60

Science

WILLIAM F. BREWER AND
PUNYASHLOKE MISHRA

The cognitive science of science studies the cognitive processes involved in carrying out science: How do scientists reason? How do scientists develop new theories? How do scientists deal with data that are inconsistent with their theories? How do scientists choose between competing theories? Research on these issues has been carried out by investigators in a number of cognitive science disciplines, particularly psychology, philosophy, and artificial intelligence. More detailed accounts of work in this area can be found in two recent conference volumes (Giere, 1992; Gholson et al., 1989).

At present the psychological and computational study of the scientific process is a relatively underdeveloped topic compared with the study of science by philosophers, historians, and sociologists. Nevertheless, the study of science and scientists within cognitive psychology and cognitive science has led to some important insights and is an exciting topic with enormous potential. Researchers in the cognitive science of science have used a wide range of methodologies. The most common approach has been laboratory studies of undergraduate participants. This type of research allows the use of laboratory control and can be used to develop the background knowledge needed to understand scientists; however, it suffers from not being directly about scientists. Actual laboratory research using scientists has been relatively rare. There have been some retrospective psychological studies of major figures in the history of science. These have the virtue of being about exemplars of exceptional scientific activity, but are limited to material that happened to find its way into the historical record. Recently there have been a few descriptive studies of working scientists. These have the virtue of ecological validity but do not allow for experimental control. There have also been a number of attempts to develop computational models of historical events in science; however, there have been few attempts to model working scientists.

Research in other areas of cognitive science informs and provides a background for the cognitive science of science. (See Articles 7, 20, and 21: CONCEPTUAL CHANGE, PROBLEM SOLVING, and REASONING.) The cognitive science perspective of this chapter means that we will not cover a number of topics that would be covered in a more inclusive discussion of the psychology of science (e.g., the personality and motivation of scientists, the social interactions of scientists, the institutions of science). We organize this chapter in terms of a simple heuristic: What do scientists do everyday in their capacity as scientists and what psychological processes are involved in those activities?

We propose that the activities of scientists can be grouped into three general categories:

1 Understanding and evaluating scientific information. Scientists spend much time reading the scientific literature and attending conferences. These activities suggest an investigation of how scientists understand and evaluate scientific data and theories.
2 Generating new scientific knowledge. Scientists design and carry out experiments and develop new theories. These activities suggest the study of scientific research strategies and the process of scientific discovery.
3 Disseminating scientific knowledge. Scientists invest considerable time writing up and giving talks about the results of their work. These activities suggest the study of the scientific writing process and the more general process of disseminating scientific information.

The above organizing heuristic provides an interesting perspective on the current literature on the cognitive psychology and cognitive science of science. Most research in this area has been directed at the issue of generating new scientific knowledge, with a particular focus on two subtopics: scientific discovery and data-gathering strategies in science. Our organizing heuristic suggests a wider range of topics for investigation.

In thinking about the topics to be covered in this chapter, we have adopted a second complementary heuristic: the focus of work in this area should be on issues relatively specific to scientists. For example, in the area of understanding scientific information, the study of the general processes involved in understanding difficult expository prose is not unique to scientists but could also be studied with other groups, such as historians or lawyers. However, understanding the role of data in theory evaluation, while not unique, seems much more specific to scientific thinking.

One final issue raised by our approach is the representativeness of the scientists discussed. Much of the thinking in the psychology of science has been driven by a small set of classic anecdotes about famous scientists: for example, the story of how Einstein was led to the development of the theory of special relativity by imagining himself traveling alongside a beam of light, or the story of how Kekulé discovered the ring structure of benzene by visualizing a snake grabbing its tail. These certainly are an interesting source of ideas; however, it seems to us that they ought to be used against a background of knowledge about how ordinary scientists work. We prefer the strategy used in areas such as the study of human memory, where individuals with exceptional memory are investigated to see how their memory performance relates to what is known about ordinary memory performance. Thus we think that the activities of exceptional scientists must be compared and contrasted with those of ordinary scientists. In the remainder of the chapter we will examine research on a number of selected topics in the cognitive science of science

Understanding and evaluating scientific information

When scientists read the work of other scientists or try to make sense of their own data, one important question that emerges is how the data are interpreted in the context of the theoretical beliefs of the scientist. One aspect of this issue is how scientists respond to anomalous data (data that go against their theories). Many popular and positivist accounts of science give primacy to data and assume that scientists rapidly change theories in the face of anomalous data. Other philosophers of science (Thomas Kuhn, Paul Feyerabend) focus on the powerful role of theory and argue that scientists often do not change their theories in the face of anomalous data.

Chinn and Brewer carried out a program of research on this issue (Brewer and Chinn, 1994). They used evidence from experimental psychology and from the history of science to argue that there are seven fundamental responses that scientists make to anomalous data. Scientists can *ignore* the data. They can *reject* the data (e.g., on methodological grounds). They can *exclude* the data from the domain of their theory by arguing that the theory is not intended to explain these particular data. They can hold the data in *abeyance* (i.e., concede that their theory cannot explain the data at present but that it will be able to in the future). They can accept the data but *reinterpret* it so that it is consistent with their theory. They can accept the data and make minor *peripheral changes* to their theory. And finally, they can accept the data and *change* their theory. Chinn and Brewer give a psychological analysis of these responses to anomalous data in terms of whether the scientist believes the data, whether the data can be explained, and whether the original theory is changed. They use this analysis to make the strong claim that these seven forms of response exhaust the psychologically plausible responses to anomalous data.

Chinn and Brewer also investigated how undergraduates respond to anomalous data. In these experiments, students read about a particular scientific theory (e.g., the meteor theory of mass extinctions in the Cretaceous period) and were then given a piece of anomalous data. Across a wide range of theories and experimental conditions, the students showed a strong tendency to hold on to their theories in the face of anomalous data by using one or more of the first five approaches listed above. Clearly these experimental studies need to be extended to examine the responses of scientists.

Thomas Kuhn's classic work on the nature of scientific revolutions has highlighted the issue of radical theory change in science. How can we understand the psychological processes involved when a scientist comes to reject one theory and replace it with another that is fundamentally different? Paul Thagard (1992) has recently developed a computational approach to these issues. He has formulated a theory of explanatory coherence which attempts to capture the relationships among concepts in a theory and the relationship between the theory and relevant data. He has implemented this theory in a computer program (ECHO) which represents theories in terms of nodes (concepts) and relations and uses a connectionist approach to calculate the degree of explanatory coherence for a theory.

Thagard has applied ECHO to a number of historical cases of scientific revolutions (e.g., Lavoisier, Darwin). He takes a cognitive approach to these historical cases and argues that a scientist choosing between two theories will choose the one with higher explanatory coherence. When presented with the historical facts in these cases of scientific revolutions, ECHO finds that the historically successful theory shows higher explanatory coherence. This is important work. However, if it is to serve as an adequate model of the psychology of the scientist, it needs further development. For example, Thagard's system currently compares two completed theories, whereas the historical record suggests that shifting from one theory to another is a slow, gradual process for the scientists involved. The system does not incorporate the psychologically important property of the theory-ladenness of data (i.e., that a scientist's theory influences his or her evaluation of the data). Finally, the system does not provide an account of what constitutes a psychologically appealing explanation for a particular datum. Although ECHO captures several important aspects of theory selection, clearly much more work is needed to provide a comprehensive model for theory choice.

Several investigators have attempted to take a cognitive science approach to the study of the work of a single famous scientist. Nancy Nersessian has studied Maxwell; Howard Gruber has investigated Darwin; and Ryan Tweney has examined Faraday. These accounts tend to be fairly detailed and complex. To give one brief example from this type of work, consider Tweney's (1992) discussion of Faraday's laboratory procedures. Faraday was an extremely prolific experimenter, and Tweney pointed out that Faraday kept extensive notebooks and developed a number of elaborate indexing schemes as memory aids, so that he could retrieve and understand the results of his own experiments.

Generating new scientific knowledge

Within the area of the cognitive study of science, the topic of research strategies for gathering data has received by far the most attention. Essentially all of this work derives from an experiment carried out by Peter Wason in 1960 (see Newstead and Evans, 1995, for reviews of this work). In Wason's experiment, undergraduate participants were told that the experimenter had a rule in mind involving three numbers; the participants were to generate number triples, and after each triple the experimenter would tell the participant if the number was correct or not. The experimenter then told the participants that the triple 2-4-6 was a correct instance of the rule. When the participants thought they knew the rule, they were to tell the experimenter, who would then inform them if they were right or not, and, if they were not right, the experiment continued. The rule used in the original and in most following experiments was: *numbers in increasing order of magnitude.* Wason intended this task to be a simulation of the scientific process of data gathering: the experimenter's rule corresponds to Nature, the undergraduates function as scientists, the generated triples correspond to designed experiments, and the feedback provided by the experimenter corresponds to data generated by the experiment.

Many of the participants in Wason's experiment never found the correct rule and instead focused on rules that were subsets of the correct rule (e.g., even numbers increasing by two). Following the work of the philosopher Karl Popper, Wason adopted the position that falsification is the normatively correct research strategy for data gathering. He argued that the students in this experiment tended to generate data consistent with their hypotheses and thus were using an unscientific research strategy.

Wason's experiment stands as a landmark in the cognitive study of science. It opened up the laboratory study of scientific research strategies. However, in retrospect, it is clear that in many important ways it was an inadequate experimental analog of scientific data gathering. For example, the *theories* in this task are analogs to scientific laws, not explanatory theories. The data never contain errors. The participants can directly ask *Nature* if their theory is correct. In this task (unlike real science), theories are easy to develop, and it is easy for participants to shift to a new theory. The correct theory, in combination with the initial seed (2-4-6), led the participants to develop theories that were subsets of the correct theory and so gave rise to a situation in which erroneous theories were highly confirmed by the data. Finally, it is not clear to us that the strong focus on disconfirmation in this literature makes sense when applied to the actual process of doing science. For example, in most real examples of science (e.g., the meteor theory of mass extinctions in the Cretaceous) the experiments that are carried out are derived from the theory (e.g., having found an iridium layer at a site in Italy, see if you can find an iridium layer in Denmark, or see if there was a large impact site that occurred at the beginning of the Cretaceous period) and have the potential to either support or disconfirm the theory. A number of these problems have been noted (and sometimes corrected) by other researchers in this area (Michael Gorman, Joshua Klayman, Michael Mahoney, Ryan Tweney).

Several investigators (Clifford Mynatt, David Klahr, Kevin Dunbar) have studied research strategies with task environments somewhat richer than the 2-4-6 task; however, it appears to us that additional work will be required to develop laboratory tasks that do a convincing job of capturing the research strategies which scientists actually use to gather data.

Quite a different approach to the study of scientific discovery has been taken by Langley, Simon, Bradshaw, and Zytkow (1987). These researchers developed computer simulations of various aspects of the process of scientific discovery. One of the most widely discussed programs that they developed is BACON, a program that carries out data-driven discovery. This program has been given data from a variety of historical cases (e.g., Kepler's third law, Black's law of temperature equilibrium) and has generated the appropriate mathematical functions.

BACON certainly provides a demonstration that it is possible to model part of the discovery process; however, one must note the limitations of this program. First, the discoveries made by the program are examples of scientific laws, not explanatory scientific theories such as the atomic theory or the theory of plate tectonics. Even with the discovery of laws, it would appear that in the historical cases the situation from the point of view of the scientist was less well structured than that given to the computer. In addition, the program has no motivated way to deal with experimental error. The occurrence of a single very erroneous data point would cause severe problems in BACON's ability to discover a relationship. Finally, it would appear that the background theoretical beliefs of scientists play a major role in most cases of scientific discovery, and this is missing from BACON. Langley and colleagues are aware of some of the limitations just discussed and have developed other programs (e.g., STAHL) which attempt to capture additional components of scientific discovery. Even though BACON and the other programs model only parts of the discovery process, they are still a very impressive achievement, since they demystify scientific discovery and show that at least some aspects of the process can be modeled by a computer.

Kevin Dunbar (1995) has recently carried out an observational study of scientists in several biological laboratories. He attempted to gather data from all the public aspects of research going on in these labs (including planning and execution of experiments, evaluation of results, group meetings, and public talks). This study provides a rich set of data on actual scientific practice. Among other things he has studied is the issue of how analogies are used in the scientific process. He finds that analogies to closely related problems were widely used to map the characteristics of previously successful experiments onto experiments with problems, in order to design new experiments.

Analogies to problems in different domains of inquiry were rarely used to drive scientific discovery, but were occasionally used to help someone understand an issue under discussion (see Article 1, ANALOGY).

Disseminating scientific knowledge

Scientists spend a significant amount of their time communicating their ideas (through journal articles, conference talks, e-mail, personal conversation, etc.). There has been relatively little research on this aspect of the scientific process. However, Greg Myers (1990) has studied the practices of working biologists engaged in different writing activities. He points out that the scientific journal article does not mirror the actual events in the laboratory that led to the article. Instead, the article is structured to give a series of logical arguments to support a theoretical position. He also notes that if the same scientist writes a popular article (e.g., for *Scientific American*) about the same experiments, the results take on, yet again, a very different textual form. In this case the findings tend to be written up as fact, without qualifiers and discussions of alternate theories. Clearly there is much additional work to be done in exploring the cognitive aspects of the rhetorical practices which scientists use in disseminating knowledge.

The research in the cognitive science of science reviewed here has already provided some interesting insights into the cognitive process underlying the activities of scientists. However, given the important role of science in modern society, it seems that the cognitive science of science has not received adequate attention. The task of understanding the cognitive processes involved in doing science is difficult, but the multiple methodologies used in cognitive science are particularly appropriate to this problem. This is an area with great potential.

References and recommended reading

Brewer, W. F. and Chinn, C. A. 1994: Scientists' responses to anomalous data: evidence from psychology, history, and philosophy of science. *PSA 1994*, 1, 304–13.

Dunbar, K. 1995: How scientists really reason: scientific reasoning in real-world laboratories. In R. J. Sternberg and J. D. Davidson (eds), *The Nature of Insight*. Cambridge, Mass.: MIT Press, 365–95.

*Gholson, B., Shadish, W. R. Jr., Neimeyer, R. A. and Houts, A. C. (eds) 1989: *Psychology of Science: Contributions to Metascience*. Cambridge: Cambridge University Press.

*Giere, R. N. (ed.) 1992: *Minnesota Studies in the Philosophy of Science*, vol. 15: *Cognitive Models of Science*. Minneapolis: University of Minnesota Press.

Langley, P., Simon, H. A., Bradshaw, G. L. and Zytkow, J. M. 1987: *Scientific Discovery: Computational Explorations of the Creative Process*. Cambridge, Mass.: MIT Press.

Myers, G. 1990: *Writing Biology: Texts in the Social Construction of Scientific Knowledge*. Madison: University of Wisconsin Press.

Newstead, S. E. and Evans, J. St B. T. (eds) 1995: *Perspectives on Thinking and Reasoning: Essays in Honour of Peter Wason*. Hove: Erlbaum.

Thagard, P. 1992: *Conceptual Revolutions*. Princeton: Princeton University Press.

Tweney, R. D. 1992: Serial and parallel processing in scientific discovery. In R. N. Giere (ed.), *Minnesota Studies in the Philosophy of Science*, vol. 15: *Cognitive Models of Science*, Minneapolis: University of Minnesota Press, 77–88.

Blackwell
Companions to
Philosophy

A Companion to Cognitive Science

Edited by

WILLIAM BECHTEL and
GEORGE GRAHAM

Advisory editors

David A. Balota
Paul G. Chapin
Michael J. Friedlander
Janet L. Kolodner

1998

BLACKWELL
Publishers