How Scientists Think in the Real World: Implications for Science Education

Kevin Dunbar
McGill University

Research on scientific thinking and science education is often based on introspections about what science is, interviews with scientists, prescriptive accounts of science, and historical data. Although each of these approaches is valuable, each lacks some key components of what scientists do, which makes it difficult to determine what scientists are being trained for and what essential thinking and reasoning tools they must have. The research reported herein sought to determine the cognitive processes underlying reasoning in science using two approaches. The first is to bring participants into the laboratory and give them scientific problems to work on. The second is to investigate real scientists as they work at their own problems. Both approaches make it possible to propose several thinking and reasoning strategies that are conducive to making discoveries. Both also make it possible to understand some of the basic cognitive mechanisms underlying scientific thinking.

How do scientists think, reason, and conduct their research? Are some thinking and reasoning strategies more conducive to making discoveries than other strategies? How do scientists represent their knowledge? Answers to these questions are of paramount importance not only in understanding what science is, but also in shaping the future of science education. Unfortunately, we know little about the basic processes that are involved in current-day scientific thinking. The goal of this research was to look at what scientists actually do in their research, which types of thinking and reasoning strategies they use, and how they change their knowledge. Over the past decade, I have been investigating scientific thinking in my laboratory and studying scientists in their own labs as they reason about their research. Eight molecular biology and immunology laboratories at two major universities have been investigated. Weekly lab meetings are one place where much reasoning occurs and new discoveries are made. We have performed extensive cognitive analyses of

---

Direct all correspondence to: Kevin Dunbar, Department of Psychology, McGill University, Stewart Biological Science Building, 1205 Docteur Penfield Avenue, Montreal, Quebec, Canada H3A 1B1 <dunbar@psych.mcgill.ca>.

Journal of Applied Developmental Psychology 21(1): 49–58
ISSN: 0193-3973

these meetings and have identified some of the key components of contemporary scientific thinking that are important in generating new concepts, modifying old concepts, and solving difficult problems. This paper outlines three important strategies that scientists use: attention to unexpected findings, analogic reasoning, and distributed reasoning.

Research on scientific thinking and science education is often based on introspections about what science is, interviews with scientists, prescriptive accounts of science, and historical data. Although each of these approaches is valuable, each lacks some key components of what scientists do, which makes it difficult to determine what scientists are being trained for and what essential thinking and reasoning tools they must have. My research has sought to determine the cognitive processes underlying reasoning in science using two approaches. The first is to bring participants into the laboratory and give them scientific problems to work on. The second is to investigate real scientists as they work at their own problems. Both approaches make it possible to propose several thinking and reasoning strategies that are conducive to making discoveries. Both also make it possible to understand some of the basic cognitive mechanisms underlying scientific thinking.

Much research on scientific thinking involves bringing participants into the cognitive laboratory and asking them to solve scientific problems (see Klahr, 1994; Tweney, Doherty, & Mynatt, 1982). In my research, I have used a number of different tasks to investigate the cognitive processes underlying scientific thinking, tasks that range from determining how a device works (Klahr & Dunbar, 1988; Klahr, Fay, & Dunbar, 1993) to determining the functions of genes in a simulated molecular genetics laboratory (Dunbar, 1993; Dunbar & Schunn, 1990; Schunn & Dunbar 1996). I have built a molecular genetics laboratory where participants can conduct simulated molecular genetics experiments on a computer (Dunbar, 1993) that are similar to the experiments molecular biologists conduct in their laboratories. After participants are taught about molecular genetics, they are given the task of discovering how sets of genes control other genes.

We have found that participants use a variety of scientific reasoning strategies, some of which lead to discoveries and others to failure. Participants who succeed in making discoveries pursue different goals from those who do not discover new concepts (Dunbar, 1993); in particular, they set a new goal of trying to explain surprising results, and they try to induce hypotheses from experimental results rather than testing hypotheses (Klahr & Dunbar, 1988). Participants also use strategies that block scientific discoveries. Those who fail to make a discovery set a goal of finding evidence consistent with their current knowledge and distort the interpretation of results to fit their current hypothesis. This "confirmation bias" is the result of setting a goal to obtain a particular finding, rather than an inability to think of alternate hypotheses. When the task is changed so that the goal of the participant is achieved, then the participant sets a new goal of accounting for surprising findings and makes the discovery (Dunbar, 1993). I have also used this task to investigate whether participants are aware of the origins of their hypothesis; Schunn and Dunbar (1996) found that participants can be "primed" to discover a new scientific concept, yet have no awareness of where their hypothesis came from. Clearly, this finding casts further doubt on scientists' retrospective accounts of how they made a discovery.

Although research on individuals has produced many rich and important theories of reasoning in general and some components of scientific reasoning in particular (see Klahr, 1994, for a recent review of this literature), there are several distinct problems with making generalizations about how scientists reason from experiments on individuals. First, science takes place in a social context; groups of scientists work on a problem in a laboratory, rather than one scientist working alone. So far, cognitive scientists have tended to investigate scientific reasoning in individuals and have ignored the social context of science. This issue is particularly important when it comes to investigating aspects of science that are important for science education. If scientific thinking in real-world science takes place in groups, we must know what types of reasoning heuristics scientists should be taught. Second, psychologists have used tasks, such as discovering an arbitrary rule, that are not “real” scientific problems (e.g., Klayman & Ha, 1987). Third, the participants that psychologists use are generally not scientists (e.g., Klahr & Dunbar, 1988). Clearly, scientists working on real scientific problems must be studied as well. Unfortunately, prior studies of scientists have given them the same simple and arbitrary concepts as nonscientists. Fourth, participants in psychology lab experiments work on problems that may last for as little as 10 minutes and involve no extensive knowledge of a scientific topic (e.g., Klayman & Ha, 1987). In scientific research, a particular problem may take months, years, or decades to solve, and the scientists have extensive knowledge of a domain.

Conducting interviews with scientists has been one alternative approach for learning more about what scientists do. A number of recent artificial intelligence models have been based on such interviews. Although interviews are a very important source of data, retrospective reports gleaned from interviews are notoriously unreliable (see Ericsson & Simon, 1993). Furthermore, research from my laboratory has shown that participants are often unaware of what leads them to form a new concept and what strategies they used (Dunbar & Schunn, 1990; Schunn & Dunbar, 1996). We found that solving one problem improved performance on an analogically similar problem, yet in retrospective interviews, the participants did not report using any information from the first problem to solve the second problem. Thus to uncover the strategies that scientists use, interviews cannot be wholly relied on. Historical approaches suffer from similar problems. Crucially, both approaches have the limitation of their data being based on *indirect* and selective access to the cognitive processes underlying scientists’ social and cognitive activities.

A DIRECT APPROACH TO UNDERSTANDING SCIENCE

The above limitations with previous research have led me to adopt a new approach to uncover the thinking and reasoning strategies involved in science: investigating real scientists working on real research. This approach entailed spending an extensive period of time in real scientific laboratories. I have now collected data from eight molecular biology and immunology laboratories at two major universities in North America. I followed all aspects of particular scientific experimental results, lab meetings, planning of further experiments, public talks, and writing of journal articles. Some of the research projects resulted in important scientific discoveries, and some did not. This provides a totally novel database with which to address fundamental questions concerning the cognitive processes involved in scientific discovery. The data have several important implications for science education. We have been able to identify certain strategies that are likely to lead to success and others that are unlikely to lead to success. Three strategies are discussed:

(a) Scientists focus on unexpected findings as a source of new experiments and theories.
(b) Analogic reasoning is an important component of generating hypotheses, designing experiments, and interpreting data.
(c) Scientists frequently engage in distributed reasoning when they encounter a problem in their research. A more extended account of the findings can be found in Dunbar (1995, 1997).

UNEXPECTED FINDINGS

Scientists often talk about the chance result, the unexpected event that leads to a discovery. Yet whether and how scientists make use of unexpected findings is still not known. Furthermore, many highly publicized events in science have had as their focal point a scientist ignoring an unexpected finding. Thus some commentators on science have argued that scientists will ignore unexpected findings, particularly if the findings are inconsistent with their theory. Other researchers have argued that a useful strategy in science is to focus on unexpected findings. According to this view, scientists work with a heuristic assumption such as: *If the finding is unexpected, then set a goal of discovering the causes of the unexpected finding* (Dunbar, 1993; Kulkarni & Simon, 1988). The ways that scientists deal with unexpected findings are potentially very important. To understand further the role of unexpected findings in contemporary science, we investigated whether unexpected findings are frequent and whether scientists do indeed pay attention to them (Dunbar, 1995, 1997, 1999).

We have been conducting extensive analyses of every finding made in the laboratories studied and what happened to them. Findings were coded as unexpected if the outcome of the experiment was different from that predicted by the scientist. We then coded what the scientists thought was the cause. We have found that unexpected findings are very frequent. For example, in just one laboratory over a 9-month period, 281 results were obtained, with 48% of the findings being expected and 52% unexpected. In our analyses of unexpected findings in other laboratories, we have also found large incidences of unexpected findings, ranging from 40% to 60%.

When an unexpected finding occurs, what do scientists do? We have categorized reactions along several dimensions and have found that scientists attribute approximately 54% of unexpected findings to an unknown origin, 30% to methodological problems, 8% to a mistake, and 8% to a new mechanism. One striking thing about these results is that all the unexpected findings received considerable attention. Furthermore, any findings that were inconsistent with the scientists' current hypothesis received more attention than findings that were merely unexpected.

We can now turn to what scientists decide to do after an unexpected result. On 37% of occasions scientists decided to replicate the experiment. One interesting point is that even replications have changes in the method, albeit minor changes. Another strategy that scientists adopted was to analyze the results of an unexpected finding further. For 38% of the unexpected findings, the scientists took an unexpected result, such as the presence of a particular protein, and conducted more detailed analyses of the protein. For 9% of the unexpected findings, the scientists adopted a new protocol. This was usually the case when the current protocol was not giving scientists the result that they needed. For 17% of the unexpected findings, the scientists decided to abandon the current protocol, either because another protocol worked better or because it just did not seem possible to obtain the result they needed.

Now we focus on the psychological processes that the scientists use when they are dealing with unexpected findings. We coded the inductions, deductions, and causal reasoning statements that the scientists make when they reason about unexpected findings. We found that more than 70% of scientists' inductions, 50% of their deductions, and 70% of their causal reasoning statements are devoted to unexpected findings. Thus unexpected findings are a primary focus of scientists' model building and causal reasoning. The scientists take an unexpected finding and attempt to build a causal model of the events underlying it. As soon as the causal model is built, the scientists can decide what step to take next. Building a causal model can involve deductions, inductions, and the addition of underlying mechanisms to produce a new model.

Overall, our results concerning unexpected findings bring to the fore several often ignored issues in science education. Most students are not explicitly taught how to deal with unexpected findings. Given the high frequency of such findings, scientists must learn to deal with them. Furthermore, unexpected findings can be the source of new hypotheses and major theoretical insights. Thus one director of another lab claimed that his *modus operandi* was to focus on unexpected findings. In his laboratory, this focus on unexpected findings that are inconsistent with current knowledge in molecular biology has led to the discovery of numerous new genes and novel biologic mechanisms (see Dunbar, 1997, for a discussion of discovery based on an unexpected finding). All of the senior scientists paid keen attention to unexpected findings, particularly those relating to control conditions. These results indicate that novice scientists must be taught to expect the unexpected and that various methodological and psychological strategies can be used when unexpected findings occur.

Analogical Reasoning

Scientists lead a complex existence. They must formulate theories, design experiments, and interpret results. One problem for scientists is that these processes are not formulaic. Scientists must go beyond what they know in all aspects of their science. For example, theory construction is often a mix of drawing inferences from old knowledge and engaging in causal reasoning, induction, and deduction. Even with a combination of all these mental processes, formulation of new theories or design of experiments can be inexorably slow. Thus although scientists frequently do build new theories or design experiments in a stepwise fashion (e.g., with an induction, followed by a causal mechanism, and then a deduction), they can use another strategy: making analogies. Scientists can use analogy to perform many steps all at once.

We found that scientists frequently use analogy when there is not a straightforward answer to their current problem. Their problem may be in formulating a theory, designing an experiment, interpreting data, or going from theory to an experiment. We have found that rather than trying various permutations on a question, the scientists search for a similar problem that has been solved and seek to import its answer to their current problem. If scientific reasoning is viewed as a search in a problem space, then analogy allows the scientist to leap to different parts of the space rather than slowly searching through it until they find a solution (see Dunbar, in press). Most views of science have treated scientific thinking and reasoning as a slow, iterative process. But we have observed that scientists frequently use analogy to solve a problem quickly. In this section, some of the mechanisms that scientists use when they make analogies are discussed, showing that certain types of analogies are important in designing experiments, whereas other are important in explaining concepts to other scientists. These different uses of analogies have several implications for the training of scientists.

Scientists have frequently referred to analogy as a source of new hypotheses. Thus Rutherford's analogy of the structure of the atom has been frequently mentioned as an example of the use of analogy in science. Many scientists and commentators on science have proposed that one way scientists generate new models and new concepts is to draw an analogy to a very different domain (see Boden, 1993). We investigated the role of analogy in science by coding the use of analogy in four laboratories (see Dunbar, 1997, for a more complete account of our work on analogy). We wanted to find out how often analogies were used, what analogies are used, and for what purposes they are used. We coded each use of an analogy at 16 meetings, citing any instance where a scientist used knowledge from a domain to fill in gaps in their knowledge of the issue they were investigating or used their prior knowledge to help others understand an issue.

Analogy is frequently used. There were 99 analogies made at the 16 meetings. Based on our coding, we found that the goals of the analogies could be grouped into classes: formulate a hypothesis, design an experiment, fix an experiment, or explain a result. We found that the type of analogy made was directly related to the goal that the scientist had. When the scientists were designing and fixing experiments, the analogies were made to very similar experiments or the same organisms (e.g., HIV to HIV). However, when the goal was to formulate hypotheses, the scientists tended to make analogies to other organisms that have been investigated (e.g., from Ebola virus to HIV). When the goal was to explain a concept to other members of a lab or a more general audience, the scientists made analogies to a very distant domain (e.g., Gallo's analogy of the HIV virus to a pearl necklace). These results were very surprising; in particular, we did not expect to see so many analogies made to other organisms. Furthermore, despite the anecdotal accounts of distant analogies having a major role in science, our findings indicate that distant analogies are primarily used to explain concepts to others rather than to formulate novel hypotheses and experiments.

As a final note on analogy, we found that scientists have little memory for the analogies that they use. When we go back and ask the scientists to remember how they generated a new concept or solved a problem at a meeting that we recorded, they have little memory of how it occurred, particularly with regard to analogies. It appears that the scientists remember the results of their reasoning rather than the small steps that they made or the different analogies that they used to make a discovery. As a result, scientists tend to undervalue the role of analogies from the same domains or to other related organisms.

Overall, our analyses suggest that analogy is a very powerful way of filling in gaps in current knowledge and suggesting experimental strategies that scientists should use. The key to successful analogic reasoning is the ability of the scientists to abstract the crucial features of the current problem and search for other problems that have been solved. The use of analogy tends to correlate with experience. Graduate students make few analogies, whereas postdoctoral fellows and professors make many. It may be the case that the graduate students do not have the requisite structural knowledge available to make analogies, or more likely, graduate students have not been educated in using analogies to solve scientific problems. It may be the case that students learn to use analogy by osmosis, and that explicit encouragement of the use of analogy would be a more effective teaching strategy.

Distributed Reasoning

One of the most significant changes in the structure of science over the past century has been a switch from science as an individual process, where one scientist conducts all the steps in a scientific project, to science as a group enterprise, where the reasoning and knowledge are distributed over many scientists. Surprisingly, little is known about the nature of the thinking and reasoning that groups of scientists use. More importantly, it is not known which aspects of scientific thinking and reasoning in groups are conducive to making discoveries and which are not. We have analyzed the distributed reasoning that occurs in several different laboratories and can now sketch out some of its basic features (see Dama & Dunbar, 1996; Dunbar, 1997).

We define distributed reasoning as any instance of reasoning where more than one person is involved in the reasoning. For example, three scientists who are analyzing data may draw three different inductions from the same data, or one scientist may provide one premise to the induction, another a second premise, and a third the conclusion. We have analyzed the distributed reasoning that occurs in five laboratories and have found that more than 50% of the reasoning that takes place at a meeting is distributed, indicating that distributed reasoning is an important component of contemporary science.

Does distributed reasoning occur for certain types of information and not others? Distributed reasoning is equally likely to occur when discussing theory, method, or results. However, our analyses indicate that most distributed reasoning is future-oriented, focusing on further developments in theory, planning of future experiments, and interpretations of future data.

An important feature of distributed reasoning is that it allows scientist to generate different representations of a problem. This can be seen clearly when we look at distributed reasoning for inductions and deductions. We have analyzed every induction and deduction made at eight meetings in one laboratory. We found that of 74 inductions, 14 involved contributions from more than one person. Interestingly, the 14 inductions resulted in 23 different conclusions. When the conclusions were analyzed, we found that many of the scientists propose alternate inductions from the data, and that much of the ensuing discussion consists of evaluating the rival inductions. This pattern is even more apparent in deductive reasoning. There were 189 deductions made; 135 deductions by one person and 54 by other members of the lab. The 54 distributed deductions resulted in 121 different conclusions! Thus given the same premises, scientists can reach radically different conclusions. The reason for this is not that scientists are illogical, but that they interpret the premises in very different manners. The different conclusions are a symptom of the differences in underlying representations of the problem that they are working on. As with induction, much of the ensuing discussion turns to evaluating different ways that the problem should be addressed. Thus distributed reasoning results in many different ways of representing a problem and can result in radically new theoretical and methodological breakthroughs.

One important question that arises is whether distributed reasoning is always beneficial. We have identified two situations where distributed reasoning is ineffective. The first is when all the members of a laboratory come from a very similar background. In this situation, problem solving by the group is no better than that by an individual. As a good case in point, one laboratory did not achieve the goal of understanding how a particular gene works. All the members of the lab came from very similar research backgrounds; they had all worked with the same organism. Furthermore, their use of analogy was equally constrained: they always made analogies to the organism that they had worked with in the past. As a consequence, the scientists were not able to import new knowledge into their research and overcome their research problems.

Distributed reasoning is also ineffective when the members of the group are very diverse and have different and competing goals. We have analyzed laboratory meetings where four laboratories have come together to discuss research projects. In this situation we have often found little agreement on many issues, and the meetings do not help to solve the scientists' problems. Another more familiar case where this happened is the Challenger disaster. In this case, there were many groups with different goals and different representations of the problems with the space shuttle, resulting in a disastrous decision.

Conclusions

This paper presents an overview of a new way of uncovering the essential thinking and reasoning components of current-day science. This method of going into the real world and analyzing the ways that scientists think and reason allows us to go beyond existing historical and anecdotal accounts. Importantly, this approach makes it possible to uncover the mental processes that scientists engage in while they conduct their research. Through this, we can identify which strategies lead to success and which do not. Although this research is at its early stages, it is now possible to make some generalizations relevant to science education, stated below as a set of rules.

(a) Follow up on surprising results. Pay particular attention to unexpected findings in the control conditions, because these results reflect either a conceptual problem or an experimental problem with your research.
(b) Engage in analogic reasoning in both formulating hypotheses and solving research problems. Use distant analogies as an explanatory device. Educators should make frequent use of distant analogies to explain things.
(c) Set up situations where the students or scientists can engage in distributed reasoning. This is particularly important in planning new research and understanding findings.
(d) Pay attention to your goals. Make sure that your current goal is not blocking you from considering alternate theories or ways of conducting experiments.

Acknowledgments: The research reported in this paper was funded by the Spencer Foundation, the Canadian National Science and Research Council, and the United States Office of Naval Research.

REFERENCES

Boden, M. (1993). *The creative mind: Myths and mechanisms.* New York: Basic Books.
Dama, M., & Dunbar, K. (1996). Distributed reasoning: When social and cognitive worlds fuse. In *Proceedings of the Eighteenth Annual Meeting of the Cognitive Science Society* (pp. 166–170). Mahwah, NJ: Erlbaum.
Dunbar, K. (1993). Concept discovery in a scientific domain. *Cognitive Science, 17*, 397–434.
Dunbar, K. (1995). How scientists really reason: Scientific reasoning in real-world laboratories. In R. J. Sternberg & J. Davidson (Eds.), *The nature of insight* (pp. 365–395). Cambridge, MA: MIT Press.
Dunbar, K. (1997). How scientists think: Online creativity and conceptual change in science. In T. B. Ward, S. M. Smith, & S. Vaid (Eds.), *Conceptual structures and processes: Emergence, discovery, and change* (pp. 461–493). Washington, DC: APA Press.
Dunbar, K. (1998). Problem solving. In W. Bechtel & T. Grahman (Eds.), *A companion to cognitive science.* New York: Blackwell.
Dunbar, K. (1999). How scientists build models: In vivo science as a window on the scientific mind. In C. Magnani, N. J. Nersessian, and P. Thagard (Eds.), *Model-based reasoning in scientific discovery* (pp. 85–99). New York: Kluwer Academic/Plenum.
Dunbar, K., & Schunn, C. D. (1990). The temporal nature of scientific discovery: The roles of priming and analogy. *Proceedings of the 12th Annual Meeting of the Cognitive Science Society* (pp. 90–100). Hillsdale, NJ: Erlbaum.
Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data* (Rev. ed.). Cambridge, MA: MIT Press.
Klahr, D. (1994). Searching for cognition in cognitive models of science. *Psycoloquy, 5*, 68.
Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science, 12*, 1–48.

Klahr, D., Fay, A. L., & Dunbar, K. (1993). Heuristics for scientific experimentation: A developmental study. *Cognitive Psychology, 25*, 111–146.

Kulkarni, D., & Simon, H. A. (1988). The processes of scientific discovery: The strategy of experimentation, *Cognitive Science, 12*, 39–176.

Schunn, C. D., & Dunbar, K. (1996). Priming, analogy, and awareness in complex reasoning. *Memory and Cognition, 24*, 271–284.

Tweney, R. D., Doherty, M. E., & Mynatt, C. R. (Eds.). (1982). *On scientific thinking.* New York: Columbia University Press.