How scientists think: On-line creativity and conceptual change in science

Kevin Dunbar, McGill University

I have now Moved to

Address all correspondence to:
Kevin Dunbar
Department of Psychological and Brain Sciences
Dartmouth College, Hanover, NH 03755, USA

Chapter appeared in T.B. Ward, S.M. Smith & J. Vaid (Eds.) Conceptual structures and processes: Emergence, discovery , and change. Washington D.C: American Psychological Association Press

How scientists think:
On-line creativity and conceptual change in science

Kevin Dunbar
McGill University

This is an investigation of "Online Creativity." I will 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 will 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 the likelihood of scientists 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.

InVivo cognition: A new way of investigating cognition

There is a large background in cognitive research on thinking, reasoning and problem solving processes that form the foundation for creative cognition (see Dunbar, in press, Holyoak 1996 for recent reviews). However, to a large extent, research on reasoning has demonstrated that subjects 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 subjects 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 what 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 InVitro and InVivo research to understand thinking. InVitro research is the standard psychological experiment where subjects 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 InVivo methodology in which online 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 InVivo cognitive research can then be used as a basis for further InVitro work in which controlled experiments are conducted. In this chapter, I will outline some of the results of my ongoing InVivo research on creative scientific thinking and relate this research back to the more common InVitro research and show that the InVivo method generates new basic models of cognitive processes and opens up avenues for new InVitro research.

Online scientific thinking

Scientific thinking is an ideal domain within 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 the individual. Contemporary science, which includes psychology, entails an experimental context that involves a group. No longer is the lone scientist under the light bulb the norm for science. Rather, groups containing members of different levels of experience and different scientific backgrounds form the basis of contemporary science. Little is known about the way in which groups reason. Scientific groups are a thus 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 on 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 around 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 takes up 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 breakthroughs, making it an ideal domain within 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 an extensive search of the literature, I identified six world-renowned scientists at a major US University. Each scientist was internationally known for conducting innovative research that frequently stretched the boundaries 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 are, what the scientists in their labs are doing and their plans for the coming year. 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 video tape, and then analyze the processes involved in the 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, reading grant proposals, and drafts of papers. My goal was to identify the points at which creative scientific thinking occurred. What I discovered was that one of the central places in which new ideas and concepts were generated was the laboratory meeting. Each laboratory has a weekly meeting that all the members of the lab attend. The senior scientist who runs the lab is present as well as the post-doctoral fellows, graduate students and technicians. Lab meetings consist of a scientist, presenting his or her latest research that is conducted with the senior scientist who is running the laboratory. Members of the lab ask questions about the research, propose new experiments, hypotheses and interpretations, often forcing the presenting scientist to reconceptualize his or her ideas. At some meetings totally new concepts are generated and modified by members of the laboratory. Often the senior scientist who runs the laboratory plays a crucial role in the development of new ideas and concepts at the meeting. The scientists' reasoning at lab meetings is often spontaneous and the online 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 are talking out loud there is an external record of thinking and reasoning. Using this method it is possible to directly monitor thinking and reasoning rather than uncovering 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 "online" thinking and reasoning without influencing the way that 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 designed a pre-present-post design for uncovering the effects of laboratory meetings on the scientists'' theories and methods: Prior to a lab meeting I interviewed the scientist to find out what their hypotheses were, what they thought the data meant, and what they were going to do next. Then I recorded the lab meeting (videotape or audiotape). After the lab meeting I interviewed the scientist to determine whether the lab meeting had an effect on their knowledge. I also interviewed the senior scientist on their conceptualization of the research project. This was a cyclical process in which I observed the scientist present work a number of times. By the end of the year, I had collected data on 19 different scientific research projects. In addition to recording laboratory meetings I conducted interviews with members of the laboratory, was given copies of grant proposals, drafts of papers, and attended lectures by the senior scientists, and attended many impromptu meetings. Thus I collected data on all aspects of scientific research with the central focus being the laboratory meetings.

The Laboratories

Data on 21 scientists in the four laboratories were collected, as well as data from the four senior scientists running the laboratories. My current analyses focus on the four senior scientists and 19 scientists in the laboratories. Twelve of the scientists were Postdoctoral fellows, five were graduate students and two were research technicians. The four laboratories that were studied were either developmental biology labs, or worked with pathogens (disease causing viruses and bacteria). Furthermore, the senior scientists varied in terms of experience. Two were full professors, one was an associate professor, and one was an assistant professor. By varying the types of sub-domains that the scientist were working in, and the level of experience of the scientists it is possible to determine whether these factors have an influence on 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 total and complete access to the day to day activities of laboratories. In addition, the laboratories were so cooperative that they frequently asked me to attend impromptu meetings and discussions within the laboratory, or they would call me to come over when they felt that interesting events were occurring in the lab.

Within each laboratory, particular research projects were selected for study on the basis of whether the research projects had just started or were about to begin. I consulted extensively with the senior scientists in choosing the research projects to investigate. Using this information, I selected the research projects to investigate. Once I had selected the projects, I then met with the senior scientists, post-docs, graduate students, and technicians that were involved in the research. All the members of the four laboratories were willing to cooperate.

Laboratory A. This laboratory is run by a senior scientist who has over 300 publications, and has won numerous awards. This laboratory has had many discoveries which have appeared on the front page of the New York Times, Science, Nature, Cell etc. His laboratory consisted of 22 post-doctoral 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 was made. One of the researchers discovered a new gene that controls cell differentiation, and the other

researcher 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 audio taping; that is, I have the moment of discovery on tape. This project will form the basis of the research discussed in the Section entitled "Anatomy of a Conceptual Change." One of the other two research projects was unsuccessful, and the other research project had not progressed significantly within the data collection period.

Laboratory B. This laboratory is run by a senior 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 is concerned with determining a general model of how certain genes control traits in a novel type of bacterium. His laboratory had 3 post docs, 5 graduate students and 1 technician. I have analyzed two of the research projects that were being 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 did made minimal progress.

Laboratory C. This laboratory is 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 post docs, 2 graduate students and 1 lab technician. I followed research projects conducted by the four post-docs. All the research projects resulted in significant breakthroughs that have been published in the major scientific journals such as Science.

Laboratory D. This laboratory is 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 had 4 post-docs, 6 graduate students, and 2 lab technicians. His current research program is centered around 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 for which 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 made by two independent transcribers with a background in molecular biology.

Coding. All coding was conducted by coding the transcriptions into a computerized database. Multiple coders are used, and reliability checks are conducted by independent coders. In this section I will provide a very general overview of the coding techniques used, I will provide a more detailed account of coding in the method parts of the 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 a meeting in the same way that statements are treated in standard protocol analyses (cf. Ericsson & Simon, 1992): We use the corpora of statements made to build a representation of scientists' mental operations. Using techniques borrowed from protocol analyses, we can aggregate individual statements by episodes, solution steps, and by processes. We can switch between different levels of analyses, depending on the questions that we are asking of the data. The MacSHAPA coding and database software system was used to code the data (Sanderson, Scott, Johnston, Mainzer, Watanabe, & James, 1993).

Summary of Results

The research reported in this chapter provide a snapshot of our current analyses and interpretation of the cognitive processes involved in creativity in science. I will address three main sources of creative cognition. First, I will present an analysis of the role of analogy. Second, I will outline our analyses of scientists treatment of unexpected findings. Third I will discuss some of our findings on distributed reasoning. Finally I will present a case study of a conceptual change that involved all three of the above 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 last 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. What the scientist does when he or she makes an analogy is to map 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. I will illustrate this discussion of analogy with an analogy that Rutherford 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 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 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 demonstrates 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 it is (a) usually from a very different domain[3], and (b) the role of analogy is to restructure the scientist's knowledge in a gestalt-like 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 broken down into a number of more detailed questions: Do scientists use analogy at all? If they do, is it the distant analogies that people have talked about in the historical creativity literature? Do less distant analogies play any role in science, as the emprical psychological work would would suggest (see Forbus et al. 1995; Holyoak & Thagard, 1994)? Does analogy work alone, or does it work on conjunction with other mental operations? Is analogy involved in scientific discoveries and conceptual change in science.

Analogy Method

The use of analogy at l6 meetings was investigated (4 meetings from each of the 4 labs). Every analogy used at each of the meetings was coded by two independent coders. Any time that 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 are provided in Table 1. 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 will be specified in the section dealing with that dimension.

Analogy Results

Frequency of Analogy Use. There were 99 analogies used in the 16 meetings (mean of 6.l analogies per meeting). The range of analogy use went from 2 to 14 analogies per meeting. All 4 labs used analogy. There were a total of 25, 30, 31, 13 analogies for Labs A, B, C, and D, respectively. Thus, analogies were frequently used at laboratory meetings.

Insert Table 1 about here

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 were coded as being "Within organism," "Other Organism," and "Non-Biological." "Within Organism" analogies are when the base and the target are from within the same organism. Example 1a from Table 1 is a "Within Organism" analogy. In this case the scientist was drawing an analogy between the way the HIV virus works in an InVivo context and how an InVitro HIV could be made by mapping from the InVivo HIV onto the InVitro HIV construct that they were making. "Other Organism" analogies are when the base and the target are from two different organisms. Example 1b from Table 1 is an analogy between the Ebola virus and the Herpes Virus. The scientist points out the differences between Ebola and Herpes to show why Ebola is a better organism to research a particular question. "Non-biological" analogies are when the base is taken from a "Non-biological" domain. In Example 1c, the scientist is highlighting the fact that a finding could be due to chance by drawing an analogy between a monkey typing Shakespeare and the Polymerase Chain Reaction (PCR) generating a chance result. Note that this type of distant analogy is the type of analogy that 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 "Non-Biological" analogies. Thus, the bulk of analogical reasoning was 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 being imposed on the data apriori. From this emerge 4 dominant categories of goals: "Formulate an 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 "explain." Usually the explanations were of methodological issues. There were 21 analogies for "design an experiment," 10 for "fix an experiment," and 23 for "formulate an hypothesis."

Next, we can look at the relation between *goals* and *range*. Table 2 shows 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 we can see that the two "non-biological" analogies were used to make explanations, they were not used to formulate hypotheses. While there were only two "non-biological" analogies in the 16 meetings coded in this paper, we do have two other "non-biological" analogies in our database, all four of

these “Non-biological” or distant analogies were used to explain a concept to members of the laboratory. Thus, “Non-Biological” or distant analogies are rare and generally used for explanations rather than the generation of new hypotheses and concepts.

Turning now to the “Within Organism” and “Other Organism” analogies and goals we can see that 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 be trying to determine the function of a gene in one organism (e.g., a gene in malaria) and draw an analogy to a gene in another organism (e.g., a similar gene in clams). If the scientists know what the gene does in one organism (e.g., in clams), they will 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 “non-biological” or distant domains, the scientists make analogies to other organisms when formulating hypotheses.

Insert Table 2 about here

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 lab or that they have read in journal articles. If we break down the range of the analogies by whether the scientists were recalling specific cases (such as specific experiments that were conducted in the past, or references to particular papers, or experiments conducted by researchers in the field) we can see that 31 of the 40 “within organism” analogies were where the scientist recalled a specific case, whereas for “other organism” analogies only 6 of the 57 analogies were the recall of specific cases. Thus, when scientists are making analogies to the same organism, they tend to recall a specific case. However, when the scientists are making analogies to a different organism, they are not recalling specific cases. We can further consider whether the cases that are recalled are from within the lab or outside the lab; 22 of the 31 “within organism” analogies that were recall of cases were recall 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.

Given that the scientists are using analogies to other organisms without recalling a specific case how do they do this? Psychological research has shown that subjects have great difficulty in going outside their current problem to make an analogy (e.g., Gick & Holyoak, 1983), yet the scientists are able to transcend this problem. How do they do it? An analysis of analogies to other organisms reveals that the scientists have 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 have an abstract knowledge of 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 will now discuss each of these ways of retrieving base analogs in turn. Scientists use homology in the following way. They determine the molecular structure of a gene by sequencing each base pair in the gene. Then, the scientists type in 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, then the scientist will infer that the gene may have that function in their organism. That is, the scientist maps the function of the homologous gene onto the gene that they are investigating.

Thus, homology allows the scientist to both retrieve analogs and to 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 give new hypotheses and new methods that the scientist can use in his or her research. The scientists 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 3 the scientists used homology to infer biological mechanisms and the methods that they should use in their experiments.

Turning now to analogies to other organisms which were not based on homology, the scientists were more likely to use biological mechanisms as a retrieval cue. The scientists might say that E.coli solves this problem 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.

Insert Table 3 about here

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 "non-biological" 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 recent historical analyses 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 have used distant analogies to *explain* a new concept to an audience rather than having a causal role in making a discovery. We are 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 "online" reasoning are easy to forget. In fact, in post lab meeting interviews 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] Furthermore, 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 will be causally involved in making a breakthrough. Again, because no one analogy made a major restructuring 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 the analogies that were involved in making a discovery are laid out, only analogies of very specific types will be seen to have played a major role in scientific reasoning and discovery. Moreover, as

we will see below, analogy is not the only mechanism that comprises conceptual change. In the next three sections you will see 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 will lay out what the complex of mechanisms are 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 conduct an experiment and the results are unexpected. In the psychological literature researchers have investigated this in terms of confirmation biases; subjects tend to seek evidence that is consistent with their hypothesis and ignore evidence that is inconsistent with the current hypothesis. Researchers have repeatedly found evidence that subjects engage in this type of behavior and have argued that scientists have similar reasoning biases (cf. Klayman & Ha, 1987; Mynatt, Tweney, & Doherty, 1982). However, before we accept the generality of the results of these types of experiments it is important to note some of the large differences between the tasks that subjects are given and what scientists do. Most importantly, 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 one task that has been widely used with subjects is the 2-4-6 task (Wason, 1960). In this task, the experimenter asks a subject to determine the rule underlying a sequence of numbers. The subject 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 subject is then told that she can generate other triads and that the experimenter will tell her whether the triad is or is not an example of the rule. Finally the subject is told that when she is sure she should state what the rule is. Many subjects tend to generate triads that are consistent with their hypotheses; they attempt to confirm their hypothesis.[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.

While 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; Kulkarni & Simon, 1987). This view of reasoning is quite different from that implied by the confirmation bias viewpoint. Presumably unexpected findings will be inconsistent with one's current hypothesis, yet if scientists are using the focus on unexpected results heuristic, they should focus on the finding rather than ignore it. Thus, we have two different conceptions of what scientists may do. Of course it may be the case that under certain circumstances the scientists focus on unexpected findings, and under other circumstances they will ignore the findings and behave like the subjects in psychology experiments (cf. Tweney, 1989). The goal of the following analyses was to investigate these questions in a real scientific environment.

Methods

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 they 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 which were consistent with the predictions were coded as "expected." A third category of finding were 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 question was how common were "expected," "unexpected," and "exploratory" findings? In four meetings there were six experiments reported with 70 conditions. There were 22 "expected," l8 "unexpected," and 30 "exploratory" findings. Clearly, unexpected findings are not uncommon. We coded all "expected" and "unexpected" findings as to whether the scientists try to explain away their results or whether they build theories with the findings. To do this, we coded the number of reasoning blocks that the scientists engaged in following an "expected" or an "unexpected" finding. A reasoning block was a group of statements which involve reasoning about a particular finding. One finding can generate many different reasoning blocks and we can use the number of reasoning blocks generated by expected and unexpected findings as an index of how much attention is given to these types of findings. In table 4 the number of reasoning blocks that the presenter generated about expected or unexpected findings is shown as well as whether the presenter tried to explain away or build theories with the unexpected findings. As can be seen from Table 4, there was more reasoning for "unexpected" than "expected" findings. Furthermore, there was little attempt to explain away the "unexpected" findings. Thus, scientists do pay attention to "unexpected" findings.

Insert Table 4 about here

The above analyses applied purely to whether the individual scientist who conducted and presented the research was likely to attend to an "unexpected" finding. We next investigated the question of whether the group also attended to "unexpected" findings. We calculated the number of reasoning blocks that anyone other than the presenter devoted to "unexpected" or "expected" findings. Again we found much more reasoning by the group for "unexpected" than for "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 that there were 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.

A further question that can be asked about the scientists' use of unexpected findings is whether there is any difference between the scientists' treatment of unexpected findings that are consistent with their hypothesis and those that are inconsistent with their hypothesis. An unexpected finding that is consistent with a scientists' hypothesis can occur when the scientist expects a certain type of result to occur, but that the size of the effect is much greater than expected. In this type of situation, the result is consistent with the hypothesis that a certain effect should occur, but is unexpected with respect to the size of the effect. 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 their expectations and that 10 unexpected findings were inconsistent with their expectations. We then coded which findings resulted in the proposal of new hypotheses. We found that 4 out of the 8 consistent findings resulted in new hypotheses and 8 out of l0 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 will have an effect

on whether an unexpected finding is attended to or not. We have found that there are two dimensions of an unexpected finding that will determine whether the unexpected finding is attended to or not. 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 does an unexpected finding occur. What we found is that when the unexpected finding occurs early and it is not a core hypothesis, the scientists ignore it. However if the unexpected finding occurs early and it 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 scientist 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 the one that subjects are in psychology experiments: The subjects are early in the project, and the hypotheses that the subjects have are not core assumptions. Thus, it is only under very restricted circumstances that we see a similarity between the subjects and the scientists.

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 seen and the more likely it is that the scientist has developed strategies or heuristics for dealing with these unexpected findings. Thus, the way that a scientist will deal with unexpected findings will be related to the scientists' having evolved specific strategies after many encounters with unexpected findings.[8] Subjects 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 will eventually attend to unexpected findings.

Distributed Reasoning

Most cognitive research on scientific reasoning deals with individuals reasoning about a problem. However, much of modern-day science is conducted by groups of scientists rather than individuals. Furthermore much of the cognitive work has demonstrated that individual subjects make many different types of reasoning errors. In this section, I will investigate whether reasoning in groups can circumvent certain individual reasoning errors. In particular I will explore the issue of distributed reasoning in science. Distributed reasoning can be defined as different members of a group reasoning about topics such as an hypothesis, experiment, methodology, or interpretation of a result with the different members of the laboratory 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 subjects display in standard psychological experiments.

One of the major problems for both individual subjects in psychology experiments and scientists confronted with new data is to determine what types of inductions to make from new data. There are infinitely may 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. We explored the role that the group as a whole made in the types of inductions that a scientist in the HIV lab made during his talk. At this talk he presented 5 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 post-doc made and modified the post-docs inductions. As can be seen from Table 5, the other members of the lab limited, expanded, replaced and discarded 7 of the inductions that the post doc made.

Insert Table 5 about here

This pattern of challenging inductions was ubiquitous across all labs and provides important information on the role of distributed reasoning. Individual subjects have great difficulties in generating alternate inductions from data, and also have great difficulties in either limiting or 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 will generate different representations. These new representations will make it possible for members of the lab to propose alternate inductions, deductions and causal explanations. Thus, distributed reasoning provides new premises and models that a particular individual might not be able to generate when reasoning alone.

A second question about distributed reasoning is how many people are involved in the reasoning. In the HIV lab we have investigated how many of the inductions and deductions are 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 are shared by more than one individual. We also found that 12% of all inductions and deductions have more than two participants. Furthermore, inductions of one individual sometimes form the basis of a deduction for other individuals.

In Summary, Distributed Reasoning consists of individual 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 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 in the group have at their disposal. The generation of different representations during distributed reasoning helps circumvent one of the major reasoning difficulties that individual subjects have.

The results of these analyses of distributed reasoning are different from the results of brainstorming experiments and creativity in groups experiments. Many studies have shown that when a group of subjects are asked to generate novel concepts, they perform no better than individual subjects. 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 the psychological experiments the subjects are not part of a group sharing common knowledge and values. Usually, subjects in psychology experiments are a random group thrown together for the purpose of the experiment. Second, the types of problems given to the subjects are often arbitrary and require little background knowledge. In the science labs investigated in this chapter, the labs consisted of a group of scientists having overlapping backgrounds and sharing sets of goals about the research. Furthermore, the different members of the lab had slightly different types of knowledge that they can 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 diversities of knowledge. The prediction is that in this type of situation groups of subjects will 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 event in science. Now I will turn to the issue of how all three aspects of scientific reasoning form a complex of mechanisms that together work 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 was when, in the sixteenth century, there was a 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 will turn to now.

Here I provide a dynamic account of a conceptual change that occurred in Lab A and use this example to show how different forms of reasoning worked together to produce entirely new concepts. In order to honor and 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 post-doctoral 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 post-doc 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 post-docs investigated the disease in yet another organism (hamsters) because the lab that the post-doc was in 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 OLA disease, the rabbit CVX disease, and the hamster CVX disease.

One May afternoon the post-doc gave a talk about his latest experiments. He began with a analogies between the human OLA 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 post-doc started to discuss some experiments in which the results were very unusual. The post-doc had conducted a straight-forward 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 were the colmenia B-cells. The post-docs 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 post-doc got to the part of his presentation where he discussed these results. He was surprised and excited by what he had 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 post-doc how it happened. The post doc said he didn't know. The director then made the question more specific. He asked the post-doc what properties were common to the colemia and CVX B-cells that allowed them entry into the heart. The post doc made an analogy to some other experiments that another post doc in the lab had conducted and induced that the CVX and colemia cells were both methylated. The Director and other post-docs 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 that 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 that they had all assumed up until that point. Previously, it had been assumed by that CVX cells only go to the heart and that colemia 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. What happened in the distributed reasoning was that the members of the lab proposed that gaining entry to 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 post-doc drew an analogy back to the human disease and mapped the mechanisms that they had 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 that had been the focus of the research, but had proposed a model that had major ramifications for whole classes of autoimmune diseases. By proposing two new mechanism the lab also had to change a number of other concepts in their knowledge of autoimmune diseases. It was at this point the everyone in the lab realized that a conceptual change had occurred and all shouted in excitement. This was followed by some further analogies in which other post-docs suggested other experiments. Finally a post-doc made an analogy to the to the methods that other researchers have used and the methods that the post-doc had used, explaining why their rivals lab had not made the discovery that 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 in the conceptual change. Second, analogy was a significant component of the conceptual change, but all the analogies that were used were either to the same organism or other organisms: Conceptual change can and does occur without distant analogies. Third, the scientists had little memory for any of the "online" analogies used at the meeting. I asked the post-doc 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, nor that distributed reasoning was involved. Thus, much of the online 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 "online" scientific thinking and reasoning reveal a number of important mechanisms underlying creative cognition. The main idea that I would like to advance 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 view of conceptual change explains why it is so hard to discover what the underpinnings of creativity are. 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 that I would like to address 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 such as the one outlined above. Given that this is the case, then the question arises as to what these scientists were doing that has made them so productive and at the forefront

of their fields? The answer lies in their choice of research topics. Each of the scientists have 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. While 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 "online" 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 and historians 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 where medical doctors are reasoning about patients (Dunbar et al. 1995; Patel, Dunbar, & Kaufman, 1995). We are now starting to investigate whether the same types of reasoning strategies occur in a business context.

The molecular biologists do have some special tools that other scientists and non-scientists do not have, such as being able to use 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 similar structural patterns to the one that 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 when the 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 make one step in drawing analogies easier 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 InVivo research makes it possible to discover fundamental mechanisms of creative cognition and how multiple cognitive processes work together to produce conceptual change. Furthermore, this InVivo approach makes it possible to both discover what aspects of InVitro research are generalizable, and suggests new types of experiments that can be conducted in the cognitive laboratory.

ACKNOWLEDGEMENTS

This research was made possible by a grant from the Spencer Foundation. The research was also supported by grant number 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 would like to thank Laura Ann Petitto for her extensive comments onan earlier draft of this manuscript. I would also like to thank Tom Ward, Keith Holyoak, and Maude A. St. Laurent for their coments on an earlier draft of this chapter. Finally, I would like to thank all the scientists who opened up their laboratories and their minds for this research.

Notes

[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 subjects in psychology experiments.

2 There is some controversy about whether the solar system analogy played a causal role in Rutherford discovering the structure of the atom. Whatever the real case maybe, 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 have drawn up 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 where the source and the taget are distant.

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

[5] Some 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 are pioneering totally new concepts, in an uncharted conceptual space. On this view we would expect to see many distant analogies relative to the other types, which was not the case.

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

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

8 Lovett & Anderson (in press) have shown that history of success has a large effect on what strategy subjects will use. In a number of experiments they have shown that subjects use both their history of success and the current problem solving context to determine the type of problem solving strategy that they use. I would like to argue that scientist use a similar set of heuristics. Whether they will use unexpected findings or not will both depend on the history of success and current context

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, NY: Basic Books, Inc.

Carey, S. (1992). In R.N. Giere (Ed.), Minnesota studies in the Philosophy of Science. Vol XV: Cognitive models of science. Minneapolis: University of Minnesota Press.

Dunbar (In Press). Problem solving. In Bechtel (Ed.) A companion to Cognitive Science. Blackwell Inc.

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

Dunbar, K. (1995). How Scientists really reason: Scientific reasoning in real-world laboratories. In R.J. Sternberg, & J. Davidson (Eds.) *Mechanisms of insight.* Cambridge: MA. MIT press. pp 365-395.

Dunbar, K. (1993). Concept discovery in a scientific domain. *Cognitive Science*, **17**, 397-434

Dunbar, K. & Baker, L. M. (1994). Goals, Analogy, and the Social Constraints of Scientific Discovery. Behavioral and Brain Sciences, 17, 538-539.

Dunbar, K., Patel, V. Baker, L. & Dama, M. (1995). Group reasoning strategies in Knowledge rich domains. Paper presented at the 36th Annual meeting of the Psychonomic Society. Los Angeles. CA. November 10th

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.* Erlbaum: Hillsdale, New Jersey.

Ericsson, K. A., & Simon H. A. (1993). Protocol analysis: Verbal reports as data (Revised Edition). Cambridge, MA: MIT Press.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou and A. Ortony, (Eds.), Similarity and analogical reasoning. New York, NY: Cambridge University Press.

Gentner, D., & Forbus, 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 Psycholog, 14, 1-38.

Holyoak, K.J. (1995). Problem solving. In E.Smith & D.Osherson Eds. An Invitation to Cognitive Science Second Edition. MIT Press.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. Cognitive Science, 13(3), 295-355

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. (l988). The processes of scientific Discovery: The strategy of experimentation. Cognitive Science, l2, l39-l76.

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

Nersessian, N. (l992). 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. Minneapolis: University of Minnesota Press.

Rhodes (l986). The Making of the Atomic Bomb. New York, NY: Simon & Schuster.

Sanderson, P.M., Scott, J.J.P., Johnston, T., Mainzer, J., Watanabe, L.M., and James, J.M. (l994). MacSHAPA and the enterprise of exploratory sequential data analysis (ESDA). International Journal of Human Computer Studies, 4l, 633-68l.

Schunn, C.D., & Dunbar, K. (in press). Priming and awareness in complex reasoning. Memory & Cognition.

Thagard, P. (l992). Conceptual Revolutions. Cambridge, MA: MIT Press.

Tweney, R.D. (l989). Five questions for computationalists. In J. Shrager, P., Langley (Eds.) Computational models of scientific discovery and theory formation. San Mateo, CA: Morgan Kaufmann.

Tweney. R.D., Doherty, M.E., & Mynatt, C.R. (Eds.) (l982). On Scientific Thinking. New York, NY: Columbia University Press.

Wason, P.C. (l960). On the failure to eliminate hypotheses in a conceptual task. Quarterly Journal of Experimental Psychology, l2, l29-l40.

Wiser, M., & Carey, S. (l983). When heat and temperature were on. In D. Gentner, & A.L. Stevens (Eds.) Mental Models. Hillsdale: NJ. LEA.

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

Table 1

Examples of Analogy types used in the laboratories

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 using a substrate that more closely mimics the invivo 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 about 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 puative homology.

1c. Non-Biological or Distant: Monkeys to PCR (Polymerase Chain Reaction) analogy
You know, just because you can see l0 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.

Table 2
Goals that scientists have for within organism, other organism, and non-biological analogies

	Within Organism	Other Organism	Non-Biological
Hypothesis	3	20	0
Design Experiment	9	12	0
Fix Experiment	5	5	0
Explain	23	20	2

Table 3.
Types of knowledge retrieved by analogies based on homology and non-homology for "other organism" analogies

	Homology	Non-Homology
Biological mechanism	10	16
Experimental method	17	7
Problems with methods	4	3

Table 4:
Number of reasoning blocks and type of reasoning about expected and unexpected findings by a presenter

	Expected	Unexpected
Theory build	33	161
Explain away	9	18

Table 5
How distributed reasoning changes the types of inductions made at a meeting in the HIV lab

Limit	3
Expand	1
Replace	2
Discard	1

.