17

How Scientists Think: On-Line Creativity and Conceptual Change in Science

Kevin Dunbar

This chapter reports an investigation of "On-line Creativity." I present a new account of the cognitive and social mechanisms underlying complex thinking of creative scientists as they work on significant problems in contemporary science. I lay out an innovative methodology that I have developed for investigating creative and complex thinking in a real-world context. Using this method, I have discovered that there are a number of strategies that are used in contemporary science that increase scientists' likelihood of making discoveries. The findings reported in this chapter provide new insights into complex scientific thinking and will dispel many of the myths surrounding the generation of new concepts and scientific discoveries.

This research was made possible by a grant from the Spencer Foundation. The research was also supported by Grant OGP0037356 from the National Sciences and Engineering Council of Canada and a leave of absence given by the Department of Psychology at McGill University. I thank Laura Ann Petitto for her extensive comments on a draft of the manuscript. I also thank Tom Ward, Keith Holyoak, and Maude A. St. Laurent for their comments on a draft of this chapter. Finally, I thank all of the scientists who opened up their laboratories and their minds for this research.

IN VIVO COGNITION: A NEW WAY OF INVESTIGATING COGNITION

There is an extensive background in cognitive research on thinking, reasoning, and problem solving processes that form the foundation for creative cognition (see Dunbar, in press; Holyoak, 1995, for recent reviews). However, to a large extent, research on reasoning has demonstrated that participants in psychology experiments make vast numbers of thinking and reasoning errors even in the most simple problems. How is creative thought even possible if people make so many reasoning errors? One problem with research on reasoning is that the concepts and stimuli that the research participants are asked to use are often arbitrary and involve no background knowledge (cf. Dunbar, 1995; Klahr & Dunbar, 1988). I have proposed that one way of determining which reasoning errors are specific and which are general is to investigate cognition in the cognitive laboratory and the real world (Dunbar, 1995). Psychologists should conduct both in vitro and in vivo research to understand thinking. *In vitro* research is the standard psychological experiment where individuals are brought into the laboratory and controlled experiments are conducted. As can be seen from the research reported in this volume, this approach yields many insights into the psychological mechanisms underlying complex thinking. The use of an *in vivo* methodology in which on-line thinking and reasoning are investigated in a real-world context yields fundamental insights into the basic cognitive mechanisms underlying complex cognition and creativity. The results of in vivo cognitive research can then be used as a basis for further in vitro work in which controlled experiments are conducted. In this chapter, I outline some of the results of my ongoing in vivo research on creative scientific thinking. I relate this research to more common in vitro research to show that the in vivo method generates new basic models of cognitive processes and opens up avenues for new in vitro research.

On-Line Scientific Thinking

Scientific thinking is an ideal domain in which to develop theories of creative cognition and complex thinking (see Klahr, 1994, for a recent review

of this literature). First, scientists are constantly adding to knowledge and, less frequently, developing new concepts and theories. Second, scientists already have a rich background of knowledge in their domain that they use as a foundation for their thought. Third, much creativity occurs in groups rather than individuals alone. Contemporary science, which includes psychology, entails an experimental context that involves a group. No longer is the lone scientist under the lightbulb the norm for science. Rather, groups containing members with different levels of experience and different scientific backgrounds form the basis of contemporary science. Little is known about the way in which groups reason. Thus, scientific groups are a very important source of creative thinking and reasoning. In sum, by investigating science as it is practiced it is possible to address key questions about the nature of thinking and creativity, uncover fundamental processes that underlie complex thinking, and suggest strategies for enhancing creative thought.

Method

The research program that I have developed centers on understanding the cognitive and social mechanisms involved in current day science. I have selected molecular biology as a scientific domain to investigate because this domain is of central importance to contemporary science. Many of the brightest and most creative minds in science are attracted to this field, and molecular biology receives a very significant proportion of funding in science and medicine. Furthermore, the field of molecular biology is undergoing an immense period of scientific discovery and breakthrough, making it an ideal domain in which to investigate creative thinking.

Having identified molecular biology as a scientific domain I then sought to identify leading laboratories in the United States that I could investigate. My goal was to investigate the thinking and reasoning strategies that leading scientists use while conducting their research. After consulting with a number of scientists, including one Nobel Prize winner, and extensively reviewing the literature, I identified six world-renowned scientists at a major U.S. university. All scientists were internationally known for conducting innovative research that frequently stretched the bound-

aries of their field. Each scientist was concerned with discovering new biological mechanisms that give fundamental insights into biology. Having identified the six laboratories, I then contacted the scientists and asked them to participate in my research. All six agreed to participate in the study. I then interviewed the scientists to determine what their current research projects were, what the scientists in their labs were doing, and what their plans for the coming year were. Following this consultation I then selected four laboratories as being most suitable for investigation.

The goal of this research was to identify the points in time at which innovative scientific thinking occurs, capture this thinking on audio- and videotape, and then analyze the processes involved in the scientists' thinking and reasoning. To this end I spent a year in the four selected molecular biology laboratories. I spent the first four months becoming familiar with the scientists in the laboratory, staying in the labs during the day, attending lab meetings, interviewing the scientists in the lab, and reading grant proposals and drafts of papers. I discovered that the laboratory meeting is one of the central places in which new ideas and concepts are generated. Each laboratory had a weekly meeting that all of the members of the lab attended. The senior scientist, who manages the lab, is present as well as the postdoctoral fellows, graduate students, and technicians. In the lab meetings, a scientist presents his or her latest research, which is conducted with the senior scientist. Members of the lab ask questions about the research and propose new experiments, hypotheses, and interpretations, often forcing the presenting scientist to reconceptualize his or her ideas. Totally new concepts are generated and modified by members of the laboratory at some of the meetings. Often the senior scientist plays a crucial role in the development of new ideas and concepts. The scientists' reasoning at lab meetings is often spontaneous, and the on-line interactions concern some of the most creative moments in science. The finding that lab meetings are a central source of creative thinking and reasoning is also important because the reasoning that occurs at these meetings occurs through presentations and spontaneous interactions in which the scientists develop their ideas. Because the scientists talk out loud during the meetings there is an external record of thinking and reasoning. Using this method it is possible to directly monitor thinking and reasoning rather than uncover reasoning through post hoc interviews, questionnaires, or think-aloud protocols. The scientists externalize much of their thinking through interactions with other scientists in the lab. Thus, by recording laboratory meetings it is possible to gain access to on-line thinking and reasoning without influencing the way the scientists think.

Following my initial data collection phase, I evaluated the best method of collecting data on scientific thinking. I found that the laboratory meetings provide a much more accurate picture of the conceptual life of a laboratory than interviews, lab books, or papers. In fact, I found that the scientists were often unable to remember the steps in the development of a particular concept. The laboratory meetings provided a far more veridical and complete record of the evolution of ideas than other sources of information. Thus, I selected the laboratory meetings as the core source of data and the interviews and papers as supplemental sources of information. Thus, the particular method that I used to collect data revolved around the discovery that the laboratory meetings are central to the conceptual life of a laboratory.

I constructed a before–during–after design for uncovering the effects of laboratory meetings on the scientists' theories and methods: Before a lab meeting I interviewed the scientists to find out what their hypotheses were, what they thought the data meant, and what they were going to do next. I then audio or videotaped the scientists during the lab meeting. After the meeting I interviewed the scientists to determine whether the lab meeting affected their knowledge. I also interviewed the senior scientists about their conceptualization of the research project. This was a cyclical process in which I observed the scientists present work a number of times. By the end of the year, I had collected data on 19 scientific research projects. In addition to recording laboratory meetings I conducted interviews with members of the laboratory, was given copies of grant proposals and drafts of papers, and attended lectures by the senior scientists and many impromptu meetings. Thus, I collected data on all aspects of scientific research with the laboratory meeting as the central focus.

The Laboratories

Data on 21 scientists in the four laboratories were collected, as well as data from the 4 senior scientists. My current analyses focus on the 4 senior scientists and 19 scientists in the laboratories. Twelve of the scientists were postdoctoral fellows, 5 were graduate students, and 2 were research technicians. The four laboratories that were studied were either developmental biology labs or just worked with pathogens (disease-causing viruses and bacteria). Furthermore, the senior scientists varied in terms of experience. Two were full professors, 1 was an associate professor, 1 one was an assistant professor. By varying the types of subdomains that the scientists work in and their level of experience, it is possible to determine whether these factors influence their research.

All the scientists allowed me free access to their laboratories, to interview anyone in the laboratory, to attend any meeting, to read and keep copies of their grant proposals (including the pink sheets), to attend their talks and lectures, and to read drafts of their papers. Thus, I was given full access to the day-to-day activities of the laboratories. In addition, the scientists frequently asked me to attend impromptu meetings and discussions, and they often called when they felt that interesting events were occurring in the lab.

I selected research projects for study on the basis of whether the project had just started or was about to begin. In addition, I consulted extensively with the senior scientists in choosing the research projects to investigate. Once I had selected the projects, I then met with the senior scientists, postdocs, graduate students, and technicians that were involved in the research. All members of the four laboratories agreed to cooperate.

Laboratory A. Laboratory A was run by a senior scientist who has over 300 publications and numerous awards. This laboratory has had many discoveries that have appeared on the front page of the *New York Times, Science, Nature, Cell,* and so forth. His laboratory consisted of 22 postdoctoral fellows, 5 graduate students, and 4 technicians. I selected four research projects to follow. Two of the four research projects were successful and led to scientific discoveries. Importantly, neither I nor the scientists involved realized that a discovery was about to be made when I started following their research. It was only after a few months of following the research projects that the discoveries were made. Thus, I had collected data before, during, and after a discovery had been made. One of the researchers discovered a new gene that controls cell differentiation, and another had discovered how certain cells proliferate into certain regions of the body. Importantly, the latter discovery actually occurred during a laboratory meeting at which I was present and was tape-recording; that is, I have the moment of discovery on tape. This project forms the basis of the research discussed in the section titled "Anatomy of a Conceptual Change." Of the two remaining projects, one was unsuccessful and the other had not progressed significantly within the data collection period.

Laboratory B. Laboratory B was run by a scientist who has made many important discoveries in molecular biology. He has numerous publications and has trained many now eminent scientists. His current research program involved determining a general model of how certain genes control traits in a novel type of bacterium. His laboratory had 3 postdocs, 5 graduate students, and 1 technician. I have analyzed two of the research projects that were conducted in his laboratory. One of the research projects has resulted in two publications; however, the scientists were unable to reach their goal of discovering the function of a component of a gene. The other project made minimal progress.

Laboratory C. Laboratory C was run by an associate professor who has made a number of important discoveries on how DNA and RNA are coded in two different types of parasites. The lab consisted of 4 postdocs, 2 graduate students, and 1 lab technician. I followed research projects conducted by the 4 postdocs. All the research projects resulted in significant breakthroughs that have been published in the major scientific journals such as *Science.*

Laboratory D. Laboratory D was run by an assistant professor who is already famous for his work on viral mechanisms and his creative approach to uncovering gene function. The laboratory consisted of 4 postdocs, 6 graduate students, and 2 lab technicians. His current research program is centered on discovering the mechanism by which certain genes in

the HIV virus allow the virus to infiltrate into the host organism. He has evolved a research program that has employed a number of novel and ingenious techniques to discover how this works. These research projects are now leading to a new model of an important component of HIV activity that has wide-ranging theoretical and practical implications for molecular biology. The director of Laboratory D also invented a new genetic technique. I was present for the implementation and development of this technique. This technique has been widely referenced and reviewed in many major scientific journals.

Data Analysis

Transcription. Transcriptions and coding were done by two independent transcribers with backgrounds in molecular biology.

Coding. All coding was conducted by coding the transcriptions into a computerized database. Multiple coders were used, and reliability checks were conducted by independent coders. In this section, I provide a very general overview of the coding techniques used. I provide a more detailed account of coding in the method discussions of other sections of this chapter. The basic unit of analysis is the statement or utterance. A statement is essentially equivalent to a clause or sentence. Statements were chosen as the basic unit of analysis as they contain a verb phrase, which in turn contains the core mental operation (proposition or idea) that the presenter is employing at the time. Thus, we treat statements at meetings in the same way that statements are treated in standard protocol analyses (cf. Ericsson & Simon, 1993). I used the corpora of statements made to build a representation of scientists' mental operations. Using techniques borrowed from protocol analyses, statements can be aggregated by episodes, solution steps, and processes. One can switch between different levels of analyses, depending on the questions that one is asking of the data. The MacSHAPA coding and database software system was used to code the data (Sanderson et al., 1994).

Summary of Results

The research reported in this chapter provides a snapshot of my current analyses and interpretation of the cognitive processes involved in creativity in science. I now address three main sources of creative cognition. First, I present an analysis of the role of analogy. Second, I outline my analyses of scientists' treatment of unexpected findings. Third, I discuss some of the findings on distributed reasoning. Finally, I present a case study of a conceptual change that involved all three of the aforementioned strategies.

ANALOGY

Analogy has been regarded as a very important psychological process involved in creative cognition and has been the focus of intense investigation over the past 15 years, culminating in a number of detailed models of the cognitive processes involved in analogical reasoning (e.g., Forbus, Gentner, & Law, 1995; Holyoak & Thagard, 1989, 1994).[1] Accounts of analogy distinguish between two components of an analogy: the target and the base. The *target* is the concept or problem that the scientist is attempting to solve or explain. The *base* is another piece of knowledge that the scientist uses to understand the target, or explain the target to others. When the scientist makes an analogy he or she maps features of the base onto features of the target. By mapping the features of the base onto the target new features of the target may be discovered, or the features of the target can be rearranged so that a new concept is invented, or the scientist can highlight a specific feature of the target for other people. To illustrate this discussion of analogy I borrow an analogy that Rutherford (Rhodes, 1986) ostensibly used in his research. When Rutherford was attempting to understand the structure of the atom he made an analogy to the solar system. In this case, the target was the atom and the base was the solar system. Rutherford ostensibly mapped the idea that the planets revolve around the sun onto the atom, and he argued that the electrons revolve around the nucleus. Thus, a number of historians have argued that by drawing an analogy to the solar system, Rutherford was able to propose a new account of the structure of the atom. By mapping the feature

[1]Many cognitive accounts of analogy start with a reference to analogy in science and have noted that the types of distant analogies alluded to in the literature on the history of science are rarely used by participants in psychology experiments.

of the planets revolving around the sun, Rutherford was able to align his data with those predicted by a solar analogy. According to this view, the analogy resulted in a major restructuring of his knowledge, and a scientific discovery was made.[2]

The Rutherford example highlights two key assumptions that researchers in the creativity literature have made about the role of analogy in science. The view of analogy in the creativity literature has been that when a scientist makes an analogy (a) the source is usually from a very different domain,[3] and (b) the role of analogy is to restructure the scientist's knowledge in a gestaltlike manner (e.g., Boden, 1993; Koestler, 1964). One of the questions I want to ask here is whether this is a valid picture of the role of analogy in science. The question can be divided into a number of more detailed questions: Do scientists use analogies at all? If they do, are they the distant analogies that have been talked about in the historical creativity literature? Do less distant analogies play any role in science, as the emprical psychological work suggests (see Forbus et al., 1995; Holyoak & Thagard, 1994)? Does analogy work alone, or does it work in conjunction with other mental operations? Is analogy involved in scientific discoveries and conceptual change in science?

Method

I investigated the use of analogy at 16 meetings (4 meetings for each of the 4 labs). All analogies were coded by two independent coders. Any time a scientist referred to another base of knowledge to either (a) explain a concept or (b) use that other base of knowledge to modify the concept, it was coded as an analogy. Three representative analogies follow:

1a. *Within organism: An HIV to HIV analogy.* "Um. In the case of HIV it's 5 bases away, umm. So, um to study RT (reaction time) using a substrate that more closely mimics the in vivo situation is difficult. Because um, number one, you will need to anneal six surface strands together. Number two, it is really doubtful that since there is only a five base pair here, where they hold this complicated structure together."

1b. *Other organism: An Ebola virus to Herpes virus analogy.* "The problem with Ebola is that it is AT rich. So you can't really do some analysis, analysis of homology with the, uh, genome because of this very AT rich, uh, richness. That would not be the case for herpes and could give a better answer for some of the putative homology."

1c. *Nonbiological or distant: Monkeys to PCR (polymerase chain reaction) analogy.* "You know, just because you can see 10 molecules that still isn't working in my book. A monkey will eventually type Shakespeare, given the opportunity. PCR is not unlike that. You do it a billion times and you probably will find one thing that happened to be right."

Note that instances where a scientist stated that *X* was like *Y* were not coded as analogies. That is, statements of similarity that neither gave explanations nor resulted in the mapping of features from the base to the target were not coded as analogies. Once the analogies were found, they were coded along a number of dimensions. The coding dimensions are specified in the section dealing with that dimension.

Results

Frequency of Analogy Use

There were 99 analogies used during the 16 meetings ($M = 6.1$ analogies per meeting). The range of analogy use was 2 to 14 analogies per meeting. All four labs used analogies. There were a total of 25, 30, 31, 13 analogies, respectively, for Labs A, B, C, and D. Thus, analogies were frequently used at laboratory meetings.

Range of Analogy Use

The range over which the analogies were used was coded. Range is an index of how far apart the base and target were for each analogy. Analogies

[2]There is some controversy about whether the solar system analogy played a causal role in Rutherford's discovery of the structure of the atom. Whatever the real case may be, my point is that researchers have used such examples to emphasize the critical revolutionary role that this particular type of distant analogy plays in scientific discovery and conceptual change.

[3]Most cognitive accounts of analogy have made no assumptions about how distant the source and the target are in science. In fact, Holyoak and Thagard (1994) have made a list of the most important analogies in science over the past 2,000 years and have found very few distant analogies. However, when researchers do allude to analogy in science they tend to give examples wherein the source and the target are distant.

were coded as being *within organism, other organism,* or *nonbiological. Within-organism* analogies are those wherein the base and the target are from within the same organism. In the previous within-organism example (1a), the scientist has drawn an analogy between the way the HIV virus works in an in vivo context and how an in vitro HIV could be made by mapping from the in vivo HIV onto the in vitro HIV construct. *Other-organism* analogies are those in which the base and the target are from two different organisms, as in Example 1b (analogy between the Ebola virus and the Herpes virus). In this analogy, the scientist points out the differences between Ebola and Herpes to show why Ebola is a better organism to research a particular question. *Nonbiological* or *distant* analogies are those in which the base is taken from a nonbiological domain. In Example 1c (Nonbiological or distant: Monkeys to PCR analogy), the scientist highlights the fact that a finding could be due to chance by drawing an analogy between a monkey typing Shakespeare and the polymerase chain reaction, generating a chance result. Note that this type of distant analogy has received the most attention in the literature.

Almost all of the 99 analogies were either within organism (40) or other organism (57). There were only 2 nonbiological analogies. Thus, the bulk of analogical reasoning happened when the base and targets were from the domain of biology. This result is very important. Most accounts of analogy in science focus on distant analogies, yet only 2 of the 99 analogies used by the scientists were of this type.

Goals and Analogy Use

Categories of goals were formulated by searching for goals in the database rather than imposing them on the data a priori. From this emerged four dominant goals: formulate a hypothesis, design an experiment, provide an explanation, and fix an experiment (when an experiment went awry the scientists often drew analogies to procedures used in other experiments and proposed replacing one step in the faulty experiment with a step from an analogically similar experiment). Almost half (45) of the analogies occurred when the goal was to provide an explanation. Usually, the explanations were of methodological issues. There were 21 analogies for

design an experiment, 10 for fix an experiment, and 23 for formulate a hypothesis.

Next, I discuss the relation between goals and range. Table 1 reports the number of analogies for each combination of goal and range. The table shows a number of interesting relations between a scientist's goals and the range over which the analogy is drawn. First, it can be seen that the two nonbiological analogies were used to make explanations; they were not used to formulate hypotheses. Although there were only two nonbiological analogies in the 16 meetings coded, there were two other nonbiological analogies in the database. All four of these nonbiological or distant analogies were used to explain a concept to members of the laboratory. Thus, nonbiological or distant analogies are rare and generally used for explanations rather than to generate new hypotheses and concepts.

I now turn to a discussion of the within-organism and other-organism analogies and goals. There was little difference in range between designing and fixing experiments. Scientists were equally likely to draw an analogy from the same organism or a different organism when designing or fixing an experiment. The major interaction of goals with range was in hypothesis generation. The scientists tended to use analogies to other organisms when formulating a new hypothesis. For example, a scientist might in attempting to determine the function of a gene in one organism (e.g., a gene in malaria) draw an analogy to a gene in another organism

Table 1

Scientists' Goals for Within-Organism, Other-Organism, and Nonbiological Analogies

Type of goal	Within organism	Other organism	Nonbiological
Hypothesis	3	20	0
Design experiment	9	12	0
Fix experiment	5	5	0
Explain	23	20	2

(e.g., a similar gene in clams). If the scientist knows what the gene does in one organism (e.g., in clams), she or he can then map the functions of that gene over to the organism that they are working on (e.g., the similar gene in malaria). Thus, rather than the source of hypotheses being analogies made to nonbiological or distant domains, when formulating hypotheses the scientists make analogies to other organisms.

How do Scientists Generate Their Analogies?

How do scientists retrieve the sources for between- and within-organism analogies? One possibility is that the scientists recall specific experiments conducted in their labs or journal articles they have read. If we break down the range of the analogies by whether the scientists were recalling specific cases when they were making analogies (such as specific experiments that were conducted in the past, references to particular research articles, or experiments conducted by researchers in the field), we can see that 31 of the 40 within-organism analogies recalled a specific case. In contrast, only 6 of the 57 other-organism analogies recalled a specific case. Thus, when scientists make analogies to the same organism, they tend to recall a specific case. However, when scientists make analogies to a different organism, they do not recall specific cases. In addition, 22 of the 31 within-organism analogies recalled cases of previous experiments conducted in the lab. Thus, when the researchers made analogies to the same organism, the bulk of the analogies were to previous experiments conducted in that lab.

How did the scientists use analogies to other organisms without recalling a specific case? Psychological research has shown that individuals have great difficulty going outside their current problem to make an analogy (e.g., Gick & Holyoak, 1983), yet the scientists were able to transcend this problem. How? An analysis of analogies to other organisms revealed that the scientists had two main ways of circumventing this analog retrieval problem. First, molecular biologists have a tool available to them that gives them another way of retrieving base analogs: homology. Second, the scientists had an abstract knowledge of the biological mechanisms that exist in other organisms. Scientists can use their knowledge of

biological mechanisms to search memory for organisms that use a particular biological mechanism.

I discuss each of the ways of retrieving base analogs in turn. Scientists use homology to determine the molecular structure of a gene by sequencing each base pair in the gene. The scientists then type the sequence of their gene into a computer and search a database of genes for a gene that has a similar coding. If the scientist finds a gene or genes with a similar sequence (i.e., a homologous gene), and the function of that gene is known, the scientist can infer that the gene may have the same function in their organism. That is, the scientist maps the function of the homologous gene onto the gene being investigated. Thus, homology allows the scientist to both retrieve analogs and propose new hypotheses about gene function. Not only does the homology allow the scientist to infer new hypotheses concerning the biological function of the gene, but the scientist can also use the methodologies that the previous researchers used when conducting their research. Importantly, the same homology can provide new hypotheses and new methods that the scientist can use in his or her research. The scientists in my study generated 31 of the 57 other-organism analogies by using homology. Thus, homology allowed them to generate other potential base analogs. As can be seen from Table 2, the scientists used homology to infer biological mechanisms and the methods that they should use in their experiments.

I now turn to analogies to other organisms that were not based on homology. The scientists were more likely to use biological mechanisms

Table 2

Types of Knowledge Retrieved by Analogies on the Basis of Homology and Nonhomology for Other-Organism Analogies

Type of knowledge	Homology	Nonhomology
Biological mechanism	10	16
Experimental method	17	7
Problems with methods	4	3

as a retrieval cue. For instance, the scientists might think that "E.coli performs a particular function by splicing the protein at the AT site; perhaps our organism splices the protein in the same way." The scientists' knowledge of biological mechanisms is often tied to particular organisms, and these organisms become part of the analogy.

Summary and Discussion of Analogy Results

Analogy was frequently used in all of the laboratories. Most of the analogies that were observed in the current study were biological. Only 2 of the 99 analogies were nonbiological or distant. These findings shed new light on the role of analogy in science. Most historical accounts of analogy in science have tended to focus on very distant analogies; yet, the results of these investigations suggest that distant analogies are not an important component of contemporary science. There are a number of reasons for the differences between these findings and those discussed in the literature on the history of science or the creativity literature. First, many of the distant analogies that scientists have mentioned in the history of science may not have had a role in the making of a discovery. In fact, a number of historical analysts have argued that the Rutherford solar system analogy, and the snake analogy mentioned by Kekulé in his discovery of the structure of the benzene ring had no role in the respective discoveries (Rhodes, 1986; Wotiz & Rudofsky, 1984).[4] The data presented here suggest that it may be the case that scientists use distant analogies to explain a new concept to an audience rather than that distant analogies have a causal role in making a discovery. I am currently monitoring the scientists' publications to see if more distant analogies seep into their accounts of their findings.[5] Second, the types of analogies that the scientists use in on-line reasoning are easy to forget. In fact, in postlab meeting interviews

[4]I thank Bill Brewer for bringing this article on Kekulé's discovery of the benzene ring to my attention.

[5]Some researchers have suggested that perhaps the evidence of distant analogies is an index of the maturation, or lack thereof, of the development of a field (with presumably a higher incidence of distant analogies occurring at the beginning of a field). However, there is nothing in my data that supports this view. Note that the scientists in my study were pioneering totally new concepts, in an uncharted conceptual space. In this view one would expect to see many distant analogies relative to the other types, which was not the case.

the scientists rarely remembered the analogies that were generated during the meeting. Thus, analogies are often used as a scaffolding that the scientists use in the construction of new theories and methodologies. Once the new concepts and methods have been advanced the analogy can be discarded. Many of these analogies will not make their way into the notebooks of the scientists, and thus the historical record will not show that the within-organism or other-organism analogies had a role in the discovery of a new concept or invention of a new method.[6]

More than one analogy may be involved in a particular discovery, and one particular analogy may not be responsible for a particular conceptual change, but a group of quite different analogies may be causally involved in making a breakthrough. Again, because no one analogy made a major restructure of knowledge, the scientists may not have recalled a particular within-organism or other-organism analogy as being a factor in the discovery. However, when all of the analogies that are involved in making a discovery are examined, only analogies of very specific types will be seen to have played a major role in scientific reasoning and discovery. Moreover, as is shown later, analogy is not the only mechanism that comprises conceptual change. In the next three sections I show that other key cognitive mechanisms produce conceptual change. Thus, analogy, while important, is but one of a complex of mechanisms that produce conceptual change. At the close of this chapter, I explain what this complex of mechanisms is and show how together they contribute to scientific discovery and conceptual change.

UNEXPECTED FINDINGS AND CONFIRMATION BIAS

There is a large literature in psychology and philosophy of science on what happens when scientists get unexpected results from their experiments. In the psychological literature researchers have investigated this in terms of

[6]My in vitro investigations of analogical reasoning also reveal that research participants have little awareness of, or memory for, the mental steps involved in making a discovery, even directly after having made a major conceptual shift (Dunbar & Schunn, 1990; Schunn & Dunbar, 1996).

confirmation biases; individuals tend to seek evidence that is consistent with their hypothesis and ignore evidence that is not. Researchers have repeatedly found evidence that research participants engage in this type of behavior and have argued that scientists have similar reasoning biases (cf. Klayman & Ha, 1987; Tweney, Doherty, & Mynatt, 1982). However, before one accepts the generality of the results of these types of experiments it is important to note some of the large differences between the tasks that research participants perform and what scientists do. Most important, there is no actual scientific knowledge involved in the psychological tasks; the to-be-discovered concepts are arbitrary, and the links between hypothesis, experiment, and data are straightforward. For example, the 2–4–6 task is one that has been widely used (Wason, 1960). In this task, the experimenter asks an individual to determine the rule underlying a sequence of numbers. The individual is given a triad of numbers, such as the numbers 2, 4, 6, and is told that this number triad is an example of the rule. The individual is then told that she or he can generate other triads and that the experimenter will determine whether the triad is an example of the rule. Finally, the individual is told that when certain she should state the rule. Many research participants tend to generate triads that are consistent with their hypotheses; they attempt to confirm their hypotheses.[7] On the basis of experiments such as these, researchers have argued that this confirmation bias is a general phenomenon that both lay people and scientists must avoid if they are to reason correctly.

Although the confirmation bias view of science has received much empirical support, another related phenomenon is the issue of unexpected findings. A number of researchers have argued that a useful strategy in science is to focus on unexpected findings. According to this view, scientists work with a heuristic such as "if the finding is unexpected, then set a goal of discovering the causes of the unexpected finding" (cf. Dunbar, 1993, 1996; Kulkarni & Simon, 1988). This view of reasoning is quite different from that implied by the confirmation bias viewpoint. According to this viewpoint,

[7]The common rule that research participants must discover is "numbers of increasing magnitude." Research participants generally propose the rule "even numbers increasing by 2" and only generate triads consistent with this rule.

when unexpected findings are inconsistent, scientists should focus on the finding rather than ignore it. Thus, there are two conceptions of what scientists may do. Of course, it may be the case that under certain circumstances the scientists may focus on unexpected findings, and under other circumstances they may ignore the findings and behave like the participants in psychology experiments (cf. Tweney, 1989). The goal of the following analyses was to investigate these questions in a real scientific environment.

Method

My graduate student Lisa Baker and I decided to investigate the role of unexpected findings by analyzing the scientists' reactions to unexpected findings at four laboratory meetings in Lab A (see also Baker & Dunbar, 1996). We chose Lab A because scientists there had obtained many expected and unexpected findings and thus provided much data to investigate these issues. We had two independent coders code every unexpected finding in each of the four lab meetings. All findings in which the scientist had previously predicted a different result, or expressed surprise at the obtained result, were coded as *unexpected.* All findings that were consistent with the predictions were coded as *expected.* A third category of findings consisted of those that occurred in exploratory experiments. Here the scientist did not have any predictions one way or the other and conducted the experiment to see what would happen. The results of these types of experiments were coded as *exploratory.*

Results

Our first step was to determine how common expected, unexpected, and exploratory findings were? In four meetings there were six experiments reported with 70 conditions. There were 22 expected, 18 unexpected, and 30 exploratory findings. Clearly, unexpected findings are common. We coded all expected and unexpected findings on the basis of whether the scientists tried to explain away their results or whether they built theories with the findings. To do this, we coded the number of reasoning blocks the scientists engaged in following both expected and unexpected findings. A reasoning block was a group of statements that in-

volved reasoning about a particular finding. One finding can generate many different reasoning blocks. The number of reasoning blocks generated by expected and unexpected findings can be used as an index of how much attention scientists give to these types of findings. There was more reasoning for unexpected (179 reasoning blocks) than expected (42) findings. Furthermore, when confronted with unexpected findings, scientists were much more likely to engage in theory building than to attempt to explain the results away. Thus, 161 reasoning blocks were concerned with theory building, and 18 reasoning blocks were concerned with attributing the result to some sort of error. Thus, scientists do pay attention to unexpected findings.

The previous analyses applied purely to whether the scientist who conducted and presented the research was likely to attend to an unexpected finding. We next investigated whether the group also attended to unexpected findings. We calculated the number of reasoning blocks that anyone other than the presenter devoted to unexpected and expected findings. Again, we found much more reasoning by the group when faced with unexpected compared with expected findings. As a measure of group attention to unexpected findings we also counted the number of interactions for expected and unexpected findings. We found 23 interactions for expected findings and 176 interactions for unexpected findings. These results indicate that the group also pays attention to unexpected findings and uses the findings to propose new hypotheses and experiments.

Another question that can be asked about the scientists' use of unexpected findings is whether there was any difference between the scientists' treatment of unexpected findings that were consistent with their hypothesis and those that were inconsistent with their hypothesis. An unexpected finding that is consistent with a scientist's hypothesis can occur, for example, when the scientist expects a certain type of result to occur, but the size of the effect is much greater than expected. In this type of situation, the result is consistent with the hypothesis but the size of the effect is unexpected. An unexpected inconsistent finding is one in which a qualitatively different type of outcome occurs. We coded the 18 unexpected findings along these dimensions and found that 8 unexpected findings were

consistent with scientists' expectations and that 10 unexpected findings were inconsistent with their expectations. We then coded the findings that resulted in the proposal of new hypotheses. We found that 4 of the 8 consistent findings resulted in new hypotheses and 8 of 10 inconsistent findings resulted in new hypotheses. These results indicate that the scientists attended to the unexpected findings even when the findings were inconsistent with their hypothesis.

We recently have been conducting new analyses of scientists' reactions to unexpected findings to determine whether the time at which an unexpected finding occurs affects whether an unexpected finding is attended to. We found that there are two dimensions of an unexpected finding that determine whether the unexpected finding is attended to. The first is whether the unexpected finding is unexpected relative to a core hypothesis in the field or to an auxiliary hypothesis that the scientist has proposed to get the experiment to work. Another dimension is how early or late in the research project the unexpected finding occurs. We found that when the unexpected finding occurs early and is not a core hypothesis, the scientists will not devote much attention to it. However, if the unexpected finding occurs early and is unexpected relative to the central assumptions of the field, the scientists will focus on the finding. When the unexpected finding occurs late in the research project the scientists will attend to it regardless of whether it is a core or an auxiliary hypothesis. Note that the situation in which the scientists ignore unexpected findings is very similar to that of individuals in psychology experiments: The individuals are early in the experiment, and the hypotheses are not core assumptions. Thus, it is only under very restricted circumstances that one finds a similarity between the results of psychology experiments on confirmation bias.

Our analyses of unexpected findings indicate that scientists do attend to unexpected and inconsistent findings. Why do the scientists attend to unexpected and often inconsistent findings? One reason is that in real science unexpected findings are frequent. The fact that unexpected findings are frequent may have a major effect on the scientists' ability to deal with these types of findings. It may be the case that the longer a scientist is in

the field, the more unexpected findings the scientist has encountered and the more likely it is that the scientist has developed strategies or heuristics for dealing with them. Thus, the way a scientist deals with unexpected findings depends on the specific strategies he or she has developed to attempt to reconcile them.[8] Participants in psychology experiments are unlikely to have developed strategies for dealing with unexpected findings and may prefer to focus on their current goal, ignoring unexpected results (as in Dunbar, 1993). As they encounter more and more evidence that is inconsistent, they are eventually forced to attend to unexpected findings.

DISTRIBUTED REASONING

Most cognitive research on scientific reasoning focuses on individuals reasoning about a problem. However, much of modern science is conducted by groups of scientists rather than individuals. Furthermore, much of the cognitive work has demonstrated that individuals make many different types of reasoning errors. In this section, I investigate whether reasoning in groups can circumvent certain individual reasoning errors. In particular I explore the issue of distributed reasoning in science. *Distributed reasoning* happens when different members of a group reason about topics such as a hypothesis, experiment, methodology, or interpretation of a result while adding new elements to the topic under discussion. The question that I will ask is whether distributed reasoning of this sort helps circumvent problems that individual participants display in standard experiments.

One of the major tasks for both individuals in psychology experiments and scientists confronted with new data is to determine what types of inductions to make from new data. There are infinitely many inductions that can be made from a set of data, and this is a potential place where different members of the group can make different inductions from the same data. To examine this, I explored the role that the group played in the types of inductions that a scientist in an HIV lab made during his talk. At this talk the scientist presented five sets of findings and made 11 inductions about the mechanisms that the HIV virus uses. The members of the lab often disagreed with the inductions that the scientist made and modified his inductions. The other members of the lab limited (3), expanded (1), replaced (2), or discarded (1) a total of 7 of the 11 inductions.

This pattern of challenging inductions was ubiquitous across all labs and provides important information about the role of distributed reasoning. Individuals have great difficulty generating alternative inductions from data and also have great difficulty limiting and expanding inductions. Distributed reasoning helps circumvent these difficulties. When distributed reasoning occurs, the group quickly focuses on the reasoning that has occurred, and the other members of the laboratory generate different representations. These new representations make it possible for them to propose alternative inductions, deductions, and causal explanations. Thus, distributed reasoning provides new premises and models that an individual may not be able to generate when reasoning alone.

Another issue relevant to distributed reasoning is the number of people involved in the reasoning. In the previously mentioned HIV lab I investigated the number of inductions and deductions that were shared. That is, how many inductions and deductions occur in which one premise is provided by one person, and another premise is provided by another person. We found that 30% of inductions and deductions were shared by more than one individual. We also found that 12% of all inductions and deductions had more than two participants. Furthermore, inductions of one individual sometimes formed the basis of a deduction for other individuals.

Distributed reasoning consists of scientists performing cognitive operations on information (e.g., induction) and then passing the results of the operation on to other scientists in the group. The other scientists then use the results of the first operation as the input to further cognitive operations. Together, the results of these cognitive operations are then used

[8] Lovett and Anderson (in press) have shown that history of success plays a role in determining what strategy research participants use to solve a problem. They have shown that research participants use both their history of success and the current problem-solving context to determine the type of problem-solving strategy to use. I argue that scientists use a similar set of heuristics. Whether they will use unexpected findings or not will depend on both the history of success and current context.

to build new cognitive representations: scientific theories and new experiments. How and when the information is passed between individuals depends on the goals of the individuals and the group, as well as the knowledge bases that the scientists have at their disposal. The generation of different representations during distributed reasoning helps scientists circumvent one of the major reasoning difficulties that individuals have: that of generating alternative hypotheses, explanations, theories, and experiments.

The results of these analyses of distributed reasoning are different from the results of brainstorming experiments and creativity in group experiments. Many studies have shown that when a group of people is asked to generate novel concepts, the group performs no better than individuals. However, in the research reported in this chapter it can be seen that groups of scientists do generate new concepts and that distributed reasoning is an important factor. The difference in findings is twofold. First, in psychology experiments the participants are not part of a group sharing common knowledge and values. Usually, individuals in psychology experiments are randomly thrown together for the purpose of the experiment. Second, the types of problems given to the participants are often arbitrary and require little background knowledge. In the science labs investigated in this chapter, the scientists had overlapping backgrounds and shared goals about the research. Furthermore, the members of the lab had slightly different types of knowledge that they could bring to bear on the problem. Taken together, these results suggest that entirely new experiments on group reasoning need to be conducted using real groups' reasoning about real problems, with significant background knowledge and diversity of knowledge. The prediction is that in this type of situation groups of individuals perform more creatively than individuals.

ANATOMY OF A CONCEPTUAL CHANGE

The account of the cognitive processes underlying scientific creativity offered so far is static. I have demonstrated that analogy is an important part of current day science, that scientists reason about unexpected findings,

and that distributed reasoning is a potentially important concept in science. I now turn to the issue of how all three aspects of scientific reasoning form a complex of mechanisms that work together to produce a conceptual change in a group of scientists at a meeting.

Many recent analyses of theory change in the history of science have focused on the notion of conceptual change (e.g., Carey, 1992; Nersessian, 1992; Thagard, 1992). *Conceptual change* has been defined as changes in scientific theories that occur when new concepts are proposed and old concepts must be radically changed or replaced to accommodate the new concepts. One example of this type of conceptual change noted in the literature is the 16th-century shift from a unitary concept of heat and temperature to two new concepts: one involving heat and one involving temperature (Wiser & Carey, 1983). It is conceptual change of this type that I now discuss.

Here I provide a dynamic account of a conceptual change that occurred in Lab A and use this example to show how in this situation different forms of reasoning worked together to produce entirely new concepts. To preserve the anonymity of this lab, I have been obliged to change the names of the diseases and the specific mechanisms involved in the diseases. Alas, I have also had to render intentionally vague specific aspects of the scientists' discussion that factored critically in the conceptual change. Nonetheless, I have tried hard to leave intact the essence of the complex of mechanisms that contributed to these scientists' conceptual change.

Let me begin with some background on the discovery. A postdoctoral fellow had recently come to a world-famous immunology lab. He had decided to investigate the way that B-cells cause a particular autoimmune disease. He had been conducting experiments in collaboration with another postdoc in another lab. Their work began with an analogy. Twenty years before, a researcher had noticed that an autoimmune disease in rabbits called CVX was very similar to a human autoimmune disease. Since then, the CVX diseases in rabbits has been used as a model for the human disease LOA. The postdocs investigated the disease in yet another organism (hamsters) because the postdoc's lab used hamsters and had facilities

that could be used to investigate the mechanisms underlying the CVX disease that few other laboratories had. Overall, the motivation for his research was based on analogies between the human LOA disease, the rabbit CVX disease, and the hamster CVX disease.

One May afternoon the postdoc gave a talk about his latest experiments. He began with analogies between the human LOA disease and the CVX disease in rabbits, noting where the similarities and differences between the two diseases arose. He then moved to analogies between the CVX disease in rabbits and in hamsters. The first set of experiments resulted in a small amount of discussion and suggestions for future experiments. Then the postdoc started to discuss some experiments in which the results were very unusual. The postdoc had conducted a straightforward experiment. He had two conditions: one that caused colmenia disease in the joints and the other that caused the CVX disease in the heart. Both the heart and the joints are immune-privileged sites that do not normally allow B-cells in. In fact, the only types of B-cells that have been found in the heart are CVX B-cells, and the only B-cells that have been found in the joints are the colmenia B-cells. The postdocs expected that the B-cells that cause the disease in the heart would go to the heart and the B-cells that cause the disease in the joints would go to the joints. Instead, they found both types of B-cells in the heart and in the joints. This was an unexpected finding. The postdoc reached the part of his presentation wherein he discussed these results. He was surprised and excited by what he found. The result was unusual. It was at this point that the conceptual change began to unfold.

The director of the lab was intrigued. He asked the postdoc how it happened. The postdoc said he did not know. The director then made the question more specific. He asked the postdoc what properties were common to the colmenia and CVX B-cells that allowed them entry into the heart. The postdoc made an analogy to some other experiments that another postdoc in the lab had conducted and induced that the CVX and colmenia cells were both methylated. The director and other postdocs in the lab then made a series of inductions and deductions that led to a causal explanation for the unexpected finding. The reasoning was distributed

over the members of the lab. However, the explanations they offered did not account for some other aspects of the findings, and another round of distributed reasoning occurred. This distributed reasoning resulted in a conceptual change: They proposed two new biological mechanisms to replace the unitary concept they had all assumed up to that point. Previously, it had been assumed that CVX cells only go to the heart and that colmenia cells only go to the joint; that is, B-cells have organ-specific attractions. The assumption was that once these cells got into the organ, they started the disease in that organ. Thus, there was one mechanism that caused both the entry into the organ and the initiation of the disease. The members' distributed reasoning led them to the conclusion that entry into the organ and the initiation of the disease were caused by two different mechanisms. They then had to propose what these mechanisms could be. They proposed two mechanisms that could together account for the CVX B-cells' causing the disease. One postdoc drew an analogy back to the human disease and mapped the mechanisms that had been proposed for the CVX disease onto the human disease. They modified their new model to fit the analogy to the human disease and thus ended up proposing a new model that not only explained the mechanisms underlying the three diseases but also had major ramifications for whole classes of autoimmune diseases.

By proposing two new mechanisms the scientists had to also change a number of other concepts in their knowledge of autoimmune diseases. It was at this point that everyone in the lab realized that a conceptual change had occurred, and they all shouted in excitement. This was followed by some further analogies in which other postdocs suggested other experiments. Finally, a postdoc made an analogy to the methods that other researchers have used and the methods that the postdoc had used, explaining why their rival's lab had not made the discovery they had just made.

This account of a conceptual change reveals some important characteristics of the mechanisms underlying conceptual change. First, there was no one reasoning mechanism underlying the conceptual change. Analogy, induction, deduction, causal reasoning, and distributed reasoning were all involved. Second, analogy was a significant component of the conceptual

change, but all of the analogies that were used were either to the same organism or to other organisms: Conceptual change can and does occur without distant analogies. Third, the scientists had little memory for any of the on-line analogies used at the meeting. I asked the postdoc who conducted the research how the "discovery" was made. I asked this question one week later, a month later, three months later, and nine months later. On none of these occasions did he recall the spontaneous analogies used, or that distributed reasoning was involved. Thus, much of the on-line cognitive processes that went into the conceptual change would have disappeared without a record if I had not taped the original meeting.

CONCLUSION: CREATIVE COGNITION IS A TINKERER

The investigation of the cognitive mechanisms involved in on-line scientific thinking and reasoning reveal a number of important mechanisms underlying creative cognition. The main idea is that no single cognitive process is responsible for creative thought. I have found that scientists use a variety of cognitive mechanisms to produce any single new concept or theory. Creative ideas and novel concepts arise through a series of small changes produced by a variety of different cognitive mechanisms. It may be the case that reasoning and conceptual change are related in much the same way that a series of minor mutations produce major changes in organisms during evolution. In conceptual change, small mutations in concepts occur due to analogy and other reasoning mechanisms. Overall, a series of small changes will produce major changes in a concept. Conceptual change, like evolutionary change, is the result of tinkering. From a psychological point of view this account of conceptual change explains why it is so hard to discover the underpinnings of creativity. The many incremental steps that are involved in creative cognition are often lost and forgotten, and the act of creation becomes a mythical entity in which the final step in the creative process is often seen as the cause of the new concept. This leads to the proposal of entities such as distant analogies and insight as more important in creativity than they really are.

A further question is whether the cognitive processes underlying creative conceptual change are different from the processes underlying simple changes in concepts. I would argue that they are not. Exactly the same types of cognitive processes that are involved in the more mundane aspects of conducting science were involved in the moments of true conceptual change outlined previously. The question then arises as to what has made these scientists so productive and what has launched them to the forefront of their fields? The answer lies in their choice of research topics. Each of the scientists has developed research programs around difficult topics for which there were few simple answers or an abundance of ready-made techniques available. To conduct their research the scientists had to invent new techniques and engage in research that was risky. Thus, the factor that unifies the creative scientists in this sample is their ability to take risks. Each of the scientists conducted both high- and low-risk experiments in their laboratories. Although taking risks does not in itself lead to success, risk taking in combination with the use of the various reasoning strategies discussed in this chapter provide the context in which discoveries can be made.

The view of creativity offered here is quite different from that offered in the creativity literature. Authors such as Boden (1993) have proposed that the main way that analogy is involved in creative discoveries in science is by having major restructuring of concepts. Here, I have argued that analogy is involved in a very different way. Many very specific analogies are made that in conjunction with other reasoning mechanisms produce both modifications in existing concepts and entirely new concepts. The reason for the difference between my conclusions and that of others in the creativity literature is the differences in methodologies used. By looking at on-line reasoning rather than scientists' patchy reconstructions of a scientific discovery or breakthrough, it is possible to discover the specific cognitive mechanisms underlying creative thought. As I have argued elsewhere in this chapter, much of the cognition involved in creative thought works as a form of scaffolding. Once a new concept is generated the cognitive scaffolding is thrown away and scientists cannot reconstruct the cognitive steps that went into the discovery. Because of this, scientists, like histori-

ans, reconstruct their creative moments, often from their lab books. Unfortunately, many of the key cognitive steps made in a discovery do not end up in the lab books. Thus, many of these reconstructions are based on partial information, and, as a result, myths surrounding the creative process develop.

An important question about the research presented in this chapter is whether the findings are generalizable to other domains. There are numerous reasons to expect that these findings are indeed generalizable. First, I have observed similar types of reasoning strategies in biology laboratories at other universities (Dunbar, Patel, Baker, & Dama, 1995). Second, we have observed similar reasoning in clinical situations wherein medical doctors reason about patients (Dunbar et al., 1995; Patel, Dunbar, & Kaufman, 1995). I am now starting to investigate whether the same types of reasoning strategies occur in a business context.

Molecular biologists have some special tools that other scientists and nonscientists do not have, such as the use of the structure of DNA to search for homologies in a database. These scientists have the advantage of a way of representing their data that makes it possible to quickly and efficiently search for analogs. By representing their knowledge in a standardized fashion and searching for structural patterns that are similar to the one they are interested in, the scientists solve the problem of how to retrieve relevant analogies. Thus, homology makes finding base analogs easier. Ultimately, using homology gives the scientists another route to access base analogs. Once the scientists retrieve these base analogs, they use the same cognitive processes for constructing analogies as they do when they search their own memories for base analogs. Can scientists in other domains retrieve analogs in a similar fashion? The answer depends on the way the knowledge in a field is codified. If knowledge is coded in a structural manner, then it should be possible for the scientists to search for analogs with a similar structure and generate new analogies. It will be interesting to see whether the new databases that have arisen in virtually all fields will allow scientists to encode structural information, thereby allowing the scientists in a field to retrieve source analogs. This would ease one step in drawing analogies and could serve as a useful aid to scientists in all fields.

Overall, the research reported in this chapter demonstrates that it is possible to investigate complex creative cognition in real-world contexts. This in vivo research makes it possible to discover fundamental mechanisms of creative cognition and how multiple cognitive processes work together to produce conceptual change. Furthermore, this in vivo approach both makes it possible to discover what aspects of in vitro research are generalizable and suggests new types of experiments that can be conducted in the cognitive laboratory.

REFERENCES

Baker, L., & Dunbar, K. (1996). *How scientists use unexpected findings.* Manuscript in preparation.

Boden, M. (1993). *The creative mind: Myths and mechanisms.* New York: Basic Books.

Carey, S. (1992). The origin and evolution of everyday concepts. In R. N. Giere (Ed.), *Minnesota studies in the philosophy of science. Vol XV: Cognitive models of science* (pp. 89–128). Minneapolis: University of Minnesota Press.

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.), *Mechanisms of insight* (pp. 365–395). Cambridge, MA: MIT Press.

Dunbar, K. (1996). *Online analogical reasoning in scientific laboratories.* Manuscript submitted for publication.

Dunbar, K. (in press). Problem solving. In W. Bechtel & G. Graham (Eds.), *A companion to cognitive science.* Cambridge, MA: Blackwell.

Dunbar, K., Patel, V., Baker, L., & Dama, M. (1995, November). *Group reasoning strategies in knowledge rich domains.* Paper presented at the 36th annual meeting of the Psychonomic Society, Los Angeles, CA.

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.

Forbus, K., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 14,* 144–206.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 14*, 1–38.

Holyoak, K. J. (1995). Problem solving. In E. Smith & D. Osherson (Eds.), An invitation to cognitive science (2nd ed., pp. 267–296). Cambridge, MA: MIT Press.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295–355.

Holyoak, K. J., & Thagard, P. (1994). *Mental leaps.* 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.

Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review, 94*, 211–228.

Koestler, A. (1964). *The act of creation.* London: Macmillian.

Kulkarni, D., & Simon, H. A. (1988). The processes of scientific discovery: The strategy of experimentation. *Cognitive Science, 12*, 139–176.

Lovett, M. C., & Anderson, J. R. (in press). History of success. *Cognitive Psychology.*

Nersessian, N. (1992). How do scientists think? Capturing the dynamics of conceptual change in science. In R. N. Giere (Ed.), *Minnesota studies in the philosophy of science. Vol. XV: Cognitive models of science* (pp. 3–44). Minneapolis: University of Minnesota Press.

Patel, V., Dunbar, K., & Kaufman, D. (1995, November). *Distributed reasoning in medical and scientific contexts.* Annual Meeting of the Psychonomic Society, Los Angeles, CA.

Rhodes, R. (1986). *The making of the atomic bomb.* New York: Simon & Schuster.

Sanderson, P. M., Scott, J. J. P., Johnston, T., Mainzer, J., Watanabe, L. M., & James, J. M. (1994). MacSHAPA and the enterprise of exploratory sequential data analysis (ESDA). *International Journal of Human Computer Studies, 41*, 633–681.

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

Thagard, P. (1992). *Conceptual revolutions.* Cambridge, MA: MIT Press.

Tweney, R. D. (1989). Five questions for computationalists. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation* (pp. 471–484). San Mateo, CA: Morgan Kaufmann.

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

Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology, 12*, 129–140.

Wiser, M., & Carey, S. (1983). When heat and temperature were on. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 267–297). Hillsdale, NJ: LEA.

Wotiz, J. H., & Rudofsky, S. (1984, August). Kekulé's dreams: Fact or fiction? *Chemistry in Britain*, pp. 720–723.