HOW SCIENTISTS REALLY REASON:
SCIENTIFIC REASONING IN REAL-WORLD LABORATORIES

KEVIN DUNBAR

Department of Psychology
McGill University

THIS VERSION COMPLETED FEBRUARY 1993

AS OF JULY 2001: I HAVE MOVED:

Address all correspondence to:
Kevin Dunbar,
Dept. of Psychological & Brain Sciences, Dartmouth College
Hanover, NH 03755

Appeared in R.J. Sternberg, and J. Davidson Eds. (1995). *Mechanisms of Insight*. Cambridge MA.: MIT Press

HOW SCIENTISTS REALLY REASON:
SCIENTIFIC REASONING IN REAL-WORLD LABORATORIES

KEVIN DUNBAR
Department of Psychology
McGill University
1205 Docteur Penfield Avenue
Montreal, Quebec, Canada H3A 1B1

1. Introduction

How do scientists think and reason? What are the psychological processes involved in scientific reasoning and discovery? These questions have been the focus of a large amount of research by cognitive scientists, historians, philosophers, sociologists and psychologists in the past forty years, and is one of the main concerns of this book. Many different approaches have been taken to answer these questions, all with their own vices and virtues. In this chapter, I will discuss two novel approaches to investigating the cognitive processes involved in scientific reasoning and discovery that I have been using in my research. These approaches are making it possible to formulate new models and theories of the cognitive and social mechanisms involved in scientific discovery. The first approach that I will discuss is one of taking a discovery from a real scientific domain, generating a task that is an analog of what the scientists had to do, giving this task to subjects, and then determining whether and how subjects make the discovery (Dunbar, 1989; 1993; Dunbar & Schunn, 1990). Because this approach is based upon a real scientific domain, rather than an arbitrary task that has a tenuous relationship to real science, it is possible to capture important components of scientific reasoning and discovery. The second approach is one of investigating real scientists working on their own research. This approach entailed actually spending an extensive period of time in real scientific laboratories. Data were collected over a one year period in four leading molecular biology laboratories following all aspects of particular scientific research projects including planning of the research, execution of the experiments, evaluation of experimental results, lab meetings, planning of further experiments, public talks, and the writing of journal articles. Some of the research projects resulted in important scientific discoveries, and some did not. This provides a totally novel database with which to address fundamental questions concerning the cognitive processes involved in scientific discovery.

Using terms borrowed from biological research, I will refer to my work on simulated scientific discoveries as "*In Vitro*" research, and my work on scientists' reasoning in real world contexts as "*In Vivo*" research. At the end of this chapter, I will argue that just as in biological research it is necessary to conduct both *In Vitro* and *In Vivo* research to fully understand a biological process, it is likewise necessary to conduct both types of methodologies in cognitive research to fully understand the cognitive processes involved in scientific reasoning and discovery.

2. *In Vitro* research: Simulating the Discovery of Genetic Control

In 1965 Jacques Monod and François Jacob were awarded the Nobel prize for discovering that there are regulator genes that control the activity of other genes. They discovered this by investigating the utilization of energy sources, such as glucose, in *E. coli*. *E. coli* need glucose to live and their most common source of glucose is lactose. When lactose is present, *E. coli* secrete betagalactosidase enzymes that break down lactose into glucose. Betagalactosidase is secreted only when lactose is present. Jacob and Monod discovered that a set of regulator genes inhibit the genes that produce betagalactosidase until betagalactosidase is needed. They proposed that there are two genes I and O that regulate the activity of the betagalactosidase producing genes and that the production of beta-gal is controlled by an *inhibitory regulation* mechanism. As can be seen from Figure 1, when no lactose is present the I gene produces an inhibitor that binds to the O gene, and this prevents the betagalactosidase genes from producing betagalactosidase. When Lactose is present, the inhibitor secreted by the I gene binds to the lactose, rather than the O gene. When this happens, the betagalactosidase genes are no longer inhibited and, consequently, they produce betagalactosidase. When all the lactose is used, the inhibitor again binds to the O gene and production of betagalactosidase stops. Monod and Jacob made this discovery using various mutations where the I, O, and betagalactosidase genes were mutated. Crucially, they initially thought that genetic control was due to genes switching on, or *activating*, other genes. It was only after a large amount of research that they discovered that the mechanism of control was inhibition. Not only was this discovery relevant to production of betagalactosidase, but it was a general model of genetic control that transformed biological research.

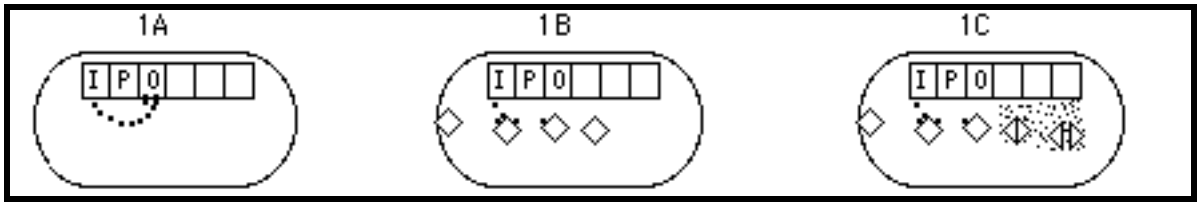

Figure 1. The cycle of inhibitory regulation of genes in *E. coli*. In Figure 1A the *E. coli* is in an inhibited state: The I gene sends an Inhibitor to the O gene, and the inhibitor binds to the O gene, this blocks production of beta-gal from the three beta-gal producing genes (the three unlabeled genes). In Figure 1B, lactose (diamonds) enters the E. coli. The inhibitor binds to the lactose and not the O gene. In Figure 1C, the beta-gal producing genes are no longer inhibited and the beta genes produce beta-gal (small dots). The beta-gal cleaves the lactose into glucose which can then be utilized as an energy source. When all the lactose has been used up the inhibitor binds to the O gene and the beta-gal genes are inhibited from producing beta-gal as in Figure 1A.

The work of Monod and Jacob provides a problem that can be adapted to the cognitive laboratory. A simulated molecular genetics laboratory was designed that made it possible for subjects to propose and test hypotheses about genetic regulation by conducting experiments using various different types of mutants. Two studies were conducted in which subjects were asked to discover how genes control other genes (see Dunbar, 1993a, for the details of these studies). In both studies subjects were taught about *activation* using one set of genes, putting them in a knowledge state similar to the one that Monod and Jacob were in prior to their discovery of inhibitory genetic control (cf. Judson, 1979). Subjects were then asked to discover how another set of genes are controlled. These genes were the I, P, and O genes in which the I and O genes function as *inhibitors*.

In Study 1, subjects had to discover that the I and O genes were *inhibitors*. Given that subjects were taught about activation, as predicted, they all began with one type of hypothesis: that is, the genes are activators that switch on enzyme production. However, subjects found no evidence that was consistent with an activation hypothesis --all experimental results were inconsistent with an activation hypothesis. At this point, the subjects employed one of two strategies for dealing with the inconsistent evidence. One strategy was to continue using their current goal of finding activation. None of the subjects using this strategy succeeded at discovering how the genes are controlled. Other subjects used a second strategy: upon noticing evidence inconsistent with an activation hypothesis, these subjects set a new goal of attempting to explain the cause of the inconsistent findings. These subjects were able to generate a new hypothesis to account for the inconsistent findings (i.e., that the I and O genes inhibit betagalactosidase production). Thus, the results of this first study indicated that subjects set a goal of finding evidence consistent with their initial hypothesis, and that this goal blocked the setting of other goals, such as discovering the cause of unexpected findings.

In Study 2, the hypothesis that maintaining one goal blocks the setting of other goals was tested. In this study, the genetic mechanism was changed so that one gene worked as an activator and the other two genes as inhibitors. In this situation, it was predicted that subjects would first set out to achieve their goal of discovering activation and then, after having achieved this goal, they would set a new goal of accounting for the findings that were inconsistent with an activation hypothesis. Once this new goal was set, subjects would be able to generate an inhibitory hypothesis. This was exactly what happened, twice as many subjects proposed inhibition than in Study 1, and more subjects reached the correct conclusion. These findings supported the hypothesis that subjects' goals determine when and how inconsistent evidence is used.

The results of these studies shed new light on a number of aspects of scientific reasoning: All subjects used inconsistent evidence to modify their hypotheses. However, subjects use of inconsistent evidence was contingent upon their current goal. Thus, the goal that subjects set was pivotal to making a discovery. When subjects maintained their initial goal they did not make a discovery. When subjects changed their goal to one of exploring the cause of unexpected and/or inconsistent findings, they then made the discovery.

The results of these studies indicate that it is possible to discover important components of scientific reasoning by taking a real scientific discovery and bringing it into a cognitive laboratory. While these types of experiments are generating new insights, it is not yet possible to determine the effects of the social context of science on the discovery process, or whether the scientific reasoning strategies that non-scientists use are used by scientists and vice versa. To achieve a more complete understanding of the specific factors that underlie scientific reasoning and discovery, other research methods are needed. In the next section, I will outline another research strategy that I have been using to investigate scientific reasoning and discovery in the real world.

3. *InVivo* Research: Real World Study of Scientists Reasoning

While research on individual subjects has produced many rich and important theories of reasoning in general and some of the components of scientific reasoning in particular, there are several distinct problems with making generalizations from experiments on individual subjects to the way in whch scientists reason. First, science takes place in a social context: groups of scientists work on a problem in a laboratory, rather than one scientist working alone. So far, cognitive psychologists have tended to investigate scientific reasoning in individuals and have ignored the social context of science. Second, psychologists have used tasks that are not 'real' scientific problems (e.g., discover an arbitrary rule: Klayman & Ha, 1987, Mynatt, Doherty, &

Tweney, 1978). Third, the subjects that psychologists use are generally non-scientists (e.g., Klahr & Dunbar, 1988). Clearly, scientists working on real scientific problems need to be studied as well. Unfortunately, when scientists have been studied they were given the same simple and arbitrary concepts that non-scientists are given (e.g., Mahoney & DeMonbruen, 1977). Fourth, subjects in psychology lab experiments work on problems that may last for as little as ten minutes and involve no extensive knowledge of a scientific topic (e.g., Klayman & Ha, 1987). In scientific research, a particular problem may take months, years, or decades to solve and the scientists have extensive knowledge of a domain.

A number of cognitive researchers have noted the limitations of the types of studies referred to above and have turned to historical data on particular scientific discoveries to provide a richer account of the scientific discovery process. Researchers using historical data have analyzed historical accounts of scientific discoveries to uncover the mechanisms involved in scientific reasoning. For example, Nersessian (1992) and others (e.g., Gooding, 1992; Holmes, 1985; Tweney, 1985) have conducted detailed analyses of diaries and notebooks that make it possible to infer some of the cognitive processes involved in particular scientific discoveries. This approach has yielded rich and important accounts of some of the cognitive components of a particular discovery. However, this method also has its drawbacks. The main limitation being that only *indirect* and selective access to the cognitive processes underlying scientists' discoveries can be obtained.

Another historical method for determining the psychological processes involved in scientific discovery that has been used is to interview scientists who have made a discovery (e.g., Giere, 1988; Karp, 1989; Mitroff, 1974). Thus, rather than relying on laboratory notebooks, researchers can interview scientists and ask about how a discovery was made. There are a number of cognitive accounts that have used this method and have provide detailed accounts of particular discoveries. For example, Karp (1989), performed a series of extensive interviews with the scientist who discovered a new mechanism of genetic control, and built a computational model of the cognitive processes that were involved in the discovery. While this is clearly a useful approach, retrospective reports are notoriously unreliable (cf. Ericsson & Simon 1982, Nisbett & Wilson, 1977). Furthermore, research from my laboratory has shown that subjects are often unaware of what leads them to make a discovery (Dunbar & Schunn, 1990). Dunbar & Schunn (1990) found that solving one problem improved performance on an analogically similar problem, yet the subjects did not report using any information from the first problem to solve the second problem: subjects did not mention the first problem either in their verbal reports while solving the second problem, or in their retrospective reports. Thus, to uncover the strategies that scientists use, retrospective reports cannot be relied upon.

A third approach to uncovering important aspects of scientific research has been the contemporary sociological approach. A number of sociologists have investigated scientists working in laboratories. These researchers have used ethnomehodological approaches, or interviews with the scientists themselves (e.g., Fujimura, 1987; Knorr-Cetina, 1983; Latour & Woolgar, 1986; Mulkay & Gilbert, 1983). While these studies have uncovered important components of the day-to-day workings of scientific laboratories they have not been concerned with uncovering the cognitive processes that are used by scientists in their day to day research. These researchers have stressed the importance of the social context of science, and have demonstrated that the social context of science has an effect on all aspects of the scientific process. However, while these studies established that social context is somehow important, exactly how the social context impacts on the scientist's knowledge remains unanswered.

To summarize, the research from my laboratory, and that of other researchers, suggests that a number of basic cognitive heuristics and operations form the foundation of scientific reasoning. However, no Cognitive Scientists have actually investigated real scientists conducting their day to day research. That is, there have been no systematic cognitive investigations of how scientists reason while conducting their research. While the standard cognitive and historical

analyses have provided rich and important accounts of the cognitive processes involved in particular discoveries, there are many crucial aspects of the scientific discovery process that it is not possible to gain access to using these methodologies. In particular, the online cognitive processes, and the social interactions that are involved in a particular discovery are not directly accessible. This suggests that alternate methodologies need to be adopted to uncover the online processes that particular scientists use. Let us now turn to a new type of cognitive research that investigates these questions.

4. How Scientists Really Think

I will now turn to a study in which I collected data on the reasoning processes and discovery heuristics that scientists used in four of the world's leading molecular biology laboratories at a major US university. The overall goals of this research were (1) to determine what types of reasoning heuristics scientists use to propose experiments, generate hypotheses, and evaluate results, (2) to determine how scientists represent their knowledge of the research projects that they are working on, (3) to uncover the cognitive processes that lead to changes in scientists' representation of their research (that is, to investigate the mechanisms involved in conceptual change and insight), (4) to discover the cognitive mechanisms that groups of scientists --rather than an individual scientist-- use to formulate experiments and hypotheses, (5) to discover whether the social context of scientific work can counteract the well-known faulty heuristics that individuals have been shown to use when reasoning scientifically, and (6) to discover whether and what the mechanisms are for the social context to influence conceptual change.

5. Method

5.1. SELECTION OF LABORATORIES

Six laboratories were identified on the basis of (1) the quality of their publications, (2) the type of research that they were conducting, (3) the fact that each laboratory had previously made discoveries that the scientific community regarded as being significant, (4) the laboratories were of different sizes, and (5) the directors of the laboratories had differing amounts of research experience.[1]

All of the six laboratories allowed me to investigate them. Four of the laboratories were judged to be most suitable and were subsequently investigated. These four laboratories varied along two dimensions. First, the laboratories were either developmental biology labs, or worked with pathogens (disease causing viruses and bacteria). Second, the laboratories were either focused at the cellular level, or at a molecular level. By selecting laboratories in this manner, it was possible to identify which aspects of the research are general, and therefore used by all 4 laboratories, and which strategies were specific to a particular field such as developmental biology or molecular biology.

For the purposes of maintaining confidentiality, the names of the scientists will not be revealed. The Laboratories will be labeled A, B, C, and D. All the scientists requested anonymity and that the results of their experiments not be divulged. To further maintain confidentiality of the data that I have obtained, many of the scientific details of the discoveries made and research projects investigated have been omitted from this paper. While the scientists did request anonymity it is important to note that all the scientists allowed free access to their laboratories, to interview anyone in the laboratory, attend any meeting, read and keep copies of their grant proposals (including the pink sheets), attend their talks and lectures, and read drafts of their papers. Thus, there was complete access to the day to day activities of laboratories. In addition,

the laboratories were so cooperative that they frequently phoned 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.

Table 1: Research areas of the four laboratories investigated

	Cell biology	Molecular
Developmental	Lab A	Lab B
Pathogens	Lab B &C	Lab D

5.1.1 Laboratory A. This laboratory is run by a senior researcher who has over 300 publications, won numerous awards, has former students who are also leading researchers in the field, and has made a number of extremely important findings that have revolutionized his field. His laboratory consisted of 22 post-doctoral fellows, 5 graduate students and 4 technicians. The director suggested following a number of research projects that he thought might lead to interesting discoveries and 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.[2] One of the other two research projects was unsuccessful, and the other research project had not progressed significantly within the eight month period.

5.1.2 Laboratory B. This laboratory is run by a senior researcher 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 followed one of the research projects that was being conducted in his laboratory (it was the only research project that was just starting). This research project has been beset by a number of problems that have meant that the researchers have made only a small amount of progress.

5.1.3 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 by studying an organism that has very unusual biological properties. He has over 60 publications and his work on RNA is regarded as seminal. 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.

5.1.4 Laboratory D. This laboratory is run by an assistant professor who is already famous for his work on viral mechanisms. He has invented a number of widely referenced techniques, is regarded as conducting some of the most innovative work on HIV (Human Immunodeficiency Virus). 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. I followed three research projects on HIV activity. These three 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 B also invented a new genetic technique. This technique is likely to end up being one of the most important inventions in the last 10 years in molecular biology and genetics.

5.2. SELECTION OF RESEARCH PROJECTS FOR INVESTIGATION

Within each laboratory particular research projects were selected for study on the basis of (a) an interview with the professor (i.e., laboratory director) about the research that was going on in his laboratory, and (b) whether the research projects had just started, or were about to begin. By selecting new research projects it was possible to investigate the cognitive components from the beginning of a scientific research project. Once the research projects were selected, I then met with the post-docs, graduate students, and technicians that were conducting the research. All the members of the four laboratories were willing to cooperate. In laboratories A, C, and D four research projects were pursued. In laboratory B, one research project was pursued, as this was the only project that was beginning.

5.3. DATA COLLECTION PROCEDURE

A pre-present-post design was used in which data were collected prior to a lab meeting (pre), during a lab meeting (present), and after the lab meeting (post). This design is similar to the pretest-posttest design used in experimental research (cf. Campbell & Stanley, 1963). The "pre-lab" meeting component consisted of an extensive initial interview in which the researcher provided background information on their research project and the rationale for conducting their research. That is, the researcher stated the theories, hypotheses, predictions, experimental results, current knowledge in the field, rival theories, relation to other research projects in the lab, and problems with the research. In addition, one or two days before a researcher was supposed to give a laboratory presentation about their research, an interview was conducted in which the researcher was asked to (i) state what research they had done, (ii) state why the researcher conducted their experiments, (iii) state the specific research question, goals, experimental design and predictions and why they did not conduct other types of possible experiments, (iv) state what their results were and any problems that occurred in conducting the experiments, (v) state what the researcher thought the experimental results meant, and (vi) state what directions the research project was going to go into next (that is, what experiments would be conducted next).

The "present" component of the procedure consisted of either video or audio taping a laboratory meeting. Notes were kept of contextual information not readily apparent in the audio or video tapes. The "post-lab" meeting component of the procedure consisted of an interview with a researcher one or two days after the laboratory meeting to ask the researcher what they were now doing, and whether the meeting had changed their plans. The same six sets of questions that were asked in the pre-lab meeting component were again asked in the post lab meeting component. This made it possible to determine the effects of the lab meeting on the researcher's representation of the research and on plans for future experiments. This pre-present-post design was repeated at least three times over an eight month period for all research projects.

By comparing the data gathered using the pre-present-post design it is possible to determine the effects of the meetings on scientists' reasoning, and on their research. Each

research project was followed for an eight month period, in which a cycle of a "pre-lab" meeting interview, tape of lab meeting (i.e., present component), and a post-lab meeting interview were conducted at least three times. All interviews and laboratory meetings were audio tape recorded and extensive notes were kept; these make it possible to understand contextually relevant information. During the last two months of the research a number of laboratory meetings were videotaped. This made it possible to get a visual representation of the data and data analysis techniques that the scientists were using as well as the social and situational factors not readily apparent in the audio tapes.

5.4. DATA ANALYSIS

5.4.1 Transcription. All data collected are transcribed and coded (i.e., audio tapes, videotapes, including notes from grants and pink sheets, drafts of papers with comments, and other relevant materials). Transcriptions are made with two independent transcribers with a background in molecular biology.

5.4.2 Coding. Following transcription, the data are coded along a set of dimensions derived from Brutlag, Galper, and Milis (1991), Dunbar (1993a), Klahr and Dunbar (1988), Stein (1992), and Ericsson and Simon (1984). The coding schemes provide converging evidence on the cognitive operations, mental representations, and social interactions that the scientists used. Once the data are coded, they are entered into a computerized database (Sanderson, Scott, Johnston, Mainzer, Watanabe, & James, 1993) with relational search capabilities that makes it possible to answer specific questions about the scientists' thinking and reasoning.

In order to give a flavor of the types of attributes that are coded, a partial listing is provided within each category below. However, the existing coding schemes are far richer than that which can be discussed here. The three major categories of knowledge that these coding schemes allow me to specify are as follows:

> *Coding of the scientists' representation of their research over time.* We are using Brutlag's 1991 scheme (Brutlag, Galper & Milis, 1991) which provides a list of attributes for molecular biological knowledge and experiments. This scheme was developed by a molecular biologist who is building computational models of molecular biological knowledge. We have adapted this scheme as a coding device that specifies the features of the scientists' representation of their knowledge. The scheme specifies the attributes of knowledge relevant to understanding DNA metabolism such as; the structure of DNA, strands, nicks, activity, specificity, activity, temperatures and pH values of reactions. The coding scheme integrates these attributes into an overall model of knowledge and experiments. This scheme makes it possible to represent the molecular biologists' knowledge, and how this knowledge changes over time. We are using another coding scheme for cellular biological knowledge.
>
> *Coding of group interactions.* We are using a coding scheme derived from work on discourse analysis and conceptual change. This coding scheme classifies the types of interactions between speakers (e.g., clarification, agreement and elaboration, disagreement, and questioning), the goals of the speaker, and the current representation of the knowledge. This makes it possible to chart the effects of the interactions on the speakers' current representation of the research project.. The coding scheme makes it possible to identify whether and when social interactions lead to conceptual changes. Using this scheme we can identify the specific types of social interactions, and the various combinations of factors that must be present for conceptual change to occur. In

addition, this coding scheme makes it possible to make predictions about whether the interaction will lead to a change in the speaker's representation and what the speaker will do.

Coding the scientists' cognitive operations. All data are being coded using standard protocol analysis techniques (cf. Ericsson and Simon, 1984; Newell & Simon, 1972) that I have used previously (Dunbar, 1993a, Klahr & Dunbar, 1988; Klahr, Dunbar, & Fay, 1989). First, a task analysis is being conducted for each research project. This task analysis determines the current state of knowledge, the goal state, and the series of cognitive operations that the scientists apply to get from their current state of knowledge to their desired state. The second step is to code the data in terms of the cognitive operations identified in the task analysis. The third step is to formulate a model of how the scientists actually combine these cognitive operations into heuristics that guide their research. This third step necessitates bringing together the coding of the scientists representation of their research, the coding of the group interactions, and the coding of the scientists' cognitive operations into one overall scheme.

6. Overall Results

A select sample of the analyses conducted on the present data are provided here (see Dunbar 1993b, c and Dunbar and Baker 1993a, b for the complete analyses). There was a large intra and inter laboratory similarity of the mental representations, experimental heuristics, and problem solving heuristics that all four laboratories used. Indeed, the analyses reveal that the basic components of the scientists' cognitive operations are surprisingly similar and differ largely in the way that these operations are combined. This high degree of regularity in the data makes it possible to apply rigorous data analysis techniques to the data and draw highly generalizable conclusions about scientific reasoning. A number of trends are emerging from the data: First, scientists make extensive use of negative evidence to discard their hypotheses. Second, the use of local analogies where knowledge is imported from the same scientific domain is a common mechanism of conceptual change. By contrast, distant analogies were used to highlight salient features of the problem that they were discussing. Third, the social context of the research produces significant changes in the representation of the problem and modulation of individual reasoning biases. We have been able to identify the particular types of social interactions and cognitive states that are present when conceptual change occurs. Overall, these results reveal that both domain specific knowledge and the social context of scientific research prevents scientists from making many of the reasoning errors identified in individual subjects in cognitive psychology laboratories.

7. Mechanisms Underlying Conceptual Change and Insight

The circumstances under which conceptual change and insights occurred will be addressed. Conceptual change and insight occurred in the face of inconsistent experimental findings, as a result of the use of analogy, in the context of group discussions, and as a consequence of surprising findings. Each of these sources will now be considered.

7.1. INCONSISTENT RESULTS AND CONCEPTUAL CHANGE

Surprisingly, results inconsistent with the scientists' current hypothesis quickly led to the discarding of hypotheses. The discarding of an hypothesis on the basis of inconsistent evidence occurred under very specific circumstances. First, inconsistent evidence tended to be used to change specific features of an hypothesis, while the overall type of hypothesis remained the same. For example, a scientist changed his hypothesis from "this particular sequence is necessary to initiate binding of the protein" to "any sequence in this region that has a base-pair mismatch will be bound to by this protein." Note that in this situation, the conceptual change that occurred was quite minimal. This type of conceptual change displayed the usual generalization, specialization heuristics that have been identified in previous work on reasoning, such as the findings obtained in my *In Vitro* work on scientific reasoning discussed in the first section of this chapter. The second type of use of inconsistent evidence was more interesting. In this case, the evidence was not only inconsistent with the current hypothesis, but was also inconsistent any hypothesis of that type, and the scientist needed to invent a totally new type of hypothesis (or concept, or frame depending on your terminology) to account for the data. This type of conceptual change rarely occurred within an individual. As in laboratory studies of cognition, individual scientists out of a group context usually attributed inconsistent evidence to error of some sort, and hoped that the finding would go away. However, when the finding was presented at a laboratory meeting, the other scientists tended to focus on the inconsistency to dissect it, and either (a) suggested alternate hypotheses, or (b) forced the scientists to think of a new hypothesis. This happened at numerous lab meetings and was one of the main mechanisms for inducing conceptual change in scientists when inconsistent evidence occurred. Often this resulted in the phenomenological experience of insight in which the scientist exclaims that they now know what was going on in their experiment. As we will see in the section on social interactions, the particular mechanics of the interaction are crucial to whether conceptual change did, or did not occur.

The way in which inconsistent evidence was treated also varied as a function of experience. Less experienced scientists were more willing to maintain a hypothesis than more experienced scientists. However, while the more experienced scientists showed mush less confirmation bias than the less experienced researchers, they often displayed what we term a "falsification bias": often they discard good data that actually confirms their hypothesis. This falsification bias appears to be the result of much experience with the negative experience of being proved wrong. We are currently simulating this falsification bias in an experiment in our laboratory (Baker & Dunbar, 1993a). These findings indicate that a crucial factor in determining whether people will maintain an hypothesis in the face of inconsistent evidence is domain-specific knowledge, rather than a reasoning bias, *per se*.

7.2. ANALOGY AND CONCEPTUAL CHANGE

We are currently coding all the uses of similarity in the corpus. We have coded all instances of where a scientists notes that something is similar, or different, from something else. We can then look at instances of analogical reasoning. A preliminary analysis of the data indicates that analogies were an important source of knowledge and conceptual change. In three of the four laboratories analogies were frequently used, ranging from 4 to 22 in any meeting. Three different classes of analogies were used. First, analogies from the same domain, in which the scientist drew an analogy from a previous experiment to their current experiment (*Local Analogies*). Second, a whole system of relationships from a similar domain was mapped onto the domain that the laboratory was working on (*Regional Analogies*). Third, a concept is mapped from a very different domain to the domain that the scientists are working on (*Long-Distance Analogies*). These different types of analogies are used under different circumstances.[3]

7.2.1 Local Analogies. This type of analogy was very frequently used. Local analogies were usually used when the experiment that a researcher was working on had problems and was not working. The researcher made an analogy to an experiment in a very similar research area, or to a similar research technique or protocol. The actual analogical mapping that occurred was to map the unsuccessful problem that they were working with to another similar experiment that was successful. The scientist would then determine what the difference was between the successful and unsuccessful experiments and substitute the different components from the successful approach into their unsuccessful approach. For example, at one meeting a scientist was having difficulty in purifying a protein and said:

> "so I had to pursue another method that would solubilize the proteins, but would also stick to the beads, and basically, this is a method by James Digby and it's also, this method is also a similar method found in Maniatis. Basically, the key step is the 8 Molar urea step. Which just, it just solubilizes everything. But anyway, this is a protocol; it basically was just followed exactly since this worked for someone else, I figured it might work for me, too."

This use of local analogies does not immediately appear to be a very sophisticated type of analogical reasoning, and certainly not the type of reasoning that has been the focus of much cognitive research. However, the use of local analogies is one of the main mechanisms for driving research forward. In the field of molecular biology, at least 60% of the experiments have technical problems that need to be resolved and local analogical reasoning is one of the main methods that the scientists used when they had problems with their experiments. This type of analogical reasoning occurred in virtually every meeting, and often numerous times in a meeting. New knowledge is added to their representation by making the analogy, and this drives the research forward.

7.2.2 *Regional Analogies*. In regional analogies the scientists mapped over entire systems of relationships from one domain to another and the two domains were different classes that shared a common superordinate category membership (e.g., both phage viruses and retroviruses were mapped onto each other and clearly both are members of the superordinate category virus). This type of analogy was not frequent, but did occur from time to time. It rarely occurred when scientists were having a problem with an experiment. Instead, this type of analogical reasoning occurred when the scientists were working on both elaborating their theory, and planning new sets of experiments. For example, one laboratory held a meeting that drew an analogy between one class of virus and another. While a considerable amount is known about certain types of viruses, little is known about many basic components of retroviruses. Furthermore, because retroviruses are considered very different from other types of viruses researchers rarely use knowledge of one to inform their research about the other. What this laboratory did was to try and map knowledge over from one class of virus to retroviruses, the goal being to (a) use this knowledge to fill in gaps in their own knowledge, by drawing sets of 1 to 1 mappings, and (b) to suggest new questions to ask about retroviruses. Thus, mapping over an entire system of relations was a very powerful tool. The finding that this was a rarely used type of analogy would be consistent with much psychological work on analogy, but the reasons may be quite different. In this case, and in the other cases in this corpus, this type of analogy tended to be used only after the scientists had already started to formulate a model of the entire process that they were investigating. Thus, the scientists then had a system of relations and mechanisms in their own domain that they could then map to another domain. Until they had built such a representation of their own domain, it would not have been possible to map over the other domain. We are currently conducting an experiment to test the hypothesis that analogical

mapping of sets of relations is most likely to occur, and lead to conceptual change, when subjects have built up a fairly detailed representation of the target domain.

7.2.3 *Long Distance Analogies*. Long Distance analogies were used, but not frequently. They were never used to solve experimental problems, or in model building. Rather, long distance analogies were used to highlight features of the research that were salient, and were usually used to bring home a point, or to educate new members of a laboratory. Thus, while distant analogies were used to change the representation of knowledge in people, it was not a driving force in making any of the discoveries observed over the year. This use of analogies often led to significant insights in the other members of the laboratory, making it clear exactly what the point was. One example of a highlighting use of a distant analogy was:

> "*Postdoc*: what goes on in the flagellar pocket is a real big question right now, and there's not much known about it. It's a very specialized domain of the plasma membrane, and it has very specialized function. What's in the flagellar pocket and what goes on in the flagellar pocket is uh, not been studied in any great depth or detail. An interesting question. Ok.
> *Professor:* It's sort of semi-closed. It's open to big molecules like LDL gets in--
> *Postdoc*: Things get in, but things don't... It's like the Hotel California - you can check in, but you can never check out."

It is important to note that this use of Long Distance analogies is quite different from that proposed by other researchers. A number of researchers have argued that many of analogies that scientists use in their publications or talks were actually causal in making particular discoveries. That is, scientists first make the analogy and then map features of the analogy over to the problem that they are investigating and make the discovery. In the corpus of data that we collected we did not find one instance of a case where a long-distance analogy led to any conceptual changes or insights on the part of a researcher. Instead, the long-distance analogies were used to highlight features of a point that a scientist makes. We are now monitoring the publications that the researchers are writing to see if long-distance analogies creep into the publications, but were not present in the discovery. If this is the case, then this would suggest that at least some of the analogies that scientists have used in their publications were not causal in making a discovery, but were added when writing up the research. Thus, the importance of long-distance analogies and their causal roles in making discoveries may have been overemphasized by some researchers.

7.3. ANALOGY USE AND SOCIAL STRUCTURE

While use of analogy was a common occurrence in the laboratories, analogical reasoning did not occur in one of the laboratories. Two questions immediately arise here. One is whether the lack of analogical reasoning had a detrimental effect on conceptual change, and the other is, why did this laboratory fail to engage in analogical reasoning? The single laboratory that did not engage in analogical reasoning did not make any real gains in their understanding of the genes that they were working on. Recall that the most common use of analogies in the laboratories was when an experiment did not work. In this situation scientists drew analogies to other experiments in an attempt to solve their problem. However in the laboratory that did not make analogies, the scientists used a different strategy when they encountered problems in their research; they manipulated experimental variables such as raising the temperature, varying chemical concentrations, and so forth, to make things work. Thus, a problem that could have been solved by making an analogy to another similar experiment (Local analogy), or to another

organism (Regional analogy) was not made --leaving some problems unsolved and others lingering for months to solve. Indeed, very similar research problems were encountered in the other laboratories, but they were solved much faster through the use of Local and Regional analogies. This finding is consistent with the hypothesis that Local and Regional analogies are a potent source of conceptual change.

Why were the members of the laboratory not making use of analogy? One aspect of the laboratory appears critical to whether analogies will be used. It is the social structure of the laboratory. All the members of this laboratory had come from highly similar backgrounds, and consequently drew from a similar knowledge base. In the other laboratories, the scientists came from widely differing backgrounds, and these different sources of knowledge were important components in the construction of analogical mappings. When all the members of the laboratory have the same knowledge at their disposal, then when a problem arises, a group of similar minded individuals will not provide more information to make analogies than a single individual.

The finding that the social structure of the laboratory has an effect on types of reasoning and conceptual change may explain why many experimental studies of reasoning by groups produce no better performance than individuals alone. In these studies, the groups of subjects are generally homogeneous with respect to background, and according to the mechanisms of conceptual change that I am invoking, should not produce conceptual change. We are currently conducting a number of experiments to test this hypothesis. These results go beyond merely stating that social structure is important. These findings indicate the groups of individuals must have different pools of knowledge to draw from to make fruitful analogies. Merely having a group of scientists working on a particular problem (i.e., social structure) will not result in the use of analogies.

7.4. ANALOGY USE AND EXPERTISE

As the above section on analogy use and social structure indicates, one of the key components in analogical reasoning is the knowledge that the laboratory has access to. Not only is the knowledge that the group has of importance, but the knowledge that an individual scientist has is central as well. We have found that the more expert the scientist is, the more analogies the scientist will make, the more similarities that he or she will note, and consequently, the more overall research success they will have. While the experts clearly have more knowledge at their disposal, they also have knowledge organized and represented in different ways from the less expert scientists. This is evident in the group interactions of the scientists with each other. An expert scientist tends to see many of the deep structural features as being very obvious and treats them almost like surface features. The novice scientists have great difficulty seeing the deep structural features. When it comes to making productive analogies it is much easier for the expert scientist: for them, the deep features are obvious and they can readily map these features onto other domains. For the novice scientists the deep features are not obvious and therefore the mappings onto other domains are difficult and non-obvious. Thus, experts make both more analogies and more productive analogies.

7.5. SOCIAL CONTEXT AND CONCEPTUAL CHANGE

In the previous sections we have seen that social factors have an effect on scientists' interpretation of inconsistent information, and on their ability to formulate and use analogies. The goal of the present analyses of social interactions is not to restate the obvious: that social interactions are important. Instead, our goal is to identify the precise mechanisms by which groups of scientist change each other's representation of knowledge. In our analysis of the laboratory meetings we are beginning to uncover a number of social mechanisms that facilitate conceptual change. To address this issue, we are analyzing sets of instances where conceptual change occurred and did not occur. Specifically, we are investigating whether conceptual change did or did not occur following (a) questions that force the scientist to think about their work at a different conceptual level, or with a different goal, (b) when the scientist was asked to engage in deductive, causal, or inductive reasoning, (c) when the researcher was asked to give more details of a particular aspect of their theory or data, (d) when the researcher's theory or data was challenged by another member of the laboratory group, and (e) questions from different types of people such as research assistants, graduate students, postdocs, professors, or a Nobel prize winner! We are only just beginning to obtain answers to these questions from our database. Our database allows us to identify particular patterns of social interactions, as well as prior knowledge states which result in conceptual change.

Analyses of the data reveal that question answering is a potent mechanism of inducing conceptual change in scientists. One type of question that produced a number of small conceptual changes in all the laboratories, and that fostered a major scientific discovery in one scientist, was to ask a question that forced the scientist to change from thinking about the research at one level to thinking about the research at another level. For example, a scientist may be conducting a series of experiments aimed at discovering the mechanism by which a certain type of lymphocyte binds to a certain type of cell. The scientist is concerned with the experimental details and particular components of the mechanism. Other scientists may ask this scientist a question about how the lymphocyte got there in the first place, rather than how does it bind. This new question forces the scientist to reorganize his knowledge, and when he does this, his original question is also answered. Thus members of a group can get a scientist to adopt new perspectives and goals that can result in a reorganization of knowledge and result in a scientific discovery. We are currently analyzing a scientific discovery that occurred under this type of questioning (Dunbar and Baker, 1993a, b).

Because we have many instances of cases in which conceptual change did, and did not occur during lab meetings, we have been able to uncover the mechanisms by which social interactions and cognitive representations interact to produce conceptual change. An analysis of the interactions surrounding the making of discoveries indicates that sequences of specific types of interactions and knowledge states occur. In particular, we have identified that when (i) surprising findings occur, (ii) the researcher believes that these findings are not due to error, and (iii) other members of the group challenge the researcher's interpretation of the findings, significant conceptual change will occur. The challenges force the scientist to look at the data with different questions and goals, thereby changing the scientist's representation of the findings. When the researcher believes that the findings are due to error, no amount of challenging, or suggestion of other explanations will result in conceptual change. We have a number of meetings, as well as professor-researcher interactions, in which no conceptual change occurred for these reasons.

Another form of questioning is one that triggers a chain of reasoning that can then result in a reconceptualization of a theory, data, or experimental design. Often the question is asked when the speaker has left out some details that they were unsure about. The speaker then engages in, for example, deductive reasoning, and often other members of the laboratory will also engage in this reasoning process, resulting in a very different conception of a problem. This often occurred when the scientist had a problem with his or her experiment. If analogical mapping did not achieve a solution, the members of the laboratory would attempt to deduce

what the source of the problem was, and then suggest a solution, thereby changing the representation of the problem. Thus certain types of social interactions that occur in a laboratory meeting have specific effects on the types of reasoning strategies that scientists use. We are now beginning to identify which combinations of social interactions and cognitive states lead to which types of reasoning.

One of the laboratories engaged in extensive group problem solving whenever a problem arose in the research. Many members of the laboratory reasoned about the research and often the results of one person's reasoning became the input to another person's reasoning. This resulted in rapid reconceptualization of problems and to significant changes in all aspects of the way the research was conducted. Situations in which group problem solving occur provide a rich example of the way that cognitive and social mechanisms interact. We have found that subgroups focus on particular features of the problem, change these features and then pass on their part of the solution to another member of the group. The researcher presenting then picks up the proposed solutions and integrates them into his or her conceptual framework, and then the group goes through another round of problem solving.

7.6. ON SERENDIPITY

A common event in all research is the presence of surprising results. Often unusual results are of no interest, other times the following up of surprising results leads to significant discoveries. Scientists frequently allude to discoveries based upon surprising findings as being due to serendipity. An example of the view that serendipity is the source of many discoveries appeared in a recent issue of the journal *Science*, where a reporter discussed a particular scientific discovery in the following way. "As with many surprising discoveries, the finding that DNA injection could get the cells of living animals to produce proteins came serendipitously" (Science, March 19, 1993, p 1691). The particular scientists in question discovered that their control condition had a much better effect than any of their experimental conditions. They then focused on the control condition and discovered a new mechanism to introduce foreign DNA into a host. While many scientists and journalists may regard certain scientific discoveries as serendipitous, the data that we have collected indicates that many findings that scientists might call serendipitous are not so. The so-called serendipitous findings are the result of careful experimentation and planning that are designed to expose novel mechanisms.

We have found that experimental results in which the control condition produces unusual results is very common and was the source of many discoveries in our corpus. One of the scientific discoveries that we recorded occurred when a control condition produced surprising results. Other discoveries in the laboratories that we have been investigating also occurred when control conditions produced surprising results, and the surprising results were followed up. Furthermore during my initial interview with the director of Laboratory A he said that one of the most important strategies that he uses in his research is to follow up surprising results. In the lab meetings he used this strategy numerous times, forcing the other scientists to focus on surprising findings, particularly when they involved the control condition, and, consequently, to gain new insights into their research.

The standard explanation for using a control condition is that it allows the scientist to determine whether the effect observed with the experimental condition is really due to the experimental manipulation, or is due to other factors. A control condition is regarded as a check on the experimental conditions. The finding that, in the laboratories that I have investigated, control conditions often generate surprising results leads us to conclude that the manner in which scientists choose control conditions and the way in which the results of controls are interpreted are crucial to understanding scientific discovery and insight. When a scientist selects a control for an experiment many factors have to be taken into account and researchers often

use more than one control. The control conditions serve functions other than checking to see that the experimental effects are real. Even when a scientist is correct that the experimental conditions should produce such and such an effect, other mechanisms may be involved, or more important mechanisms may be involved that the control condition can uncover.

This type of analysis suggests that far from being the result of serendipity, the use of surprising findings, (e.g., in a control condition), makes it possible to uncover hidden mechanisms in an orderly manner. The scientist carefully constructs controls that serve the function of both checking the experimental conditions and exposing hidden mechanisms, should they be there. In the data that we have collected the scientist is usually looking for the desired results in the experimental conditions and to do this the scientist has formulated a rich set of hypotheses and mechanisms that could account for a wide variety of possible findings. When the control conditions produce unusual results the scientist is already considering a host of potential mechanisms and thus a surprising finding allows the scientist to focus on the aspects of his or her current conceptual structure that need to be changed or rejected. The surprising finding is genuinely surprising, but the use of controls and an already richly articulated conceptual structure makes it possible to make sense of the findings and propose novel theories. The manner in which experiments are constructed minimizes the role of serendipity to the extent that when surprising results do occur the scientist already has a constrained set of active hypotheses and mechanisms that can be used to interpret the findings.

7.7. RISK!

One of the most intriguing aspects of this research has been the scientists assessment of risk in their research. Most of the research scientists engaged in two or more research projects. The scientists tended to pick one high risk and one low risk project to work on. The scientists categorized projects as high risk if they rated a research project as having a low probability of working out, but had the prospect of being an important discovery. They rated a project as low risk if they could see that the project had a high probability of success. Often the low risk projects were not regarded as ones that could produce important discoveries. Given that post-docs tended to be concerned with getting a job in the near future, they were often reluctant to engage in high risk projects as the high risk projects would not result in any publications and hence no job. The laboratory directors were often much more enthusiastic about high risk projects as their goals were more long term than the post docs. Furthermore, given that there were many combinations of high and low risk projects occurring in a laboratory at any one time, the probability of one of these projects working was fairly high. However, by getting the post docs to conduct combinations of high and low risk projects the directors helped ensure that the researchers would at least make a small discovery that would lead to a publication, and facilitate their own more long term goals.[4]

7.8. HOW TO MAKE A DISCOVERY

The findings discussed in this chapter have clear implications for the conduct of successful research. The following heuristics have been identified as being potentially important in making discoveries: (1) Members of a research group should have different, but overlapping research backgrounds. This will foster group problem solving and analogical reasoning. (2) Analogical reasoning should be engaged in when problems arise in the research. In particular, the scientists should engage in making both local and regional analogies. (3) Researchers should be encouraged to engage in combinations of high and low risk projects. This increases the probability that each scientist will have achieved a tangible result. (4) Take note of surprising

results. Use the surprising results to generate new hypotheses and research programs. (5) Provide opportunities for the members of the research group to interact and discuss the research by having overlapping research projects and breaking the lab up into smaller groups working on similar problems.

8. Conclusion

In this chapter I have provided an overview of two approaches that I have used to investigate the cognitive processes involved in scientific reasoning and discovery. The first approach--investigating aspects of scientific reasoning--is conducted in the laboratory. Using this approach, the researcher has experimental control over aspects of the discovery process. As in other biological sciences, this *In Vitro* approach makes it possible to isolate aspects of the reasoning process and to tease apart particular mechanisms. For example, the *In Vitro* research that I have discussed shows that the goals that subjects set are crucial to understanding scientific reasoning. Previous research using the *In Vitro* approach has also identified important components of scientific reasoning (e.g., Klayman and Ha, 1987; Klahr & Dunbar, 1988). However, as in other biological sciences, it is also necessary to investigate the processes of interest in their real-world context.

Indeed, in this chapter I have argued that *In Vivo* investigations of the cognitive processes involved in real-world scientific reasoning and conceptual change are also needed. Investigations of real-world reasoning are crucial as they reveal novel mechanisms of reasoning that would be impossible to uncover using *In Vitro* methods. For example, by using this *In Vivo* methodology, entirely new insights were uncovered regarding the ways that (i) analogies are used and their role in conceptual change, and (ii) the mechanisms underlying conceptual change in the social context of science. Further, use of this *In Vivo* approach to cognition demonstrates that it is possible to investigate complex cognitive processes in the real world. In summary, I argue here that use of the *In Vivo* approach is vital to achieving a more complete understanding of the specific mechanisms that underlie scientific reasoning and discovery.

The work discussed in this chapter shows that some of the mechanisms that have been found in the cognitive laboratory can be seen to be at work in the real world, and more importantly, that a new range of mechanisms can be uncovered by investigating real-world scientific contexts. Thus, I advance the claim that both *InVivo* and *In Vitro* investigations are necessary to understand cognition and conceptual change. As in the Biological sciences, the results of *In Vivo* work can be brought into the laboratory and analyzed using *In Vitro* methods. This cross fertilization of the two approaches ensures that neither approach becomes paradigm bound.

Acknowledgment

This research was supported by a grant from the Spencer Foundation, 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 also like to thank Laura Ann Petitto for her comments on earlier drafts of this manuscript.

Notes

1. Clearly, I had to become an expert in molecular biology, which I studied for five years in preparation for this research.

2. Note that I use the term "discovery" in the manner used in Cognitive Science. Sociologists and historians of science have argued that a finding is not a discovery until the scientists have convinced other scientists that their finding is a discovery. With this view in mind, I am continuing to investigate these scientists to see whether and *how* their findings become accepted by the scientific community as a discovery.

3. It is important to note that we regard these different types of analogies as being along a continuum, rather than being discrete classes of analogy. A more complete account of our findings on analogical reasoning appears in Dunbar 1993b

4. See Dunbar (1993c) for a detailed account of the role of risk assessment in scientific research.

References

Baker, L., & Dunbar, K. (1993b). Falsification bias in Expert reasoners. Research in progress

Brutlag, D.L., Galper, A.R., & Milis, D.H. (1991). Knowledge-based simulation of DNA metabolism: prediction of enzyme action. CABIOS, 7, 9-19.

Campbell, D.T., & Stanley, J.C. (1963). Experimental and quasi-experimental designs for research. Houghton Mifflin Company.

Dunbar, K. (1993a). Scientific reasoning strategies for concept discovery in a complex domain. *Cognitive Science*

Dunbar, K.(1993b). Real-world analogical reasoning. Manuscript in preparation

Dunbar, K. (1993c). Cognitive components of scientific discoveries that are determined by Risk management and social constraints. Manuscript in preparation.

Dunbar, K., & Baker, L. (1993a). Anatomy of a scientific discovery. Manuscript in preparation

Dunbar, K., & Baker, L. (1993b). Group Problem solving: What is it, and when is it successful? Manuscript in preparation

Dunbar, K. (1989). Scientific reasoning strategies in a simulated molecular genetics environment. Proceedings of the 11th annual meeting of the Cognitive science Society. Erlbaum: Hillsdale, New Jersey.

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. (1984). *Protocol Analysis: Verbal Reports as Data* . Cambridge, MA: MIT Press.

Fujimura, J.H. (1987) Constructing 'Do-able' Problems in cancer research: Articulating alignment. *Social Studies of Science*, 17, 257-293.

Giere, R.N. (1988) *Explaining science: A cognitive approach*. Chicago, IL: University of Chicago press.

Gooding, D. (1992). The procedural turn. In R.N. Giere (Ed.), *Minnesota studies in the philosophy of Science. Vol. XV: Cognitive models of Science*. Minneapolis: University of Minnesota Press.

Holmes, L. (1985) *Lavoisier and the chemistry of life: An exploration of Scientific creativity*. Madison: University of Wisconsin Press.

Judson, H.F. (1979). *The Eighth day of creation*. New York, NY: Simon & Shuster.

Karp, P.D. (1989) Hypothesis formation as Design. In Shrager, J. & Langley, P. (Eds.),*Computational models of discovery and Theory formation.* Morgan Kaufmann, San Francisco, CA. pp 275-315.

Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science*. 12, 1-48.

Klahr, D., Dunbar, K., & Fay, A. L. (1989) Designing good experiments to test bad hypotheses. In Shrager, J. & Langley, P. (Eds.),*Computational models of discovery and Theory formation*. Morgan Kaufmann, San Francisco, CA. pp. 355-402.

Knorr-Cetina, K.D. (1983). The ethnographic study of scientific work: Towards a constructivist interpretation of science. In K.D. Knorr-Cetina & M.J. Mulkay (Eds.), *Science observed: Perspectives on the social studies of science.* Beverley Hills CA: Sage

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

Latour, B., & Woolgar, S. (1986). *Laboratory life: The construction of scientific facts.* Princeton: Princeton University Press.

Mahoney, M.J., & DeMonbruen, B.G. (1977). Psychology of the scientist: An analysis of problem solving bias. *Cognitive Therapy and research*, 1, 229-238.

Mitroff, I.I. (1974). *The subjective side of science.* Amsterdam: Elsevier

Mulkay, M., & Gilbert, G.N. (1983). Scientist's theory talk. *Canadian Journal of Sociology*, 8, 179-197

Mynatt, C.R., Doherty, M.E., & Tweney, R.D. (1978). Consequences of confirmation and disconfirmation in a simulated research environment. *Quarterly Journal of Experimental Psychology, 30* , 395-406.

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.* Minneapolis: University of Minnesota Press.

Nisbett, R.E., & Wilson, T.D. (1977). Telling more than we can know: verbal reports on mental processes. *Psychological Review*, 84, 231-259.

Sanderson, P.M., Scott, J.J.P., Johnston, T., Mainzer, J., Watanabe, L.M., and James, J.M. (1993). MacSHAPA and the enterprise of exploratory sequential data analysis. Manuscript submitted for publication.

Stein N. (1992). A taxonomy of discourse processes. Manuscript in preparation

Thagard, P. (1993) Hundred best Analogies. *Canadian Journal of Artificial Intelligence*

Tweney, R.D. (1985) Faraday's discovery of induction: A cognitive approach. In D. Gooding & F. James (Eds.), *Faraday rediscovered*. New York: Stockton Press.

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