

**Looking Down, Around, and Up:
Mechanistic Explanation in Psychology**

William Bechtel

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

Abstract

Accounts of mechanistic explanation have emphasized the importance of *looking down*—decomposing a mechanism into its parts and operations. Using research on visual processing as an exemplar, I illustrate how productive such research has been. But once multiple components of a mechanism have been identified, researchers also need to figure out how it is organized—they must *look around* and determine how to recompose the mechanism. Although researchers often begin by trying to recompose the mechanism in terms of sequential operations, they frequently find that the components of a mechanism interact in complex ways involving positive and negative feedback and that the organization often exhibits highly interactive local networks linked by a few long-range connections (small-worlds organization) and power law distributions of connections. The mechanisms are themselves active systems that are perturbed by inputs but not set in motion by them. Researchers also need to *look up*—situate a mechanism in its context, which may be a larger mechanism that modulates its behavior. When looking down is combined with looking around and up, mechanistic research results in an integrated, multi-level perspective.

The question of the appropriate level for explaining mental phenomena has been an enduring one in philosophy of psychology. On the one hand, many philosophers have argued for the autonomy of the special sciences (Fodor, 1974). Psychological explanations, they have claimed, need (and can) go no deeper than mental states characterized in intentional idioms such as *believes* and *desires*. On the other hand, reductionists of various stripes have argued for the primacy or sole legitimacy of neuroscience explanations. At the extreme, John Bickle (2003), advocating what he calls *ruthless reduction*, contends that all explanation of mental phenomena must be in cellular and molecular terms. Recently a number of philosophers focusing on biological sciences such as biochemistry, molecular and cell biology, and neuroscience, have argued that explanation in those disciplines commonly takes the form of describing the mechanism responsible for a given phenomenon (Bechtel & Richardson, 1993; Bechtel & Abrahamsen, 2005; Craver, 2007; Darden, 2006; Glennan, 1996; Machamer, Darden, & Craver, 2000; Thagard, 2003). I have recently argued that a similar approach to explanation, albeit one that often focuses on information processing mechanisms, is characteristic of psychology and cognitive neuroscience (Bechtel, 2008). On first appearances, mechanistic explanation seems to embrace the reductionistic perspective. However, as I will argue in this paper, mechanistic explanation requires consideration not just of the parts and operations in the mechanism but of the organization within the mechanism and the environment in which the mechanism is situated.

Accordingly, mechanistic explanation in psychology requires not just looking down (decomposing the mechanism), but also looking around (recomposing the mechanism) and looking up (situating the mechanism).

Despite its prominence in biology and related disciplines, mechanistic explanation was largely ignored by philosophers of science in the 20th century, who focused instead on the centrality of laws in explanation (Hempel, 1965). If mechanistic explanation was discussed at all, it was viewed as a species of nomological explanation (Nagel, 1961). But biologists as well as psychologists seldom avert to laws in giving explanations (Bechtel, 2006; R. Cummins, 2000) and when they do, the laws tend to be those of physics or chemistry (Weber, 2005). Instead, they appeal to a mechanism to explain a phenomenon of interest, where a mechanism is understood as an organized set of parts that perform different operations which are orchestrated so as to realize in the appropriate context the phenomenon in question. The explanation itself consists of representing the mechanism and showing how it realizes the phenomenon (often by simulating its functioning).

For purposes of this paper I will focus on visual perception as the phenomenon of interest (for development of similar points, developed using circadian rhythm research as an exemplar, see Bechtel, in press-a; Bechtel & Abrahamsen, in press-b). I begin by offering a brief account of how neuroscience research over the past century has resulted in a relatively detailed schema of the mechanism of mammalian visual processing.¹ I then turn to the attempts to understand the organization of the various neural components involved in vision. Although investigators often present the visual system as largely a feedforward system that begins with sensory information and ends up with extraction of pertinent information about the visual world,² it is widely recognized that there are as many backwards as forwards projections. These feed all the way back to the earliest stages of visual processing, rendering the visual system (and the brain more generally) as a highly interactive dynamical system. Finally, I will emphasize the critical need to understand the context in which visual mechanisms operate as the behavior of visual mechanisms is often highly conditioned on the regularities and complexities of the environment in which perception occurs.

1. Looking Down: The Components of Mechanisms

A critical aspect of mechanistic explanation is identifying the parts of a mechanism and determining what operation each performs. Although this is not sufficient for providing a mechanistic explanation, it is the aspect of mechanistic explanation which often witnesses the greatest success. In part the reason for this is that there are strategies for empirically operating on mechanisms to secure clues about their parts and what they do. The parts themselves can often be identified by their physical characteristics. But not everything that can be differentiated

¹ Machamer, Darden, and Craver distinguish mechanism sketches (accounts with known gaps between operations) and schemas (accounts exhibiting productive continuity but with many details remaining to be specified). The current accounts of visual processing as shown in Figure 3b below suggest productive continuity, but it is recognized that many other brain areas figure in the pathways of visual processing even though the operations they perform are not known.

² Theorists such as Marr (1982) construed the task of vision as constructing a representation of the three-dimensional world, the evidence is compelling that organisms don't reconstruct the visual scene, but actively sample it to extract that information that is pertinent to action (Churchland, Ramachandran, & Sejnowski, 1994).

physically is a part in the sense relevant to mechanistic explanation—that is, a part that performs an operation (a working part). Since the late 17th century a great deal of effort has gone into delineating the parts of the brain. Effort first focused on the sulci and gyri of the convoluted cerebral cortex of higher mammals, but these turned out not to be working parts but merely the products of the way cortex folds. The regions Brodmann (1909/1994) and other neuroanatomists demarcated at the beginning of the 19th century on the basis of cytoarchitectural criteria provided a more useful starting point, which has been elaborated on by other brain mappers in the 20th century (Mundale, 1998). As the objective of brain mappers is to identify working parts, they typically appeal to criteria such as connectivity and the presence of topographical maps in delineating brain areas, and appeal not infrequently to the operations performed by brain regions (Felleman & van Essen, 1991).

The more challenging endeavor is that of determining the operations performed by parts (which may sometimes be identified without knowing what physical parts are involved). Three strategies have played an especially important role in the developing an understanding of operations performed in the brain:

1. lesioning brain areas (or studying naturally occurring brain lesions) and identifying changes in the way the whole mechanism functions
2. stimulating brain areas and determining the behaviors they elicit
3. recording from brain areas as stimuli are presented to the system.

None of these techniques on its own can provide a complete understanding of the mechanism (in fact, each on its own can easily produce misleading ideas), but each can provide important insights and evidence about the parts and their operations (Bechtel, in press-b).

A common first strategy in attempting to understand how a mechanism performs a behavior is to seek out one part of the mechanism that is responsible for the phenomenon of interest. On the basis of his experience with patients who developed blindness after stroke induced damage to the occipital lobe, Bartolomeo Panizza (1856) proposed that it was the locus of visual processing. Although his findings attracted little attention, shortly afterwards other researchers, most notably Hermann Munk (1881), on the basis of ablation experiments with animals, and Theodor Meynert (1870), on the basis of tracing neural projections from the retina first to the lateral geniculate nucleus of the thalamus and then to the occipital lobe, made a strong case for localizing visual processing in the occipital lobe. Although these claims were contested by David Ferrier (1876) who, on the basis of his own ablation experiments as well as stimulation experiments, argued for a locus in the angular gyrus, they received overwhelming empirical support by the end of the century. Further evidence in support of the claim that the occipital cortex was the locus of visual processing was provided by the discovery of topographical maps of the visual field in the occipital lobe. Salomen Henschen (1893) advanced the first map (one reversed to that which was subsequently accepted) on the basis of tracing projections from the retina to the occipital lobe, but the project was significantly advanced by the analysis of visual deficits following bullet wounds by Tatsuji Inouye (1909) during the Russo-Japanese war and by Gordon Holmes (1919) during World War I. When techniques were finally developed that allowed recording from electrodes inserted next to neurons in living organisms, Talbot and Marshall (1941) confirmed the map by demonstrating electrical response in cells in the appropriate region of occipital cortex in anesthetized cats and monkey when a stimulus was presented in a given region of the visual field (Figure 1).

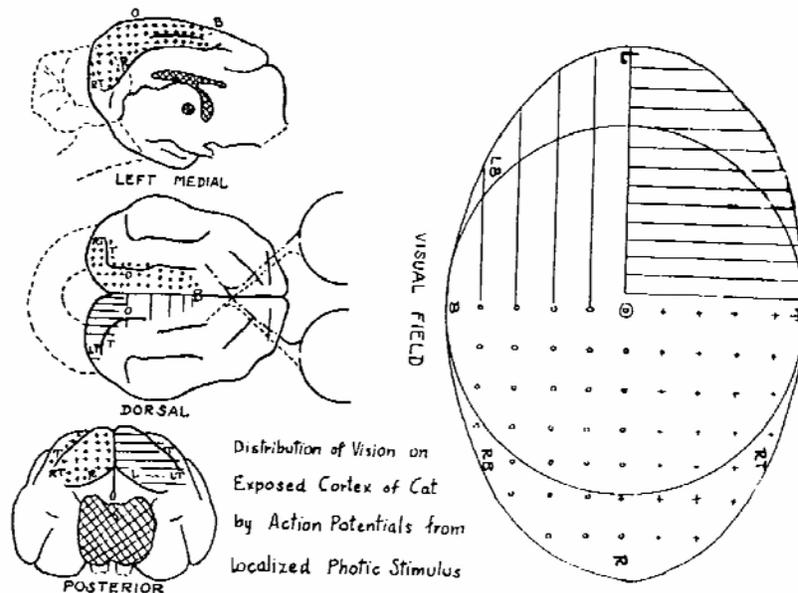


Figure 1. Talbot and Marshall's (1941) projection of areas of the visual field on to primary visual cortex in the cat based on recording from cells.

Localizing a whole process in a part of a system (what Richardson and I called *simple* localization) does not itself constitute an advance in explaining how the part contributes to the phenomenon. Advance requires decomposing the phenomenon into component operations associated with different parts (a *complex* localization). Starting from a simple localization, there are two common ways to develop a complex localization. One is to discover multiple sub-parts within the initially identified part and link them with different operations. The other is to discover that the initial part only performs one of the operations required to produce the phenomenon and then search for other parts that also contribute. In the case of vision, the latter strategy played the major role. Using single-cell recording to determine the nature of stimuli that would produce responses in cells in Brodmann's area 17 in the occipital cortex in anesthetized cats and monkeys, Hubel and Wiesel (1962, 1968) determined that specific cells therein tended to be responsive either to stationary, oriented bars (rectangles) of light or bars moving in a particular direction.³ Some cells responded when the brightness was different on different sides of a line; Hubel and Wiesel described such stimuli as *edges* and over time the term *edge detector* came to be applied as a general label for these cells. Hubel and Wiesel labeled cells that responded to edges at specific locations as *simple cells* and those that responded to edges at different locations in their receptive fields or moving in specific directions across the receptive field as *complex cells*. They also determined that simple and complex cells tended to appear in different cortical layers. By gradually inserting electrodes obliquely to the cortical surface and recording from cells at successive depths, they further discovered a complex organization

³ Previous researchers had succeeded in identifying cells that responded most strongly to contrast between in brightness between the center of their receptive fields and the surround. These included cells in the optic nerve of frogs (Hartline, 1938) and retinal ganglion cells in frogs (Barlow, 1953) and cats (Kuffler, 1953).

amongst cells in which successive neurons responded to edges of gradually changing orientations and often showed preference for stimuli present to one eye or the other.

Hubel and Wiesel's probes of area 17 with electrodes while varying sensory stimuli had revealed a quite unexpected feature of the visual processing mechanism: whereas area 17 had been regarded as the locus of visual processing, it appeared only to detect bars or edges that contrasted with the visual background when they were either stationary or moving. This implied that other brain areas must also be involved in visual processing since seeing involves more than registering edges. Accordingly, Hubel and Wiesel end their 1968 paper by commenting:

Specialized as the cells of 17 are, compared with rods and cones, they must, nevertheless, still represent a very elementary stage in the handling of complex forms, occupied as they are with a relatively simple region-by-region analysis of retinal contours. How this information is used at later stages in the visual path is far from clear, and represents one of the most tantalizing problems for the future" (p. 242).

Hubel and Wiesel thereby effectively set aside Lashley's (1950) proclamation that "visual habits are dependent upon striate cortex [area 17] and upon no other part of the cerebral cortex."⁴

The first suggestion of other specifically visual areas involved the discovery of additional topographical maps that preserved spatial relations in the visual field. Alan Cowey (1964) discovered the first additional area in Brodmann's area 18, which is immediately adjacent to area 17. This became known as V2 (visual area 2), while area 17 was designated V1. Subsequently, it was determined that V2 cells responded to illusory contours (such as created in the Kanizsa illusion—Figure 2) to which V1 cells did not respond. Using single-cell recording, Hubel and Wiesel (1965) confirmed the topographical organization of this area and identified yet a third area, V3, in Brodmann's area 19. By tracing neuron degeneration from these areas, Semir Zeki (1969, 1971) identified two additional areas, known as V4 and V5 (or MT). Although the topological maps became less distinct in these areas, Zeki claimed that cells in V4 responded specifically to the wave length of visual stimuli while those in V5 responded to their motion. It was later determined that in fact these cells detected not just wave-length and motion but perceived color and perceived motion. In a particularly elegant study, Anthony Movshon, William Newsome, and their colleagues demonstrated that by recording from MT cells they could predict a monkey's response to stimuli whose direction of motion was ambiguous and that by microstimulating these cells they could bias that response (Britten, Shadlen, Newsome, & Movshon, 1992). This provided very compelling evidence that MT cells were the locus at which motion was detected.

⁴ Lashley as well as several other brain investigators in the first half of the 20th century embraced a holistic view where what mattered for most of the brain was not specific neural tissue but how much of it there was. Thus, while he gave the name *prestriate region* to the area in front of striate cortex, he construed it as a general association area not limited to specific types of information processing.

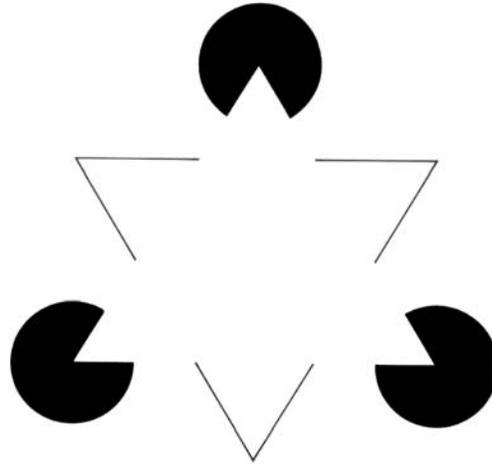


Figure 2. Illusory triangle developed by Gaetano Kanizsa. The triangle generated by the cut outs in the pac man figures is seen as bright white against a less white background. V2 cells respond to the illusory boundary of the triangle to which V1 cells do not.

Even after the operations performed in extrastriate areas had been taken into account, researchers recognized that they still could not explain vision. Zeki commented in his 1974 paper: “The picture that is beginning to emerge, therefore, is one of a mosaic of areas, each with a different functional emphasis. Presumably the visual information analysed in detail in these areas is then assembled at an even more central cortical area” (p. 569). Already in the late 19th century Edward Schäfer (1888) had determined that monkeys with lesions in the temporal lobe could see (they responded to visual stimuli) but could not recognize what they saw. In addition, Heinrich Lissauer (1890) identified in human patients a deficit in recognizing what one saw, a deficit that, following Klüver and Bucy (1938), is known as *visual agnosia*. Morten Mishkin (1966) developed an elaborate preparation in which a monkey’s striate and inferotemporal cortexes both functioned but in which connections between them were cut, and traced the deficit of visual agnosia to inferotemporal lobe damage. After several unsuccessful attempts, Charles Gross and his colleagues succeeded in identifying cells in inferotemporal cortex that responded strongly to specific shapes such as hands, faces, or trees (Gross, Rocha-Miranda, & Bender, 1972). Following a lag of more than a decade in which no further results were published, an explosion of papers appeared reporting locations in inferotemporal cortex responsive to specific types of object stimuli (see Tanaka, 1996, for a review).

The late 19th century also witnessed reports of deficits affecting the ability to locate objects correctly in space following damage to the angular gyrus in the posterior parietal cortex (Ferrier & Yeo, 1884; Brown & Schäfer, 1888), and early in the 20th century Rezső Bálint (1909) described a stroke patient who could not use visual information to guide the motion of his right hand. Initially researchers were unable to elicit responses from parietal lobe neurons during visual tasks, but once techniques were developed for recording from awake, behaving animals, Juhani Hyvärinen and Antti Poranen (1974) found cells that responded when a monkey was required to visually orient towards a stimulus or reach for it. Subsequent research, especially by Richard Andersen and his colleagues (Andersen, Essick, & Siegel, 1985) identified neurons in the lateral intraparietal area that support transformations between locations specified in retinal

coordinates and head-based coordinates, a computation required in order to guide eye movements to objects.

By the 1990s investigators had identified 32 different brain areas that responded specifically to visual stimuli, indicating that they figured in the visual processing mechanism (Felleman & van Essen, 1991). Most of these are areas in occipital, temporal, and parietal cortex, although some are in frontal cortex. Moreover, the patterns of connectivity between these areas could be ascertained (on average, each area was connected to about 10 others) and feedforward, feedback, and collateral connections distinguished so as to demarcate a processing hierarchy. For about a third of these areas there was sufficient information to make strong claims about the kinds of information processing they performed. Figure 3a from van Essen and Gallant (1994) shows these areas on a flattened map of monkey cortex. Operations have been associated with many of these areas (but not for those shown in medium grey), and Figure 3b uses visual icons to indicate these operations. Although substantial details remain to be filled in, and further research is likely to lead to modifications in the sketch already provided, the results of looking downward into the visual system has been the identification of many of the parts and operations of the visual mechanism in primates. Looking down has clearly been enormously productive.

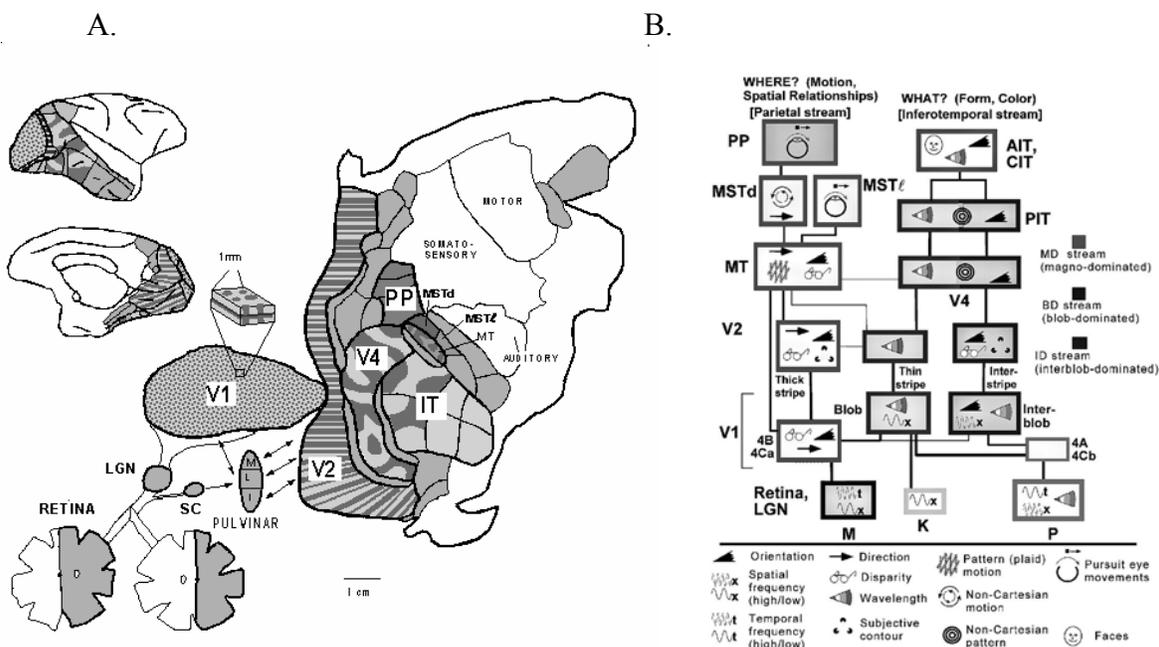


Figure 3. A. Van Essen and Gallant's flat map of the right hemisphere of the macaque identifying major subcortical and cortical visual processing areas. The blob and interblob structure as well as layer 4B are differentiated on the expanded representation of V1 while the thin, thick, and interstripe regions of V2 are shown. At the upper left are lateral and medial views of the right hemisphere showing where the respective cortical areas are on the three-dimensional cortex. B. van Essen and Gallant's iconic representation of the organization of multiple processing streams in the macaque, with boxes indicating different processing areas and icons representing the types of information to which cells in each area are responsive. M, P, and K refer to the magnocellular, parvocellular, and K streams identified in the retina and LGN. AIT, CIT, and PIT refer to the anterior, central, and posterior portions of the IT complex. The thickness of lines indicates the relative degree of connectedness of the various areas. Note that there are connections between the What and Where streams at a variety of levels.

2. Looking Around: The Organization of Mechanisms

An assumption commonly made at the outset of research on a given phenomenon is the simple localizationist assumption that a specific part of the system is responsible for that phenomenon. For example, Lashley assumed that striate cortex performed all specifically visual processing. Even as more visual areas were identified, they were often characterized in terms of the types of visual information they carried, often providing the misleading suggestion that they alone were responsible for the specific information processing. This is not a unique feature of brain research. Following Mendel's lead, genetics, for example, began by characterizing genes for the traits in which they were expressed and one still finds references to the genes for specific traits, including genes for specific disorders. An implicit assumption is that these brain areas or genes themselves are responsible for a given mental activity or trait and that the whole system is just a collection or aggregation of independent factors. Following Fodor's (1983) characterization of cognitive input systems as informationally encapsulated modules, evolutionary psychology construed the mind as a aggregation of specialized processors that evolved separately (Barkow, Cosmides, & Tooby, 1992; D. D. Cummins & Allen, 1998).

Just as simple localization often turns out to be false, so does the aggregative view of organization. A first step beyond independent aggregated components often involves proposing that some components depend upon others and trying to fit the components into a linear pathway of processing components. The model here is the assembly lines employed in human manufacturing in which components are added to the product one at a time. Thus, once researchers recognized that fermentation was not a simple reaction due to a single enzyme (which Buchner had called *zymase*) and that there were multiple chemical intermediates produced in the conversion of sugar to alcohol, biochemists tried to organize the processes into a linear pathway of successive transformations of a substrate. The main reactions of fermentation are still represented in such a fashion despite the fact that some of the necessary metabolites required in early reactions (ATP) are generated by reactions later in the pathway. In other cases, such as the reactions of oxidative metabolism, it proved impossible to link the reactions into a linear pathway, leading Thunberg and later Krebs to propose cyclic organization, a mode of organization of major importance in biology (Bechtel & Abrahamsen, in press-a).

After Hubel and Wiesel determined that LGN cells responded in a center-surround fashion whereas simple cells in V1 responded to oriented edges at specific locations and complex cells responded to oriented edges at different locations in the receptive field or edges moving in a specific direction, they proposed processing procedures whereby a given simple cell would sum inputs from several LGN cells whose center lay along a line with a given orientation and respond when these inputs exceeded a threshold (Figure 4a). They further proposed that a complex cell that fired whenever a line with an appropriate orientation anywhere in its visual field received inputs from several simple cells and responded when any of its inputs was active (see Figure 4b). In terms of logic, the simple cells functioned as *and-gates* and the complex cells as *or-gates*.

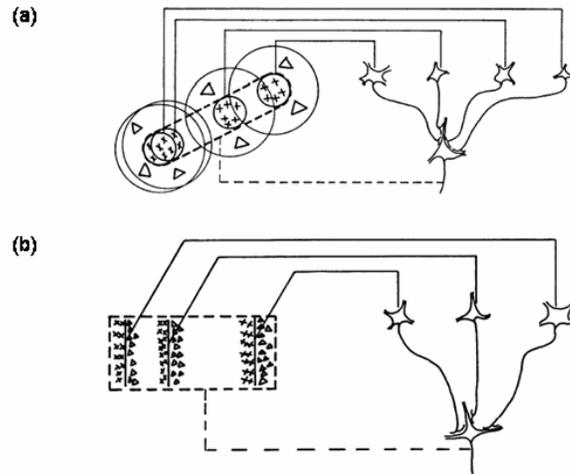


Figure 4. A. Hubel and Wiesel's proposals for wiring diagrams for (a) simple and (b) complex cells. Simple cells receive their inputs from LGN cells with on centers that are aligned so that when a bar of light crosses the receptive fields of all the linked LGN cells, the simple cell fires. Complex cells receive their input from simple cells which are responsive to edges oriented in the appropriate way at any point in their receptive fields. Reprinted from Hubel, D. H., & Wiesel, T. N. (1962), Receptive fields, binocular interaction and functional architecture in the cat's visual cortex, *Journal of Physiology*, 160, p. 142, text figure 19, and p. 143, text figure 20.

Hubel and Wiesel's model is a linear one, and although it provides a suggestive account of how simple cells compute edges, there are a number of limitations researchers have identified with such linear models. For example, while simple cells saturate, so that their response levels off as the contrast between center and surround becomes great, a linear model would generate a continually increasing response over the entire range. Further, when two patterns are overlaid, one representing the cell's preferred stimulus, a simple cell's response is diminished, whereas a linear model predicts it would respond normally to the preferred stimulus. This led Heeger, Simoncelli, and Movshon (1996) to develop non-linear models that normalize the predicted response by dividing the summed input by a factor representing the pooled response of nearby neurons. They then extended the application of these non-linear models to MT cells which respond not just to individual components of motion, but to complex patterns. According to their model, a given MT cell sums inputs from V1 cells compatible with the composite motion and normalizes its response by dividing the summed input by the pooled responses of other MT cells.

The general idea underlying these different proposals is that successive processing regions in the brain extract information by being appropriately linked to those in other areas which have executed simpler stages of information extraction. This in fact fits the pattern of discovery recounted in the previous section where researchers, recognizing the limited nature of processing in a given brain area, proceeded forward in the brain to find other areas that carried out additional levels of processing. Already with the differentiation of V4 and MT, this research pointed to a bifurcation in the types of information being processed, with V4 processing static information pertinent to identification of the stimulus and MT processing information about its motion. Subsequent research in temporal and parietal cortex maintained this differentiation and, as shown in Figure 5, in the early 1980s Ungerleider and Mishkin (1982) proposed that visual processing was segregated into two pathways, one projecting ventrally into the temporal cortex

that processes information about the identity of an object and another projecting dorsally into the parietal cortex that processes information about where an object was.⁵

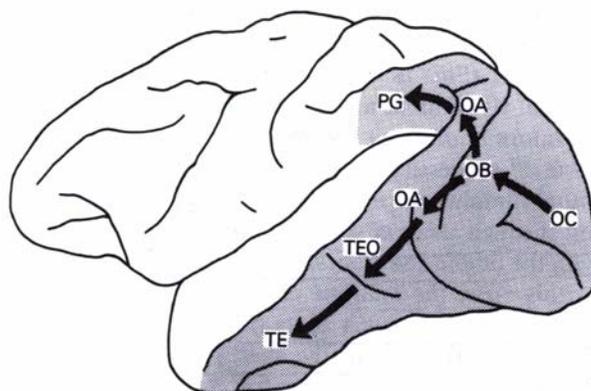


Figure 5. Two pathways of visual processing in the rhesus monkey proposed by Mishkin, Ungerleider, and Macko (1983). Each begins in area OC (primary visual cortex or V1) and projects into extrastriate areas OB (V2) and OA (V3, V4, and MT). What they called the *what* pathway then projects ventrally into inferior temporal cortex (areas TEO and TE), whereas the *where* pathway projects dorsally into inferior parietal cortex (area PG).

Although Ungerleider and Mishkin's account proposed that the two pathways originated in a common processing system in V1, Livingstone and Hubel (1984) soon advanced a model that differentiated the pathways all the way back to the retina. Earlier research (Enroth-Cugell & Robson, 1966; Dreher, Fukada, & Rodieck, 1976) had identified different populations of cells linked in pathways in the retina and LGN of the cat and Old World monkeys. Livingstone and Hubel, relying on a stain for cytochrome oxidase, differentiated populations of cells in V1 that were orientation sensitive and wavelength sensitive. Other researchers, however, challenged that sharpness of the differentiation of the pathways. While Figure 4b, from van Essen and Gallant, reveals an overall division into *what* and *where* processing, it also indicates a number of points where there are projections from one pathway into the other. Van Essen and Gallant also note that there are extensive interconnections between the pathways at the level of the LGN, and these enable processing in both pathways to continue even if the supposedly specific subcortical input to one pathway is removed. Accordingly, they characterize the organization not in terms of pathways but in terms of streams that partly interconnect. During the 1990s further research on patients with damage to the parietal lobe also lead to recharacterizing the information processing it performs: instead of characterizing it as simply processing information about location, Milner and Goodale (1995) characterize it as extracting information specific to orienting one's body so as to act.

As I have emphasized, Ungerleider and Mishkin's proposals and their descendants construe processing as feedforward: each successive brain area performs another step in information processing. Researchers have long recognized, though, that most feedforward projections are

⁵ Working with hamsters, Gerald Schneider (1967) had proposed a similar pair of pathways, with object discrimination requiring a pathway from geniculate areas to cortex and location processing relying on tectofugal regions. Colin Trevarthen (1968) proposed a similar division of pathways for primates. The difference in Ungerleider and Mishkin's proposal is that both pathways are cortical.

accompanied by feedback or recurrent projections. While there are differences between the layers in cortex from which feedforward and recurrent projections originate and where they connect, which enable the identification of a hierarchy of processing areas (Felleman & van Essen, 1991), the frequency of feedback and collateral projections suggests that feedforward accounts fail to provide a sufficient account of how information is processed in the visual system.

Understanding the behavior of systems with complex interactions is extremely challenging for humans. When we try to understand a process involving multiple operations, we typically attempt to follow operations step-by-step in a feedforward fashion. Such reasoning is facilitated by our conscious thought, in which one thought succeeds another in linear fashion. This makes it very difficult to understand systems in which an operation that we take to be later in a process has effects on those that we represent as earlier. As a result, one of the simplest forms of organization in which processes later in a pathway turn off or down regulate those earlier, negative feedback, proved extremely challenging to human theorists. The principle appears to have been first employed by Ktesibios of Alexandria around 270 BCE as part of his design for a water clock. To ensure constant flow of water into a vesicle in which time was measured by the height of the water, Ktesibios inserted a second vesicle into the supply pathway with a float that would rise and block the flow of new water into it whenever the water in the vesicle reached the target height and drop to admit new water whenever the water in the vesicle dropped below that height. Although this was an ingenious solution to a specific problem, the underlying principle of negative feedback was not recognized as applicable to other domains and, as Mayr (1970) documents, had to be rediscovered context by context. After James Watt rediscovered the principle in 1788 in designing a governor for a steam engine and James Clerk Maxwell (1868) developed a mathematical analysis of such governors, the principle came to be more widely recognized. Finally, in the middle of the 20th century it was championed as a fundamental principle in the design of biological and social as well as engineered systems by the cybernetics movement (Wiener, 1948) and construed as providing a foundation for resuscitating notions such as *purpose* and *teleology* in biology (Rosenblueth, Wiener, & Bigelow, 1943).

Even as negative feedback began to gain acceptance as an organizational principle, positive feedback continued to be rejected as leading only to run away processes. For example, when Boris Belousov identified a cyclic process that generated an oscillatory reaction, his paper was rejected as describing something impossible. Eventually, however, theorists discovered how coupled autocatalytic reactions could generate systems that oscillated until their supply of reactants was exhausted (Field, Koros, & Noyes, 1972). The discovery of modes of organization such as autocatalytic cycles in the later part of the 20th century challenged human intuitions about what can be accomplished in systems with non-linear modes of organization and revealed how little we understand about what different modes of organization make possible. Two advances at the end of the century reveal some of this potential. Duncan Watts (Watts & Strogatz, 1998), drawing upon earlier social psychology experiments by Stanley Milgram and the subsequent popular lore about six degrees of separation (according to which on average a given human being can be linked to any other human being through a chain of five acquaintances), showed that one could achieve both the high-clustering of components as found in regular lattices and the short path-length found in random networks by adding just a few long-range connections to a regular lattice. The resulting “small worlds” organization is one that may be optimal for information

processing and that can be realized via Hebbian learning (Gong & van Leeuwen, 2004). Most models for network design have assumed that the number of connections per unit is distributed roughly randomly, but in many natural systems, including biochemical systems, the distribution follows a power law such that most units have very few connections while a few are very highly connected (Barabási & Albert, 1999). Such networks are robust against the loss of most components.

The discovery of different modes of organization with effects very different from those found in linearly-organized systems has provided new tools for understanding a long-noted but previously hard to explain capacity of living systems—the ability to be self-active and self-maintaining. This capacity begins to seem far less peculiar when one considers how a mechanism might be organized to accomplish one of the most basic requirements of living systems—to maintain themselves as systems far from thermodynamic equilibrium despite the tendency of such systems to dissipate. Doing so requires that organisms recruit matter and energy from their environment and utilize it to remake themselves, achieving what Ruiz-Mirazo, Peretó, & Moreno (2004, p. 330) characterize as an autonomous system:

a far-from-equilibrium system that constitutes and maintains itself establishing an organizational identity of its own, a functionally integrated (homeostatic and active) unit based on a set of endergonic-exergonic couplings between internal self-constructing processes, as well as with other processes of interaction with its environment.

A consequence of being so organized as to maintain themselves by continually recruiting matter and energy is that living systems, starting with individual cells, are inherently active, not merely responsive, systems.

The extensive interconnection of brain areas with feedforward, recurrent, and collateral projections suggests that it too might be a system that is inherently active, generating patterns of neural activity on its own (Buzsáki, 2006). If it is an active system, then it may be better to think of sensory inputs as perturbing the brain's internal dynamics rather than as providing inputs to which the brain passively responds. A few theorists have been advocating such active models of perception (Ballard, 1991; Churchland, Ramachandran, & Sejnowski, 1994), but it remains challenging to comprehend how such a system might function.

One very suggestive approach has been advanced by Cees van Leeuwen and his collaborators (van Leeuwen, Steyvers, & Nooter, 1997), who develop a neural network model that begins with active units. By using the logistic map function

$$a_{t+1} = A a_t (1 - a_t)$$

as the basis for the activation function for their units, they model units which both maintain activity and have the potential for internally driven variation. Figure 6 shows the logistic map function for varying values of the parameter A; for values of $A < 3.0$ a unit eventually settles into a fixed activation, but for most values between 3.0 and 3.6 it enters into periodic oscillations and for most values above 3.6 it exhibits deterministic chaos. The regions of transition between stable oscillations and chaos turn out to be the most interesting. By coupling such units, van Leeuwen has developed networks of coupled oscillators that are capable of synchronizing their behavior for periods before spontaneously desynchronizing their behavior. The coupling is accomplished by basing the netinput to a unit on its own previous activation and that of the units with which it is coupled

$$netinput_x = \sum C a_y + (1 - C) a_x.$$

The coupling parameter C determines how much the unit is affected by other units. The netinput is then supplied to the logistic map function

$$a_{x,t+1} = A netinput_{x,t} (1 - netinput_{x,t}).$$

The behavior of spontaneous synchronization and desynchronization between units is determined by the parameters A and C . van Leeuwen treats A as an input to each unit (which is construed as having a distinctive receptive field) and C as an adjustable weight between units. When an unambiguous visual input is provided, the dynamics are in a stable regime and the relevant units in the network synchronize their activity (construed as interpreting the stimulus). When the input is ambiguous, however, the synchronization is temporary and the internal dynamics of being in a chaotic regime result in periods of desynchronization followed by a different synchronization (construed as an alternative interpretation of the stimulus).

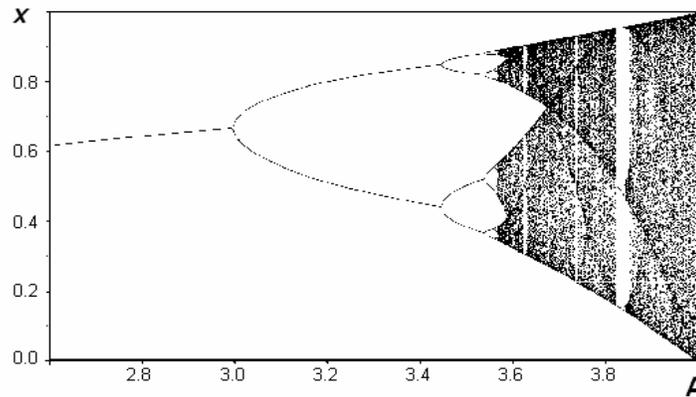


Figure 6. Plot of the logistic function for values of A between 2.6 and 4.0.

Van Leeuwen's network provides an illustrative model of how dynamics in a network can enable it to be an internally active system whose behavior is largely determined by internal dynamics, but is nonetheless capable of being modulated by external inputs. The interconnectivity in the primate visual system is many times more complicated than in van Leeuwen's networks and provides the potential for much more complex internally driven behavior, including the internally controlled sampling of visual input that researchers advocating active vision have emphasized is characteristic of our visual behavior. Researchers are only beginning to develop the tools for understanding the range of complex dynamics that may arise in the visual system and what such dynamics makes possible; the task of understanding the significance of the organization in visual mechanisms remains largely a project for the future.

3. Looking Up: Situated Mechanisms

In the previous section I emphasized the potential for internally driven activity in mechanisms with appropriate modes of organization. Although this renders the resulting mechanisms no longer purely reactive systems, they remain subject to modulation from the environment. This means that understanding the behavior of a mechanism requires focusing not just on its lower level constitution and its organization, but also on the specific character of the inputs it receives from its environment. This is clearly true of the visual system, whose primary function is to secure for organisms information about their distal environments, much of which is directly relevant to acting in that environment.

Although the dependency of the behavior of mechanisms on environmental conditions is widely recognized, there is a high risk in mechanistic research of underestimating the significance of this environment. Many mechanisms in their natural settings are complexly related to their environment but in order to understand their operation it is common to extract the mechanism from its natural environment and examine its responsiveness in artificial, experimental environments. Since typically in the natural environment too many factors are operative at once to enable researchers to determine their individual effects on the mechanism's behavior, such experimentation is critically important. But researchers need also to attend to how the experimental context is eliminating inputs that may be critical to the particular behaviors they want to explain. The importance of context for visual tasks is easily demonstrated; in Figure 7 the middle letters in both the words *the* and *cat* are written the same but are interpreted very differently. By considering the letter alone as the stimulus one would not be able to explain how people respond to it.



Figure 7. The middle letter in each word is printed the same and without context would be ambiguous. In context, however, it is viewed as an *H* in the first word and as an *A* in the second. Figure after Selfridge (1955).

James J. Gibson (1979) and researchers in the tradition of ecological psychology which he initiated have long emphasized the importance of carefully analyzing the visual context of perception. Against the view that perception requires complex information processing to build up from single stimuli to an understanding of the environment, Gibson emphasized the richness of *information in the light* that specified features of the distal environment critical to the functioning of the organism. One feature that Gibson emphasized is that perceivers are typically not stationary, but moving in their environment. As a result, important elements of the information they acquire through vision is dependent on their own motion. Accordingly, Gibson focused on *optic flow*, the manner in which the visual scene changes as a result of the relative movement of the perceiver and objects in the environment.

It is often challenging to determine the specific information to which a visual system is responding. I noted in the context of discussing what stimulus would drive individual cells in various visual areas that results were dependent on the particular stimuli researchers thought to test. Even when it is clear that an organism acquires information about a feature of its environment, it remains a challenge to figure out the source from which they procure the knowledge. Baseball players, for example, are highly skilled at catching fly balls, which requires both correct positioning and timing. Gannets that dive to catch fish in the water reliably fold their wings at just the right point to enter the water (folding late would cause serious damage to the wings, whereas folding earlier would cost control of flight). David Lee (1981) identified a simple factor, *tau* (the inverse of the rate of expansion of the object in an organism's perceptual field), that specifies time to contact. Tanaka and Saito's (1989) subsequent discovery of cells in area MSTd that respond to the expansion or contraction of stimuli in the visual field provided information about the mechanism that could enable organisms to utilize such information.

Behind providing input to a mechanism, the environment often has important systemic properties that are missed if one simply looks at inputs to the mechanism sequentially. That is, the environment is not just a collection of stimuli, but of objects behaving in systematic ways. Simple examples involve objects behaving in accord with classical physical principles—unsupported objects fall with a constantly increasing velocity. If one observed an object falling at a constant or declining velocity, one would be alerted to the fact that something unusual was affecting the process. These regularities enable organisms to develop expectations from simple cues. Sometimes these expectations are foiled, as in the room designed by Adelbert Ames (1951) in which the wall and windows opposite the restricted viewing point are trapezoids, not rectangles. The result is that people moving about the room seem to grow or shrink. The difficulty in resisting the illusion produced by the Ames room indicates how perceivers rely on such expectations about their environment.

Many mechanisms, including cognitive mechanisms, are capable of taking advantage of the regularities in the environment to advance their own ends. If the mechanism can utilize resources provided by its environment, it does not need to carry out all operations on its own. The common practices of performing calculations by creating and responding to inscriptions on a page (or using tools such as the abacas) illustrate how we utilize such resources. This practice is becoming even more common in our highly technological world which provides many prosthetic devices for performing our cognitive work (Clark, 2003). Moreover, it is not just our physical environment, but also our social environment, in which there are systemic relations that provide resources to cognitive agents (e.g., the Pedel enters the room and announces “hora est”, thereby signaling to the committee and candidate that the dissertation defense has ended and obviating the need for anyone within to keep track of time).

The behavior of mechanisms is highly dependent on conditions in their environments, including any regularities that occur there. But these are not discovered by looking inside the mechanism to the parts and operations or how these are organized. They must be discovered by examining the environment in which the mechanism operates and employing tools appropriate for such inquiry.

4. Conclusion: Mechanism as an Integrated, Multi-level Perspective

The quest to explain phenomena by identifying responsible mechanisms involves an inherent reductionistic commitment—such research decomposes the mechanism into its parts and their operations. But unlike more traditional philosophical accounts of reduction, the mechanistic perspective is not exclusively reductionistic, for it requires also recomposing the mechanism by taking into account the organization among the parts of the mechanism and situating the mechanism in its environment. Studying the parts and operations, organization, and situatedness of a mechanism requires different sets of investigatory tools. Understanding mechanisms requires a pluralistic approach that looks not just down but around and up.

One of the sharpest distinctions between the mechanistic perspective on reduction and more traditional accounts involves the attitude about how complete the knowledge available at any given level is. In nomological accounts of reduction, the lowest level is often viewed as itself providing a complete account of everything that happens: if one provides the appropriate initial

conditions, then one can derive from lower-level laws all the consequences. Appealing to higher-level explanations then appears problematic since any causal processes at higher levels would overdetermine events already determined by lower-level processes. In particular, appealing to mental causes for events determined by brain processes appears problematic (Kim, 1998). From a mechanist perspective, there is no basis for assuming that one could provide a complete account of the functioning of the mechanism in terms of the parts alone. The behavior of the mechanism depends not just on the parts but how they are organized and the context in which they are situated. In biological systems, even the behavior of the parts themselves is often affected by the organization and environment in which they function and learning about such behavior requires studying the part in such a context. Although researchers can sometimes predict how a part will behave in a new context given what they have found in other contexts, such predictions are often wrong when the parts are at all complex and plastic. Far from being complete, lower-level accounts are often extremely fragmentary, providing accounts of parts of mechanisms only under the conditions that researchers have had reason and tools to consider.

Concerns about causal overdetermination have been invoked specifically in criticism of appeals to *downward* or *top-down* causation, such as advocated by Donald Campbell (1974). There is something right about the objections to top-down causation even if lower-level accounts are not complete. Changes in the components of a mechanism are typically the result of interactions between them and other components of the mechanism. Appeal to higher-level causes seems to overdetermine the effects. Nonetheless, the point behind the appeals to top-down causation is also legitimate—components of a mechanism are affected by factors that causally impact the mechanism as a whole. Carl Craver and I (2007) have argued that these two perspectives can be reconciled if, instead of appealing to relations between levels as causal, we understand them only as constitutive and think in terms of *mechanistically mediated effects*. In cases that appear to involve top-down causation, we identify the mechanism with its total constitution and realize that when the mechanism is affected by things external, so are some of its parts. The account also works for apparent cases of bottom-up causation: the changes in the parts may generate a cascade of changes within the mechanism, resulting in the mechanism as a whole being differently constituted and having different effects on things external to it. The constitutive relations within a mechanism mediate between causal processes within parts of the mechanism, those between the parts, and those between the mechanism and its environment. And, as I have been stressing throughout this paper, each of these requires independent investigation with the appropriate tools and techniques, requiring mechanistic researchers to look down, around, and up.

References

- Ames, A. (1951). Visual perception and the rotating trapezoidal window. *Psychological Monographs*, 65(7), 1-32.
- Andersen, R. A., Essick, G. K., & Siegel, R. M. (1985). Encoding of spatial location by posterior parietal neurons. *Science*, 230, 456-458.
- Bálint, R. (1909). Seelenlähmung des Schauens, optische Ataxie, räumliche Störung der Aufmerksamkeit. *Monatschrift für Psychiatrie und Neurologie*, 25, 5-81.
- Ballard, D. H. (1991). Animate vision. *Artificial Intelligence*, 48, 57-86.
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286, 509-512.
- Barkow, J. H., Cosmides, L., & Tooby, J. (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. New York: Oxford.
- Barlow, H. B. (1953). Summation and inhibition in the frog's retina. *Journal of Physiology*, 119, 69-88.
- Bechtel, W. (2006). *Discovering cell mechanisms: The creation of modern cell biology*. Cambridge: Cambridge University Press.
- Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. London: Routledge.
- Bechtel, W. (in press-a). The downs and ups of mechanistic research: Circadian rhythm research as an exemplar. *Erkenntnis*.
- Bechtel, W. (in press-b). The epistemology of evidence in cognitive neuroscience. In R. Skipper, C. Allen, R. A. Ankeny, C. F. Craver, L. Darden, G. Mikkelsen & R. C. Richardson (Eds.), *Philosophy and the Life Sciences: A Reader*. Cambridge, MA: MIT Press.
- Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36, 421-441.
- Bechtel, W., & Abrahamsen, A. (in press-a). Complex biological mechanisms: Cyclic, oscillatory, and autonomous. In C. A. Hooker (Ed.), *Philosophy of complex systems. Handbook of the philosophy of science, Volume 10*. New York: Elsevier.
- Bechtel, W., & Abrahamsen, A. (in press-b). Decomposing, recomposing, and situating circadian mechanisms: Three tasks in developing mechanistic explanations. In H. Leitgeb & A. Hieke (Eds.), *Reduction and elimination in philosophy of mind and philosophy of neuroscience*. Frankfurt: Ontos Verlag.
- Bechtel, W., & Richardson, R. C. (1993). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton, NJ: Princeton University Press.
- Bickle, J. (2003). *Philosophy and neuroscience: A ruthlessly reductive account*. Dordrecht: Kluwer.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, 12, 4745-4765.
- Brodmann, K. (1909/1994). *Vergleichende Lokalisationslehre der Grosshirnrinde* (L. J. Garvey, Trans.). Leipzig: J. A. Barth.
- Brown, S., & Schäfer, E. A. (1888). An investigation into the functions of the occipital and temporal lobes of the monkey's brain. *Philosophical Transactions of the Royal Society of London*, 179, 303-327.

- Buzsáki, G. (2006). *Rhythms of the brain*. Oxford: Oxford University Press.
- Campbell, D. T. (1974). 'Downward causation' in hierarchically organised biological systems. In F. J. Ayala & Dobzhansky (Eds.), *Studies in the philosophy of biology*: Macmillan Press Ltd.
- Churchland, P. S., Ramachandran, V. S., & Sejnowski, T. J. (1994). A critique of pure vision. In C. Koch & J. L. Davis (Eds.), *Large-scale neuronal theories of the brain* (pp. 23-60). Cambridge: MIT Press.
- Clark, A. (2003). *Naturally-born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford: Oxford University Press.
- Cowey, A. (1964). Projection of the retina on to striate and prestriate cortex in the squirrel monkey *Saimiri Sciureus*. *Journal of Neurophysiology*, 27, 366-393.
- Craver, C. (2007). *Explaining the brain: What a science of the mind-brain could be*. New York: Oxford University Press.
- Craver, C., & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology and Philosophy*, 22, 547-563.
- Cummins, D. D., & Allen, C. (1998). Introduction. In D. D. Cummins & C. Allen (Eds.), *The evolution of mind* (pp. 3-8). Oxford: Oxford University Press.
- Cummins, R. (2000). "How does it work?" versus "what are the laws?": Two conceptions of psychological explanation. In F. Keil & R. Wilson (Eds.), *Explanation and cognition* (pp. 117-144). Cambridge, MA: MIT Press.
- Darden, L. (2006). *Reasoning in biological discoveries: Essays on mechanisms, interfield relations, and anomaly resolution*. Cambridge: Cambridge University Press.
- Dreher, B., Fukada, Y., & Rodieck, R. W. (1976). Identification, classification and anatomical segregation of cells with X-like and Y-like properties in the lateral geniculate nucleus of old-world primates. *Journal of Physiology*, 258, 433-452.
- Enroth-Cugell, C., & Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *Journal of Physiology*, 187, 517-552.
- Felleman, D. J., & van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1-47.
- Ferrier, D. (1876). *The functions of the brain*. London: Smith, Elder, and Company.
- Ferrier, D., & Yeo, G. F. (1884). A record of the experiments on the effects of lesions of different regions of the cerebral hemispheres. *Philosophical Transactions of the Royal Society of London*, 175, 479-564.
- Field, R. J., Koros, E., & Noyes, R. M. (1972). Oscillations in chemical systems. II. Thorough analysis of temporal oscillation in the Bromate–Cerium–Malonic acid system. *Journal of the American Chemical Society*, 94, 8649–8664.
- Fodor, J. A. (1974). Special sciences (Or: Disunity of science as a working hypothesis). *Synthese*, 28, 97-115.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Glennan, S. (1996). Mechanisms and the nature of causation. *Erkenntnis*, 44, 50-71.
- Gong, P., & van Leeuwen, C. (2004). Evolution to a small-world network with chaotic units. *Europhysics Letters*, 67, 328-333.
- Gross, C. G., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. *Journal of Neurophysiology*, 35, 96-111.

- Hartline, H. K. (1938). The response of single optic nerve fibers of the vertebrate eye to illumination of the retina. *American Journal of Physiology*, *113*, 400-415.
- Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (1996). Computational models of cortical visual processing. *Proceedings of the National Academy of Sciences (USA)*, *93*, 623-627.
- Hempel, C. G. (1965). Aspects of scientific explanation. In C. G. Hempel (Ed.), *Aspects of scientific explanation and other essays in the philosophy of science* (pp. 331-496). New York: Macmillan.
- Henschen, S. E. (1893). On the visual path and centre. *Brain*, *16*, 170-180.
- Holmes, G. M. (1919). Lecture I--The cortical localization of vision. *British Medical Journal*, *2*, 193-199.
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, *160*, 106-154.
- Hubel, D. H., & Wiesel, T. N. (1965). Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat. *Journal of Neurophysiology*, *195*, 229-289.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, *195*, 215-243.
- Hyvärinen, J., & Poranen, A. (1974). Function of the parietal associative area 7 as revealed from cellular discharges in alert monkeys. *Brain*, *97*, 673-692.
- Inouye, T. (1909). *Die Sehstörungen bei Schussverletzungen der kortikalen Sehsphäre nach Beobachtungen an Verwundeten der letzten japanischen Kriege*. Leipzig: Engelmann.
- Kim, J. (1998). *Mind in a physical world*. Cambridge, MA: MIT Press.
- Klüver, H., & Bucy, P. (1938). An analysis of certain effects of bilateral temporal lobectomy in the rhesus monkey, with special reference to "psychic blindness". *The Journal of Psychology*, *5*, 33-54.
- Kuffler, S. W. (1953). Discharge patterns and functional organization of mammalian retina. *Journal of Neurophysiology*, *16*, 37-68.
- Lashley, K. S. (1950). In search of the engram. *Symposia of the Society for Experimental Biology, IV. Physiological Mechanisms in Animal Behaviour*, 454-482.
- Lee, D., & Redish, P. E. (1981). Plummeting gannets: A paradigm of ecological optics. *Nature*, *293*, 293-294.
- Lissauer, H. (1890). Ein Fall von Seelenblindheit nebst einem Beitrag zur Theorie derselben. *Archiv für Psychiatrie*, *21*, 222-270.
- Livingstone, M. S., & Hubel, D. H. (1984). Anatomy and physiology of a color system in the primate visual cortex. *Journal of Neuroscience*, *4*, 309-356.
- Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, *67*, 1-25.
- Marr, D. C. (1982). *Vision: A computation investigation into the human representational system and processing of visual information*. San Francisco: Freeman.
- Maxwell, J. C. (1868). On governors. *Proceedings of the Royal Society of London*, *16*, 270-283.
- Mayr, O. (1970). *The origins of feedback control*. Cambridge, MA: MIT Press.
- Meynert, T. (1870). Beiträge zur Kenntniss der centralen Projection der Sinnesoberflächen. *Sitzungsberichte der Kaiserlichen Akademie der Wissenschaften, Wien. Mathematisch-Naturwissenschaftliche Classe*, *60*, 547-562.
- Milner, A. D., & Goodale, M. G. (1995). *The visual brain in action*. Oxford: Oxford University Press.

- Mishkin, M. (1966). Visual mechanisms beyond the striate cortex. In R. W. Russell (Ed.), *Frontiers in physiological psychology*. New York: Academic.
- Mundale, J. (1998). Brain mapping. In W. Bechtel & G. Graham (Eds.), *A companion to cognitive science*. Oxford: Basil Blackwell.
- Munk, H. (1881). *Über die Funktionen der Grosshirnrinde*. Berlin: A. Hirschwald.
- Nagel, E. (1961). *The structure of science*. New York: Harcourt, Brace.
- Panizza, B. (1856). Osservazioni sul nervo ottico. *Memoria, Istituto Lombardo di Scienze, Lettere e Arte*, 5, 375-390.
- Rosenblueth, A., Wiener, N., & Bigelow, J. (1943). Behavior, purpose, and teleology. *Philosophy of Science*, 10, 18-24.
- Ruiz-Mirazo, K., Peretó, J., & Moreno, A. (2004). A universal definition of life: Autonomy and open-ended evolution. *Origins of Life and Evolution of the Biosphere*, 34, 323-346.
- Schäfer, E. A. (1888). Experiments on special sense localisations in the cortex cerebri of the monkey. *Brain*, 10, 362-380.
- Schneider, G. E. (1967). Contrasting visuomotor functions of tectum and cortex in the golden hamster. *Psychologische Forschung*, 31, 52-62.
- Selfridge, O. G. (1955). Pattern recognition and modern computers. In *Proceedings of the Western Joint Computer Conference* (pp. 94-97). New York: Institute of Electrical and Electronics Engineers.
- Talbot, S. A., & Marshall, W. H. (1941). Physiological studies on neural mechanisms of visual localization and discrimination. *American Journal of Ophthalmology*, 24, 1255-1263.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience*, 19, 109-140.
- Tanaka, K., & Saito, A. (1989). Analysis of motion of the visual field by direction, expansion/contraction, and rotation cells clustered in the dorsal part of the medial superior temporal area of the macaque monkey. *Journal of Neurophysiology*, 62, 626-641.
- Thagard, P. (2003). Pathways to biomedical discovery. *Philosophy of Science*, 70, 235-254.
- Trevarthen, C. (1968). Two mechanisms of vision in primates. *Psychologische Forschung*, 31, 299-337.
- Ungerleider, L. G., & Mishkin, M. (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.
- van Essen, D. C., & Gallant, J. L. (1994). Neural mechanisms of form and motion processing in the primate visual system. *Neuron*, 13, 1-10.
- van Leeuwen, C., Steyvers, M., & Nooter, M. (1997). Stability and intermittency in large-scale coupled oscillator models for perceptual segmentation. *Journal of Mathematical Psychology*, 41, 319-344.
- Watts, D., & Strogatz, S. (1998). Collective dynamics of small worlds. *Nature*, 393(440-442).
- Weber, M. (2005). *Philosophy of experimental biology*. Cambridge: Cambridge University Press.
- Wiener, N. (1948). *Cybernetics: Or, control and communication in the animal and the machine*. New York: Wiley.
- Zeki, S. M. (1969). Representation of central visual fields in prestriate cortex monkey. *Brain Research*, 14, 271-291.
- Zeki, S. M. (1971). Cortical projections from two prestriate areas in the monkey. *Brain Research*, 34, 19-35.