex. *J Cogn Neurosci* 1997; 9: 624–647
- 53 Cordes D, Haughton VM, Arfanakis K *et al*. Mapping functionally related regions of brain with functional connectivity MR imaging. *Am J Neuroradiol* 2000; 21: 1636–1644
- 54 Greicius MD, Krasnow B, Reiss AL *et al*. Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proc Natl Acad Sci USA* 2003; 100: 253–258
- 55 Honey CJ, Sporns O, Cammoun L *et al*. Predicting human resting-state functional connectivity from structural connectivity. *Proc Natl Acad Sci USA* 2009; 106: 2035–2040
- 56 Greicius MD, Supekar K, Menon V *et al*. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cereb Cortex* 2009; 19: 72–78
- 57 van den Heuvel MP, Mandi RC, Kahn RS *et al*. Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Hum Brain Map* 2009; 30: 3127–3141
- 58 He Y, Wang J, Wang L *et al*. uncovering intrinsic modular organization of spontaneous brain activity in humans. *PLoS ONE* 2009; 4: e5226
- 59 Buckner RL, Sepulcre J, Talukdar T *et al*. Cortical hubs revealed by intrinsic functional connectivity: Mapping, assessment of stability, and relation to Alzheimer's disease. *J Neurosci* 2009; 29: 1860–1873
- 60 Sporns O, Kötter R. Motifs in brain networks. *PLoS Biol* 2004; 2: e369
- 61 Sporns O. *Networks of the brain*. Cambridge, MA: MIT Press; 2010
- 62 Vicente R, Gollo LL, Mirasso CR *et al*. Dynamical relaying can yield zero time lag neuronal synchrony despite long conduction delays. *Proc Natl Acad Sci USA* 2008; 105: 17157–17162
- 63 Sporns O, Honey CJ, Kötter R. Identification and classification of hubs in brain networks. *PLoS ONE* 2007; 2: e1049
- 64 Bechtel W. The endogenously active brain: The need for an alternative cognitive architecture. *Philosophia Scientia* (in press)
- 65 Andreasen NC, O'Leary DS, Cizadlo T *et al*. Remembering the past: two facets of episodic memory explored with positron emission tomography. *Am J Psychiatry* 1995; 152: 1576–1585
- 66 Antrobus JS, Singer JL, Goldstein S *et al*. Mindwandering and cognitive structure. *Trans N Y Acad Sci* 1970; 32: 242–252
- 67 Buckner RL, Andrews-Hanna JR, Schacter DL. The brain's default network: Anatomy, function, and relevance to disease. *Ann N Y Acad Sci* 2008; 1124: (The Year in Cognitive Neuroscience 2008) 1–38
- 68 Broyd SJ, Demanuele C, Debener S *et al*. Default-mode brain dysfunction in mental disorders: a systematic review. *Neurosci Biobehav Rev* 2009; 33: 279–296
- 69 Menon V. Large-scale brain networks and psychopathology: a unifying triple network model. *Trends Cogn Sci* 2011; 15: 483–506
- 70 Buckner RL, Synder AZ, Shannon BJ *et al*. Molecular, structural, and functional characterization of Alzheimer's disease: evidence for a relationship between default activity, amyloid, and memory. *J Neurosci* 2005; 25: 7709–7717
- 71 Lustig C, Synder AZ, Bhakta M *et al*. Functional deactivations: change with age and dementia of the Alzheimer type. *Proc Natl Acad Sci USA* 2003; 100: 14504–14509
- 72 Greicius MD, Srivastava G, Reiss AL *et al*. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. *Proc Natl Acad Sci USA* 2004; 101: 4637–4642
- 73 Supekar K, Menon V, Rubin D *et al*. Network analysis of intrinsic functional brain connectivity in Alzheimer's disease. *PLoS Comput Biol* 2008; 4: e1000100
- 74 Yao Z, Zhang Y, Lin L *et al*. Abnormal cortical networks in mild cognitive impairment and Alzheimer's disease. *PLoS Comput Biol* 2010; 6: e1001006

- 75 Garrity AG, Pearlson GD, McKiernan K *et al.* Aberrant "default mode" functional connectivity in schizophrenia. *Am J Psychiatry* 2007; 164: 450–457
- 76 Bassett DS, Bullmore E, Verchinski BA *et al.* Hierarchical organization of human cortical networks in health and schizophrenia. *J Neurosci* 2008; 28: 9239–9248
- 77 van den Heuvel MP, Mandi RC, Stam CJ *et al.* Aberrant frontal and temporal complex network structure in schizophrenia: a graph theoretical analysis. *J Neurosci* 2010; 30: 15915–15926
- 78 White TP, Joseph V, Francis ST *et al.* Aberrant salience network (bilateral insula and anterior cingulate cortex) connectivity during information processing in schizophrenia. *Schizophr Res* 2010; 123: 105–115
- 79 Kennedy DP, Redcay E, Courchesne E. Failing to deactivate: resting functional abnormalities in autism. *Proc Natl Acad Sci USA* 2006; 103: 8275–8280
- 80 Iacoboni M. Failure to deactivate in autism: the co-constitution of self and other. *Trends Cogn Sci* 2006; 10: 431–433
- 81 Greicius MD, Flores BH, Menon V *et al.* Resting-state functional connectivity in major depression: abnormally increased contributions from subgenual cingulate cortex and thalamus. *Biol Psychiatry* 2007; 62: 429–437