

Mental Mechanisms: What are the Operations?

William Bechtel (bill@mechanism.ucsd.edu)

Department of Philosophy and Science Studies Program,
University of California, San Diego, La Jolla, CA 92093-0119

Abstract

An important goal of the cognitive sciences is to explain mental phenomena in terms of mechanisms. A mechanistic explanation requires characterizing the operations of the mechanism's parts. The challenge in characterizing such operations is illustrated by an example from biology in which some investigators tried to characterize the internal operations in the same terms as the overall physiological system while others appealed to elemental chemistry. Before biochemistry became successful, researchers had to identify operations at a new level of organization—operations over molecular groups. Existing attempts at mechanistic explanation in cognitive science are in a situation comparable to the earlier stage in the biological example, drawing their inspiration either from overall psychological activities or from low-level neural processes, neither of which is likely to provide a successful account of the operations in mental mechanisms.

Keywords: Mechanism; explanation; operations; levels of organization; connectionism; symbolic models

Introduction

Philosophical accounts of cognitive science still commonly treat explanation as a matter of subsuming descriptions of cognitive phenomena under laws in the manner characterized by the deductive-nomological (D-N) model of explanation (Hempel, 1965; Suppe, 1977). Cognitive science explanations, however, typically do not appeal to laws. Cummins (2000), for example, argues that laws, or what are commonly called *effects* in psychology, do not explain, but are what require explanation. Instead of appealing to laws, explanations in cognitive science, as in most of the life sciences, appeal to mechanisms. Recently several philosophers of biology have attempted to articulate the conception of mechanism and mechanistic explanation that figures in these sciences. In the first part of this paper I extend this account to cognitive science.

In order to construct successful mechanistic explanations, a discipline requires an understanding of the types of operations out of which explanatory accounts can be constructed. Historically developing the appropriate account of operations for a given inquiry has proven challenging. In section two, I illustrate this with an example from biochemistry in which initial attempts to identify the operations that figure in physiological processes focused either on too low or too high a level of organization. Only once investigators learned to identify operations at the appropriate level of organization did biochemistry develop into the successful science that we know today. In the final section I will argue that in fact cognitive science is in the position biochemistry was prior to the articulation of an appropriate catalog of operations—it too is looking for operations at too low or too high

a level of organization and awaits the discovery of the nature of operations at the level required for successful explanation.

Mechanistic Explanations

As philosophers of science increasingly focused on particular sciences in the 1970s, philosophers focusing on biology noted the paucity of laws in biology. Some viewed this as a shortcoming of biology (Rosenberg, 1985, 1994), while others maintained that explanation in biology often takes a different form—articulation of mechanisms. Although there is much commonality in the accounts of mechanism that have been advanced (see, for example, Bechtel & Richardson, 1993; Glennan, 2002, 1996; Machamer, Darden, & Craver, 2000), there are differences in terminology, scope, and emphasis. On my account

A mechanism is a structure performing a function in virtue of its components parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena (Bechtel & Abrahamsen, 2005).

A mechanism on this account is a system operating in nature whereas a mechanistic explanation is an epistemic product. To arrive at a mechanistic explanation, scientists must represent (sometimes verbally, but often visually in diagrams) the component parts and their operations and the ways in which they are organized.

A central feature of such mechanistic explanations is that they *decompose* a system that is responsible for a phenomenon into component parts and component operations. (Other features of mechanisms, such as the critical role played by organization and the fact that mechanisms are often constrained by their environments, will not be discussed in this paper.) The parts and operations into which a mechanism is decomposed are closely related: the relevant parts are those that perform operations and hence are *working parts*, or are operated on by working parts. But it is important to distinguish parts understood structurally from operations understood functionally. Although a full understanding of a mechanism requires both the structural and functional perspective, different investigatory techniques are required to establish structural and functional properties of components. As a result, a given group of researchers may only be able to secure evidence about only one or the other. It is then often a challenge to link parts with operations (an activity I refer to elsewhere as *localization*).

The cognitive sciences in general have been in the position of attempting to develop functional decompositions without the benefit of techniques to decompose the brain structurally in relevant ways. In neuroscience, researchers

involved in brain mapping have developed tools for identifying what they hope will turn out to be working parts of the brain, although lacking techniques to link the areas they delineated with cognitive operations. Brodmann (1909/1994), for example, used a variety of cytoarchitectural criteria to differentiate brain areas and expressed optimism that subsequent investigations would be able to relate these areas to mental operations. Contemporary brain mappers (Felleman & van Essen, 1991) apply additional cytoarchitectural tools and have, in domains such as vision, begun to link these to cognitive operations (Bechtel & McCauley, 1999; Bechtel, 2001).

The process of decomposing a mechanism is iterative—the working parts of a mechanism are themselves often mechanisms, and these can in turn be decomposed into their working parts. This iterative process results in the differentiation of successive levels of organization. Such levels are characterized locally within the mechanism, not globally, and are constituted by the parts and operations that are orchestrated to produce the phenomenon in question. (Levels of organization so conceived are very different from levels of analysis advanced by David Marr 1982; for discussion, see Bechtel, 1994). From the point of view of mechanistic explanation, it is important to stress that decomposition of a mechanism proceeds in stages, and only the operations one level down are *directly* relevant in accounting for a given phenomenon. Yet lower levels are relevant for addressing different questions, ones about how the parts perform their operations. Thus, an important step in developing a mechanistic explanation is to identify the operations that constitute a level one down from the phenomenon to be explained.

Discovering the relevant parts and operations into which to decompose a mechanism is often a challenging project. A well-functioning mechanism typically does not reveal either its parts or their operations and experimental interventions are required to reveal them. However, experimental strategies alone do not reveal the appropriate way to decompose the mechanism into operations—that requires developing a conceptual framework that identifies types of operations. Until such a framework is developed, researchers often proceed by trying to characterize internal operations by analogy with what the whole mechanism does or by reaching to a much lower level of organization at which other investigators have already identified a set of operations. In the next section I develop an example from the history of biology that illustrates the problem and the strategies for dealing with it.

Identifying Operations: A Biological Example

The 19th century witnessed a sustained attempt to identify the chemical operations involved in physiological processes such as fermentation and respiration. The chemical revolution at the end of the 18th century resulted in the identification of carbon, hydrogen, oxygen, and nitrogen as the principal elemental constituents of living organisms (Berthollet, 1780; Lavoisier, 1781). With this foundation, investigators began trying to characterize physiological processes in

terms of changes in elemental composition (see Holmes, 1963). For example, Lavoisier (1789) himself characterized fermentation as involving the oxygenation of carbon in part of a sugar molecule, producing carbon dioxide, at the expense of the deoxygenation of the remainder, which resulted in alcohol. The fact that most chemical reactions required in living organisms do not occur freely in the environment led Berzelius (1836) to introduce the concept of chemical catalyst to designate the parts responsible for these operations.

The goal of much early 19th research on the chemical processes occurring in living organisms was to understand nutritional requirements. Direct appeal to elemental composition, however, did not provide a useful way of conceptualizing foodstuffs and how they figured in the animal economy. A more productive approach was developed by Prout (1827), who classified the nutrients required by animals into three classes: saccharine (carbohydrates), oleaginous (fats), and albuminous (proteins). Prout noted that there were only minor differences between the chemical composition of nutrients animals took in from plants and the compounds that comprised the fluids and solids of their bodies. Perhaps the most celebrated chemist of the first half of the 19th century, Justus Liebig drew upon this idea to formulate a central part of his synthetic and highly speculative account of the chemical processes of animals in his *Animal Chemistry* (1842). Since animal tissue was largely comprised of proteins, he proposed that animals simply incorporated proteins from plants into their tissues, whereas they oxidized the carbohydrates and fats in their diet to generate heat. When insufficient oxygen was available for oxidizing carbohydrates, Liebig proposed that animals converted them to fat and stored them. He conjectured that the proteins incorporated into the animal body were broken down and waste products excreted when work was performed. New proteins were thus continually required in animal diets to rebuild animal tissues. With these key ideas, Liebig articulated a general scheme, which he filled in with detailed formulae, that described the chemical reactions occurring in animals.

Liebig's proposal was soon subjected to empirical investigation. One kind of investigation took the form of feeding experiments in which the intake in various food groups, waste products generated, and work performed were measured. Results such as those of Fick and Wislicenus (1866), who made such measurements on themselves as they undertook a climb of Mt. Faulhorn in the Swiss Alps, failed to support Liebig's claim that all energy used resulted from breakdown of protein. Going inside the organism, Claude Bernard (1848) traced chemical changes through the animal body and discovered glycogenesis occurring in the liver. This challenged Liebig's contention that all chemical reactions in animals were catabolic.

Demonstrations by various researchers in the 1830s that fermentation seemed to require living yeast cells, a finding confirmed by Pasteur (1860), cast a general pall over attempts to explain physiological processes chemically. Those who persevered in the project recognized the limitations of trying to explain these reactions in terms of changes in ele-

mental composition. Organic chemists in the later decades of the 19th century determined that chemical compounds were not just composed of atoms but were structured. A consequent was that not every chemical formula designating a combination of elements corresponded to actually occurring substances. This meant it was necessary to take chemical structure into account in explaining physiological processes, thereby moving beyond the project of elemental analysis.

An alternative strategy to building up from elemental composition was to start with a compound such as glucose and try to break it down chemically. In the case of glucose, researchers applied various alkalis to it in the attempt to decompose it into component compounds. Three such compounds were identified—methylglyoxal, glyceraldehyde, and dihydroxyacetone. The identification of these compounds raised the question of whether they might be intermediaries in fermentation. At the end of the 19th century the pursuit of chemical investigations of fermentation were rejuvenated by the serendipitous discovery by Eduard Buchner (1897) that cell extracts in which all whole cells had been removed could still perform fermentation. Although Buchner construed the process as a single reaction transforming glucose to alcohol, which he attributed to an enzyme he named *zymase*, other investigators began to pursue the question of whether methylglyoxal, glyceraldehyde, and dihydroxyacetone might be intermediates in fermentation. What is particularly interesting is how researchers characterized these investigations. They asked whether methylglyoxal, for example, would *ferment* as rapidly as sugar. Abandoning the attempt to explain the processes in elemental terms, they now could only use the same vocabulary as applied to the overall process to describe the possible component operations.

The challenge confronting those seeking to provide chemical explanations of basic physiological processes was to characterize the component operations (reactions) at an appropriate level of organization. Elemental composition was too low a level at which to characterize changes while decomposition into fermentations simply invoked the vocabulary designed to describe the overall behavior to describe component operations. Fortunately for these researchers, at about this same time organic chemistry provided a new framework. Their efforts to determine the structure of organic compounds revealed that they were comprised of groups of molecules such as amino (NH_3^+), carboxyl (COO^-), hydroxyl (OH), and phosphate (PO_4^{--}) groups that were bound to a carbon ring backbone (Holmes, 1992). Reactions would involve whole groups being added, deleted, or moved on the backbone—reactions such as deaminations (removal of an amino group), carboxylations (addition of a carboxyl group), dehydroxylation (removal of an hydroxyl group), phosphorylations (addition of a phosphate group), etc.

This focus on molecular groups provided the basis for conceptualizing types of reactions at a level above that of elemental composition and provided the resource biochem-

istry needed to begin working out the intermediate steps in numerous physiological processes. This view of reaction pathways through such reactions together with the proposal that these reactions were catalyzed by enzymes provided the guiding assumptions of the newly emerging discipline of biochemistry. For example, one of the best known biochemical pathways, the citric acid or Krebs cycle, consists of successive steps involving oxidations (removal of pairs of hydrogen atoms, which are transferred to NAD^+ or FAD), hydrations and dehydrations (adding or removing H_2O groups), decarboxylations, addition or removal of sulfhydryl-CoA groups, etc. The challenge for biochemists now was to piece together pathways of reactions on molecular groups that would generate the end product from the initial metabolite and to secure evidence for each reaction. (An important aspect of this task, which I am not emphasizing here, was the discovery of models of organization, such as cyclic pathways, that related component operations.)

The Challenge of Identifying Cognitive Operations

Like investigators who tried to characterize physiological processes as fermentations, many investigators in cognitive science have tried to characterize cognitive operations using the same types of idioms as are used to describe the activities cognitive agents perform. For example, differentiating encoding, storage, and retrieval as operations in memory is to conceptualize the internal processes involved in memory in terms of what cognitive agents do. This approach is especially apparent in symbolic or symbol manipulation accounts of cognitive activities. In such accounts, mental operations are viewed as transformations on symbol structures where these symbol structures are construed as being much like sentences in a natural or a formal language. Fodor (1975) quite appropriately characterized symbolic theorists as committed to “a language of thought.” The operations in turn are much like those humans themselves perform when they carry out a task such as writing a manuscript—typing words and phrases, reading them back, altering some, etc. The main difference is that these symbols are thought to be encoded in some way inside a person’s brain and the operations of reading and writing are internal operations, not operations on paper.

In this regard, it is interesting to note that Turing (1936; see also Post, 1936), in proposing the Turing machine as a computational device, was explicitly trying to model human computers—humans whose occupation was to carry out complex mathematical computations. Human computers read and write symbols on a page and apply rules to transform symbols. The finite state device in a Turing machine plays the role of the human and the tape functions as the external memory. When the Turing machine is then invoked by advocates of the symbolic account as the exemplar of the kind of device the mind is taken to be, the external memory and the finite state device are moved inside the head. In this way an activity performed by humans provided the model for operations occurring in their minds. The explanatory

strategy is comparable to that of physiological chemists' invoking fermentations as intermediate processes in alcoholic fermentation. The component operations within the posited mechanism are of the same sort as the behaviors of the mechanism itself.

One of the powerful early tools for constructing symbolic AI models, Newell and Simon's method of protocol analysis (Newell & Simon, 1972), made modeling internal mental operations on agent level behaviors almost inevitable. They required participants to talk aloud as they solved problems, such as the Tower of Hanoi problem, in order to elicit the steps participants employed in solving such problems. These operations then became the building blocks of their computational models. The process did not stop there—the programs were further tested against human performance data. But the overall operations invoked were ones subjects reported performing on the external problem. The production system architecture, which became the foundation for some of the most powerful computational models of human performance (Rosenbloom, Laird, & Newell, 1993; Anderson & Lebiere, 1998) developed out of this perspective. The fundamental idea of this architecture is that just as human agents have a variety of strategies that can be elicited by the problems they are trying to solve (and partial solutions already obtained), their minds are assumed to be equipped with productions that fire when appropriate symbol strings are active in working memory.

The appeal to operations comparable to those performed by human agents is not just characteristic of AI, but also of cognitive psychology. Early cognitive research in psycholinguistics provides an illustrative example. Psychologists extended Chomsky's (1957) proposals for generative grammar, developed initially simply to provide compact accounts of the structure of language itself, to characterize the operations performed when people comprehend or construct sentences. Sentences whose grammatical analysis involved more transformations were hypothesized to require additional mental operations and were found to require more time to process than sentences requiring fewer operations. This evidence was taken to show that the grammatical transformations were also psychologically real (Miller, 1962; see Abrahamsen, 1987). Early research on memory exhibited a very similar character. Sternberg (1966) compared different models of memory search, which all assumed that memory involved the storage of symbolic structures and mentally scanning them, that predicted different patterns of reaction times and argued that the model that fit best characterized actual human mental operations.

It is possible that internal mental operations do have the same character as activities performed by human agents, but if so this is a very unusual case in the history of science. Typically the operations within a mechanism that enables it to perform its behaviors are different in kind from those behaviors. The ability of mechanisms to perform behaviors different from those that their component parts perform is what makes mechanistic explanations so powerful. Organization is the key to achieving this—the operations of the

component parts are orchestrated to work together to accomplish something more than the parts alone can. (Engineering provides an illuminating example: an engineer typically solves a problem by taking existing components and organizing them to function together. For discovering such organization she can win a patent.) Although evolutionary arguments are subject to much abuse, a minimal appeal to evolution enables us to note that distinctive human behaviors largely originate through reorganization of components found in the brains of our close primate relatives. It is operations performed as well in these other species that are organized in novel ways that permits human performance. It seems peculiar to propose that symbol-processing components would have evolved in species that themselves had yet to develop the capacity to manipulate symbols and then became the foundation for our ability to engage in symbol processing behavior.

If not from characterizations of the behavior of humans, where else can investigators draw insights as to the nature of internal mental operations? The prime alternative to which theorists have appealed is the brain. Such was the origin of connectionist approach to cognitive modeling. Neurons are rather explicitly the model for units in a connectionist network and the connections between them are modeled, albeit very loosely, on axons and dendrites (McCulloch & Pitts, 1943; Pitts & McCulloch, 1947; Rosenblatt, 1962). When connectionist modeling was re-energized in the 1980s, the neural plausibility of connectionist networks was one of their touted virtues (Rumelhart & McClelland, 1986; McClelland & Rumelhart, 1986; Smolensky, 1988).

While avoiding the problem of appealing to the activities of cognitive agents for their models of cognitive operations, connectionist models are likely to face the same risk as accounts of physiological processes that appealed to elemental chemical changes. Although it is certainly true that changes in elemental composition of substrates occur in physiological processes, the relevant operations involved higher-level molecular units. Likewise, mental operations involve neurons, but the operations themselves likely involve operations involving parts at a higher level than individual neurons.

The challenge is to identify what sorts of basic operations parts above the level of individual neurons might perform. In studying perceptual processing, systems-level neuroscientists have been able to develop some clues. There the pertinent parts are not individual neurons but brain areas comprised of neural columns. Investigators characterize areas such as V1, V4, and MT as extracting different types of information from the input signal (edges of objects, shape and color, motion) and making it available to areas downstream for further analysis (van Essen & Gallant, 1994; see Bechtel, 2001, for analysis and an account of the history of development of these accounts). Development of such accounts in the case of vision was facilitated by both a fruitful technique (single-cell recording) and the fact that researchers can control the processing occurring in the brain by modulating sensory input. Although single-cell recording actually records from individual neurons, it revealed that

neurons in a particular area all processed similar types of information from different parts of the visual field. As well, within each region there was internal structure: neurons are organized into columns involving layers of connected units that process information from the same part of the visual field and tend to project inhibitory processes to units in other columns. To determine what sort of information a given area extracted researchers could vary the stimuli and correlate inputs with responses. In many respects the kinds of information that visual areas extract are what one might expect from characterizing performance at the behavioral level—people see colors, shape, motion, etc. But the details are often surprising. The shapes detected, for example, are frequently not simple Cartesian shapes but rather more complex forms and the motion registered is not just linear but circular (van Essen & Gallant, 1994).

What has emerged as the dominate approach for linking mental processes with brain activity is functional neuroimaging in which investigators measure blood flow changes as subjects perform tasks. But, as Petersen and Fiez (1993) made very clear, the object in such research is not to localize tasks, although finding increased blood flow in only one or a small number of brain areas as subjects performed tasks in early imaging studies fueled such interpretations. As imaging techniques matured, neuroimaging has begun to identify multiple brain areas characterized as networks engaged in performing the task. But what does each area do? Here neuroimaging confronts the same problem I have been focusing on in this paper—characterizing the component operations. In this regard, neuroimaging is dependent upon progress in cognitive science in developing mechanistic models employing plausible component mental operations.

Biochemistry was fortunate in that structural information about organic molecules provided it with information about higher-level parts on which enzymes operated. Cognitive science and cognitive neuroscience are unlikely to be able to directly use information about the brain in identifying operations, increasing the challenge of discovering the nature of psychological operations. I foresee two strategies that may help guide the discovery of appropriate psychological operations. One is determining, through techniques such as neuroimaging, that the same brain areas are involved in multiple tasks, and then trying to identify what might be common requirements of the different tasks. The other involves using comparative psychology to discover ways in which related species use areas homologous to those in our brains to perform very different tasks. Deacon (1989; 1997) adopts this approach, pointing to evolutionary changes through which control of the vocal apparatus changed in humans from our primate ancestors. Rather than being under the control of limbic system areas, our vocal apparatus came under the control of prefrontal areas that in monkeys are involved in inhibiting previously learned associations and developing new ones. Deacon suggests such inhibition might represent an early form of negation and play an important role in establishing a referential symbolic system in which symbols have internal relations to one another as well

as referential relations to things in the world. This provides a tentative hypothesis as to the operations involved in performing linguistic tasks.

Conclusion

I have argued that an essential part of developing the sort of mechanistic explanations of cognitive phenomena to which cognitive science aspires is identifying the types of operations that parts of the mechanism perform. Appealing to an example from biology, I have illustrated how researchers often appeal either to operations at the same level as the overall mechanism or to operations at too low a level and how progress required discovering operations at the appropriate level of organization. I have argued that cognitive science is confronting the same problem and, moreover, is in the same situation as the early researchers trying to explain physiological processes chemically. There is, however, no simple discovery procedure for the types of operations at a given level of organization in a mechanism. Ultimately, there may be no alternative for cognitive scientists but to employ accounts of operations drawn from what are likely too high or too low a level while awaiting inspired theorizing. If I am right, though, such a theoretical advance is essential if cognitive science is to succeed in the search for mechanisms.

References

- Abrahamsen, A. A. (1987). Bridging boundaries versus breaking boundaries: Psycholinguistics in perspective. *Synthese*, 72(3), 355-388.
- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.
- Bechtel, W. (1994). Levels of description and explanation in cognitive science. *Minds and Machines*, 4, 1-25.
- Bechtel, W. (2001). Decomposing and localizing vision: An exemplar for cognitive neuroscience. In W. Bechtel, P. Mandik, J. Mundale, & R. S. Stufflebeam (Eds.), *Philosophy and the neurosciences: A reader* (pp. 225-249). Oxford: Basil Blackwell.
- Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences*.
- Bechtel, W., & McCauley, R. N. (1999). Heuristic identity theory (or back to the future): The mind-body problem against the background of research strategies in cognitive neuroscience. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the 21st Annual Meeting of the Cognitive Science Society* (pp. 67-72). Mahwah, NJ: Lawrence Erlbaum Associates.
- Bechtel, W., & Richardson, R. C. (1993). *Discovering complexity: Decomposition and localization as scientific research strategies*. Princeton, NJ: Princeton University Press.
- Bernard, C. (1848). De l'origine du sucre dans l'économie animale. *Archives générales de médecine*, 18, 303-319.
- Berthollet, C. L. (1780). Recherches sur la nature des substances animales et sur leurs rapports avec les substances

- végétales. *Mémoires de l'Académie royale des sciences*, 120-125.
- Berzelius, J. J. (1836). Einige KIdeen über bei der Bildung organischer Verbindungen in der lebenden Naturwirksame, aber bisher nicht bemerkte Kraft. *Jahres-Bericht über die Fortschritte der Chemie*, 15, 237-245.
- Brodmann, K. (1909/1994). *Vergleichende Lokalisationslehre der Grosshirnrinde* (L. J. Garvey, Trans.). Leipzig: J. A. Barth.
- Buchner, E. (1897). Alkoholische Gärung ohne Hefezellen (Vorläufige Mittheilung). *Berichte der deutschen chemischen Gesellschaft*, 30, 117-124.
- Chomsky, N. (1957). *Syntactic structures*. The Hague: Mouton.
- Cummins, R. (2000). "How Does It Work?" versus "What Are the Laws?": Two Conceptions of Psychological. In R. Wilson (Ed.), *Explanation and cognition* (pp. 117-144). Cambridge, MA: MIT Press.
- Deacon, T. W. (1989). The neural circuitry underlying primate calls and human language. *Human Evolution*, 4(5), 367-401.
- Deacon, T. W. (1997). *The symbolic species*. New York: Norton.
- Felleman, D. J., & van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1-47.
- Fick, A. E., & Wislicenus, J. (1866). On the origin of muscular power. *Philosophical Magazine & Journal of Science London*, 4th ser., 31, 485-503.
- Fodor, J. A. (1975). *The language of thought*. New York: Crowell.
- Glennan, S. (1996). Mechanisms and the nature of causation. *Erkenntnis*, 44, 50-71.
- Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science*, 69, S342-S353.
- 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.
- Holmes, F. L. (1963). Elementary analysis and the origins of physiological chemistry. *Isis*, 54, 50-81.
- Holmes, F. L. (1992). *Between biology and medicine: The formation of intermediary metabolism*. Berkeley, CA: Office for History of Science and Technology, University of California at Berkeley.
- Lavoisier, A. L. (1781). Mémoire sur la formation de l'acide nommé air fixe ou acide crayeux, que je désignerai désormais sous le nom d'acide du charbon. *Mémoires de l'Académie royale des sciences*, 448-458.
- Lavoisier, A. L. (1789). *Traité élémentaire de chimie, présenté dans un ordre nouveau et d'après les découvertes modernes*. Paris: Cuchet.
- Liebig, J. (1842). *Animal chemistry: or organic chemistry in its application to physiology and pathology*. Cambridge: John Owen.
- 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.
- McClelland, J. L., & Rumelhart, D. E. (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2. Psychological and biological models*. Cambridge, MA: MIT Press.
- McCulloch, W. S., & Pitts, W. H. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 7, 115-133.
- Miller, G. A. (1962). Some psychological studies of grammar. *American Psychologist*, 17, 748-762.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Pasteur, L. (1860). Mémoire sur la fermentation alcoolique. *Annales de Chimie*, 3e Ser, 58, 323-426.
- Petersen, S. E., & Fiez, J. A. (1993). The processing of single words studied with positron emission tomography. *Annual Review of Neuroscience*, 16, 509-530.
- Pitts, W. H., & McCulloch, W. S. (1947). How we know universals: The perception of auditory and visual forms. *Bulletin of Mathematical Biophysics*, 9, 127-147.
- Post, E. L. (1936). Finite combinatorial processes - Formulation I. *Journal of Symbolic Logic*, 1, 103-105.
- Prout, W. (1827). On the ultimate composition of simple alimentary substances; with some preliminary remarks on the analysis of organised bodies in general. *Philosophical Transactions of the Royal Society of London*, 117, 355-388.
- Rosenberg, A. (1985). *The structure of biological science*. Cambridge: Cambridge University Press.
- Rosenberg, A. (1994). *Instrumental biology, or, The disunity of science*. Chicago: University of Chicago Press.
- Rosenblatt, F. (1962). *Principles of neurodynamics; perception and the theory of brain mechanisms*. Washington: Spartan Books.
- Rosenbloom, P. S., Laird, J. E., & Newell, A. (Eds.). (1993). *The Soar papers: research on integrated intelligence*. Cambridge, MA: MIT Press.
- Rumelhart, D. L., & McClelland, J. L. (1986). *Explorations in the microstructure of cognition. Volume 1. Foundations*. Cambridge, MA: Bradford Books, MIT Press.
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11, 1-74.
- Sternberg, S. (1966). High-speed scanning in human memory. *Science*, 153, 652-654.
- Suppe, F. (1977). The search for philosophical understanding of scientific theories. In F. Suppe (Ed.), *The Structure of Scientific Theories* (pp. 3-241). Urbana: University of Illinois Press.
- Turing, A. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society, second series*, 42, 230-265.
- van Essen, D. C., & Gallant, J. L. (1994). Neural mechanisms of form and motion processing in the primate visual system. *Neuron*, 13, 1-10.