How Scientists Build Models:
InVivo Science as a Window on the Scientific Mind

Kevin Dunbar, ~~McGill University~~

AS of JULY 2001
at Dartmouth College, Hanover NH 03755

This chapter appeared in : Magnani, L., Nersessian, N., & Thagard, P. (1999)
Model-based reasoning in scientific discovery. Plenum Press, pages 89-98

How Scientists Build Models
InVivo Science As A Window On The Scientific Mind

Kevin Dunbar, McGill University

Abstract

How do scientists think, reason and generate new models and theories? How do scientists represent their knowledge? Answers to these questions are of paramount importance not only in understanding what science is, but also in assessing different theories of science. Surprisingly, we know little about the basic processes that are involved in current-day scientific thinking. The goal of my research is to look at what scientists actually do in their research, what types of thinking and reasoning strategies they use, and how they change their knowledge. Over the past decade, I have been investigating scientific thinking by scientists in their own labs, reasoning about their research and by conducting experiments on scientific thinking and model building in my own laboratory. The labs that I have been investigating are molecular biology and immunology laboratories in the U.S., Canada, and Italy. I have found that one place where much reasoning and new discoveries are made is at weekly lab meetings. We have performed extensive cognitive analyses of these meetings and have identified some of the key components of contemporary scientific thinking that are important in generating new models, modifying old models and solving difficult problems. In this paper I will outline four activities that are important in model building: Analogical Reasoning, Attention to Unexpected findings, Experimental Design, and Distributed Reasoning.

Introduction: A new way of investigating science-The InVivo Approach

Kevin Dunbar, McGill University

What scientists do and how they think has been a surprisingly difficult issue to investigate. Researchers have used notebooks, interviews, diaries, historical reconstructions and colleagues' accounts to obtain a picture of the cognitive processes that underlie scientific thinking and model building. Using these approaches, researchers have built many important accounts of what happens in science. However, over the past decade we have been pursuing a different approach to investigating the ways that scientists think and reason. We have been investigating the thinking and reasoning that occurs "live" at laboratory meetings in molecular biology and immunology laboratories (Dunbar, 1995, 1997, 1999). What we have done is to videotape and audiotape scientists as they think and reason in their own labs. We then analyze, sentence by sentence, the types of thinking and reasoning that the scientists use when formulating theories, analyzing data, designing experiments, and building models. We have investigated leading laboratories in The United States, Canada, and Italy. Each laboratory we have followed from three months to a year, taping the weekly laboratory meetings and supplementing the meetings with interviews and other documents such as grant proposals, drafts of papers, and one-on-one meetings. Using this approach we have been able to build detailed models of the cognitive underpinnings of the scientific mind. In particular, we have been investigating the roles of Analogy, Unexpected findings, and Group Reasoning in contemporary science.

Why use the lab meeting as a source of data on scientists? I found that the most representative cross-section of the ways scientists think and reason is the weekly laboratory meeting. It is in the weekly lab meeting that scientists present their data, reason about findings and propose new experiments and theories. Thus, by analyzing the reasoning that scientists use at their lab meetings it was possible to obtain a clear vista into how scientists really reason. Furthermore, much of what happens at meetings consists of spontaneous reasoning, and sometimes scientific discoveries and breakthroughs are made at these meetings. Using this method it is possible to directly monitor thinking and reasoning rather than uncovering reasoning through after-the-fact 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. This approach allows us to see science as it unfolds rather than the selective information that the scientists put in their notebooks or reconstruct when attempting to remember what happened in their lab. The lab meetings therefore provide a unique opportunity to see science "live." Borrowing terminology from biology I have called this the "InVivo" cognitive approach (Dunbar 1993, 1995, 1997, 1999a). Another approach that we have used is the "InVitro" Cognitive approach where we conduct experiments on scientific thinking in our own cognitive laboratory. By using a combination of InVivo and InVitro approaches it is possible gain a clearer picture of the cognitive processes involved in science.

A typical lab meeting

The laboratories that we have been investigating consist of a senior scientist, three or four post-doctoral fellows, five or six graduate students, and one or two technicians. Each week one of the members of the lab will present the results of their latest experiments. The presentations consist of a brief rationale, many data slides, and what they are going to do next. While this is the basic plan of a meeting, some tend to focus more on theory, others on method, others on data, and many on future experiments. Thus, the meetings provide a nice cross section of the things that scientists do and provide a window into the workings of the scientific mind. While the lab meeting does have a presenter, the format of the meetings is very informal. Much of the lab meeting consists of many scientists reasoning about all aspects of the enterprise, models are discussed, diagrams drawn, inductions made, deductions given, competing models discussed, experiments designed and dissected. In addition, the possibility of alternate models, methodological errors, and feasibility of various approaches is discussed. From the point of view of a cognitive scientist, this type of data is unrivalled in its richness. Here we can see all the basic reasoning processes at work and how different cognition really works in science. These meetings are supplemented by interviews both before and after the meetings that provide detailed background information.

Using this InVivo approach we have uncovered a number of basic features of the way scientists think, reason and build models. In the next sections I will provide an overview of four important aspects of building scientific models –Analogy, Unexpected findings, Distributed Reasoning, and the cognitive processes subserving experimental design. While each of these different aspects of model building can be discussed separately, it is important to note that scientists freely move between analogy, unexpected findings, distributed reasoning, and designing experiments. Thus, while I have four different sections on these aspects of scientific model building, I do not mean to imply that each process works alone, or moves along in a serial manner.

Analogy

Analogy has been regarded as an important cognitive component of scientific thinking for many decades. Furthermore Analogical Reasoning 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 ain making an analogy (e.g., Forbus, Gentner, & Law, 1995; Holyoak & Thagard, 1995, 1997). A number of cognitive accounts of analogy have noted that many different types of analogy are involved in science (see Gentner, 1997; Holyoak & Thagard, 1995). Accounts of analogy distinguish between two components of an analogy; the target and the source. The target is the concept or problem that the scientist is attempting to solve or explain. The source 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 source onto features of the target. By mapping the features of the source 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 source 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.

The Rutherford example and other examples of analogy in the history of science raise a number of important questions regarding the role of analogy in science. Do scientists use analogy at all in their day-to-day science? 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 empirical psychological work would suggest (see Forbus et al. 1995; Holyoak & Thagard, 1995)? Does analogy work alone, or does it work on conjunction with other mental operations? Finally, is analogy involved in scientific discoveries and model building in science?

We sought to investigate the role of analogy in science by coding the use of analogy in 4 laboratories. We coded each use of an analogy. Any instance where a scientist drew knowledge from one domain and used that knowledge to either fill in gaps in their knowledge of the issue they were investigating, or used their prior knowledge to help others understand an issue. What we have found is that scientists frequently use analogy when there is not a straightforward answer to their current problem. Their problem may be in formulating a theory, designing an experiment, interpreting data, or in going from theory to an experiment. We found that analogy is frequently used. There were 99 analogies made at 16 meetings. We coded the goals that the scientists had and found that the goals that the scientists had when they made the analogies could be grouped into four classes: Formulate an Hypothesis, Design an Experiment, Fix an Experiment, or Explain a result. We found that the type of analogy that the scientists made were directly related to their goal. When the scientists were designing and fixing experiments, the analogies were made to very similar experiments or the same organisms. For example, the scientists might draw an analogy from one gene on the HIV virus to another gene on the virus. These types of analogies accounted for around 45% of the total number of analogies that the scientists made. However, when the goal was one of formulating hypotheses the scientists tended to make analogies to other organisms that have similar underlying structure, or might have similar underlying biological mechanisms. For example, from the Ebola virus to the HIV virus. These types of analogies were frequent, accounting for over 50% of the total number of analogies. When the goal is to explain a concept to other members of a lab or a more general audience, the scientists make analogies to a very distant domain . For example, one scientist attempted to explain the way a vacuole works by drawing an analogy to the pop song “Hotel California.” These “distant analogies” were very rare accounting for 2% of the total number of analogies.

These results were very surprising; we did not expect to see so many analogies made to other organisms. Most research on analogy has shown that people tend to make analogies based on superficial features, rather than deep structural features. In the case of science, the scientists have both a knowledge of underlying biological mechanisms, and access to structural information such as the homology of a gene in one organism to a similar gene in another organism. This knowledge makes it possible for the scientists to go beyond superficial characteristics in making an analogy. Furthermore, despite the anecdotal accounts of distant analogies having a major role in science, our findings indicate that distant analogies are rarely used, and are primarily used to explain concepts to others, rather than as a source of novel hypotheses and experiments. In particular, we found that much analogy use revolved around accounting for unexpected findings. I will discuss this use of analogy in the next section.

Overall our research on analogy paints a very different picture of analogy use in science than what the anecdotal literature would suggest and is more consistent with the psychological literature (e.g., Holyoak & Thagard, 1995). . Rather than new theories and hypotheses being generated by making analogies to very different domains, new theories are generated by making analogies to related domains. One final note on analogy is that we have found that scientists have little memory for the analogies that they use. When we go back and ask the scientists to remember how they generated a new concept or solved a problem, at a meeting that we recorded, they have little memory of how it occurred. Thus, biographies and anecdotes do not mention many of the analogies that went into making a scientific discovery.

The Unexpected: A Common Occurrence!

One of the most frequently mentioned aspects of scientific discovery is that a finding was due to chance or was unexpected. The recent discoveries of Naked DNA and Buckey balls, not to mention penicillin, nylon, and gravity itself are among the many significant discoveries that have been attributed to the unexpected. Many scientists have claimed that their discovery was "fortuitous," "lucky," or due to "pure chance." However, when we analyze the strategies that scientists use in their research, we can see that they structure their research to take advantage of unexpected findings, and that they actually conduct experiments that lead to unexpected findings that they can then exploit. Thus rather than being the victims of chance, scientists are ready to use chance events. Thus we can modify Pasteur’s famous phrase of “chance favors the prepared mind” to “the prepared mind favors chance.”

To answer the question of how scientist deal with the unexpected, we have been conducting extensive analyses of every finding that scientists make and what happens to the finding. An unexpected finding occurs where the outcome of an experiment is different from that predicted by the scientist. The analysis of the 12 laboratory meetings yielded 28 research projects, with 165 experiments, and they reasoned about 417 results. We found that the first step that the scientists take is to classify their data. All findings -both expected and unexpected- were classified in one of two main ways. The first was to classify the finding without reference to the other conditions in the experiment. Here, the scientist says immediately what a particular result is by looking at that result alone. For example, a scientist might look at a newly obtained gel and say “This is P30.” Clearly this way of classifying the data refers to knowledge that the scientists already have and can recognize relatively easily. The second way of classifying their data is by comparing one result to another result from a control condition and inferring what their result is by making this comparison. For example the scientists might say that “This band is larger than the control one so this is P30.” Interestingly, only 15% of the findings were classified alone, whereas 85% of their findings were classified by comparing results from one condition with results from other conditions. This holds for both expected and unexpected findings. Thus, classification of data is not a simple process. Instead, it usually consists of comparisons between experimental conditions and the numerous other conditions that scientists use in their experiments.

Experimental results were coded as unexpected if the scientist claimed that the result was not what she or he expected to find. When we divided their results into expected and unexpected findings, we found that over half of their findings were unexpected (223 out of 417 results). Thus, rather than being a rare event, the unexpected finding was a regular occurrence that the scientists reasoned about. The large number of unexpected findings is important: It is not the case that scientists can take any unexpected finding and make a discovery merely by focussing on it. Rather, the scientists have to evaluate which findings are due to methodological error, faulty assumptions, chance events, or to new mechanisms.

Once a finding was classified, expected and unexpected results were treated in different ways. Expected results usually led to the next step in a sequence of experiments, whereas unexpected findings led to either replication, change in the protocol, or use of an entirely new protocol. Interestingly, scientists initially proposed a methodological explanation for their unexpected findings. For 223 unexpected findings 196 methodological and 27 theoretical explanations were offered. Thus, the first strategy that the scientists used was to blame the method. One reasoning strategy that they used to support their methodological explanations was to draw analogies to other experiments that have yielded similar results under similar conditions, usually with the same organism. For example, if one scientist obtained an unexpected finding, another scientist might draw an analogy to another experiment using the same organism in their lab and say "Oh. I also incubated the cells at 37 degrees and failed to get digestion of trypsin, but when I incubated the cells at zero degrees I got it, maybe you should do that." The types of analogies that the scientists used at this point in their research were what we term "local," that is, the analogies were to very similar experiments, usually with the same organism and within their own lab. Using local analogies is the first type of analogical reasoning that scientists use when they obtain unexpected findings and is an important part of dealing with such findings. Scientists rarely mention these types of analogies in their autobiographies as they appear to be uninteresting, yet these types of analogies are one of the first reasoning mechanisms that are used in dealing with unexpected findings.

The way that the scientists reason about their data changed when they obtained a *series* of unexpected findings. This usually occurred when the scientist continued to obtain unexpected findings despite having modified the method, attempted to replicate the finding, or when they obtained a whole series of unexpected findings in related conditions. It is at this point that a major shift in the reasoning occurs; the scientists begin to offer new more general models, hypotheses, or theoretical explanations. In fact, Model building is most important and most common when a series of unexpected findings are made. Thus, for repeated unexpected findings 51 theoretical explanations were offered as compared with 33 methodological ones. These models usually covered anywhere from 2-5 unexpected findings. The types of reasoning that occur at this point were very different from the reasoning that scientists initially used when they obtained isolated unexpected findings. Three different types of reasoning processes were used to reason about a series of findings that were repeatedly unexpected. First, was to draw analogies to different types of mechanisms and models in other organisms rather than making analogies to the same organism. This also involves making analogies to research conducted outside their lab. The scientists switched from using local analogies to more distant analogies. For example, a scientist working on a novel type of bacterium might say that "*IF3* in *ecoli* works like this, maybe your gene is doing the same thing." Thus, the way that analogies are used dramatically changes now.

A second type of model building that occurs following a series of unexpected findings is that the scientists attempt to generalize over the series of findings. Often the scientists will search for common features of the unexpected finding and attempt to propose a general model that can explain these results. For example, in one immunology lab two different types of T-cells went to two different places that neither cell should have been able to enter. The scientists asked what was common to the two different types of cells and inferred that it was the cells had a particular property. Starting with this generalization, the scientists then attempted to build a causal chain of events that could provide a complete model. The causal models that the scientists build usually specify a starting state, a process that transforms the starting state to another state, what the components of this process are and what is the end state. Often this model building phase goes through a number of cycles where a number of small models are linked together to provide a larger model of the entire biological process that is being explained. This model building process makes reference to findings from the scientists' own lab, other people's labs, and general knowledge of the domain. Note that many reasoning processes are involved here: causal reasoning, generalization, analogy, deduction and visual reasoning are often all present. Inevitably, there are numerous gaps in the models that the scientists propose. Often the gaps are crucial. The scientists may think that some mystery protein is causing something to happen. This will lead to the proposal of new experiments to find the protein. New experiments will also be proposed to test the predictions of the model.

How do we understand the role of the unexpected in science? Again it is important to note that unexpected findings are common and that scientists expect the unexpected. Scientists use the unexpected in many different ways. First, as we will see in the next section, scientists build very complex experiments with control conditions that often capture unknown and unexpected processes. Second, scientists attend to both expected and unexpected findings. Initially, they gave methodological explanations for their findings, and support these explanations by making analogies to very similar results obtained in other experiments. Third, if they continue to see the same type of unexpected result despite changes in methodology, they form new models and theories using more distant analogies and group reasoning. Thus, it is only when the scientists pass a threshold in unexpected findings that they propose general model and discoveries begin to unfold.

Model Building and Experimental Design

A crucial component of model building in science is the experiment itself. While some analyses have been conducted on the design of experiments over the past twenty years, relatively little is known on what really goes on when an experiment is designed (but see Friedland & Isakawi, 1985; Gallison, 1987; Hacking, 1983; Kulkarni & Simon, 1988). Instead, many researchers have assumed that an experiment is virtually entirely predetermined by the scientists' theoretical stance, or that the apparatus that the scientists use is so permeated by the theoretical assumptions that the scientists hold, that the experiment is again merely an appendage a theory. This viewpoint leads to the conclusion that experimental design is of little interest. While these views of the experiment are certainly interesting, we wanted to go beyond them to see what scientists actually do when an experiment is designed. Lisa Baker and I have conducted a series of analyses, and yes, experiments to determine what the cognitive components of the experimental design process are (Baker & Dunbar, 1996, 1999). It turns out that while theory is usually the starting point for experiments, much of the decision making that takes place while designing an experiment takes place independent of the hypothesis being tested.

Our first investigations of the experimental design process consisted of a qualitative analysis of designing an experiment in an immunology laboratory (Baker & Dunbar, 1996). What we found was that experimental design consists of a number of basic processes. The first component is choosing an overall method for investigating the hypothesis or question that the scientist has. It is at this point that the relation between theory and experiment is strongest. After selecting an overall type of methodology, the scientist’s concern is with getting the method to produce an interpretable result. While the subsequent steps can follow any order, I will describe them in a serial fashion. Following the choice of a particular methodology the scientists must fill in the details of the experiment. What they do is to unpack the design, asking questions about which components are needed and what their values should be. The scientists evaluate each component of the experiment locally while at the same time attempting to coordinate each local component of the experiment to see if it will really work. What we found was that four sets of criteria are used when they evaluate various components of their experiment: ensuring a robust internal structure to the experiment, optimizing the likelihood experiments will work, performing cost/benefits analyses on possible design components, and ensuring acceptance of results by the scientific community.

Our qualitative analyses of experimental design revealed that while the experiments that the biologists conduct are very complex, the scientists appeared to be working with a general model of experimental design that guided what they did. One key aspect of the experiments is the choice of conditions. Scientists design experiments that have both experimental conditions and control conditions. The experimental conditions usually consist of a manipulation of the variable that the scientist is interested in. For example if the scientists thinks that a particular protein is responsible for a cell having a particular function, the scientist might add this protein to cells and to see if the cells gain this function. However, the scientists also have numerous control conditions in their experiments. One type of control that they use is "known controls” (Baker & Dunbar 1996, 1999).

“Known controls" consist of conditions that have previously been used and validated and are standardized. These known controls are important for comparisons. What these controls do is to ensure that the technique really works and can allow the scientist to conclude that they are conducting the experiment in an appropriate manner. In the complex experiments that the scientists conduct, many things can go wrong. By having numerous known controls the scientists can ensure that the various steps in the experiment really are working. Another type of control is a "baseline control" in which something is taken away from the experiment or not added to an experiment. This "baseline control" tests the hypothesis that doing something very similar to the experimental condition does not produce the result that the experimental condition produces. When we analyzed the types of experiments that scientists conducted we found that "known controls" were particularly important in controlling for error.

Following our initial analyses we decided to further probe the use of control conditions and the relationship of controls to error (Baker & Dunbar 1999). We analyzed the design of experiments at four meetings in two Canadian immunology labs. We coded the goal that the scientist had (testing an hypothesis, or anticipating error) and we also coded the type of control conditions used ("baseline control "or "known control"). What we found was that most "baseline controls" were proposed when testing an hypothesis, and most "known controls" were used when ruling out potential errors in use of the techniques. This result is interesting as it shows that the use of controls, particularly 'known controls" is one way that scientists attempt to determine whether methodological errors have occurred in their experiments. One important point to note is that these controls are usually put into the experiment before the experiment is conducted. The scientists are attempting to minimize error before they conduct an experiment. Thus the scientists assume that an experiment is an error prone process. We also saw this in the previous section on unexpected findings where scientists usually interpret the unexpected findings as being due to error. Clearly, the possibility of error underlies the models that the scientists use and their initial reaction to unusual data.

Having found that scientists spend much of their time both anticipating and dealing with error, we decided to go back into the cognitive laboratory and conduct our own experiments on the ways that scientists deal with error (Baker & Dunbar 1999). In one experiment we asked immunology and molecular biology students to design experiments that would test an hypothesis regarding a particular gene function. Once the students had designed their experiments we told them that the professor who wanted the experiment done was afraid that there was an error in the experiment. Thus, initially we gave the students the goal of testing a hypothesis and then we switched the students to the goal of worrying about the possibility of error in their experiments. Would the science students, like the scientists, use "baseline controls" when testing an hypothesis and "known controls" when anticipating error? Yes. We found that both immunology and developmental biology students used "baseline controls" when testing hypotheses and "known controls "when anticipating error. These results indicate that using "known controls" is a common strategy that scientists use for anticipating error and interpreting results once they obtain data. Interestingly, non-science students did not use this "known controls" strategy.

Overall, our work on experimental design indicates that while experiments are usually designed in the context of theoretical questions, much of the experimental design process is constrained by methodological considerations. While it is also true that the methods and instruments that the scientists use are predicated on theory, these background assumptions are really background assumptions that are rarely considered. What is important here is that experimental design is a complex cognitive process in which the scientists must navigate through many obstacles to get a result. The process by which the scientists do this is appears to be a search in a problem space (see Klahr & Dunbar, 1988). One important side effect of the use of complex experiments with many controls is that these controls, particularly baseline controls, can often reveal hitherto unknown phenomena that are at the root of many discoveries attributed to unexpected findings. The use of multiple conditions thus serves a dual function – guarding against error and generating potentially important unexpected findings that can underlie later model building.

Distributed Reasoning: An important source of new models in Science

The image of science that we are all familiar with is one of the lone scientist toiling away under a naked light bulb for long hours. Suddenly, inspiration strikes and the scientist has made a discovery. This is the image of science that has motivated much research in cognitive science. Its key features are that the scientist works alone and that scientific discovery occurs in a flash of insight. How representative of contemporary science is this view? Our research and that of other cognitive scientists such as Paul Thagard (Thagard, 1997) is revealing that reasoning in science, particularly at the critical moments of hypothesis formation, experimental design, data interpretation, and discovery is by groups of scientists and not individual scientists. We call this type of reasoning distributed reasoning (Dama & Dunbar 1996, 1999; Dunbar, 1995, 1997).

What is distributed reasoning? There are many differing uses of this term. We define distributed reasoning as reasoning in which more than one person is contributing to a reasoning episode. For example, when a scientist obtains an unexpected finding the scientist and other members of the laboratory may induce a new concept to explain the finding. The different members of the lab provide different elements to the induction of the concept. In other words, induction is distributed over individuals. We have analyzed the distributed reasoning that occurs in five laboratories and have found that over 50% of the reasoning that takes place at meetings is distributed. Clearly, distributed reasoning is an important component of science. Are there precipitating events for distributed reasoning to occur? What do scientists do during distributed reasoning? Does distributed reasoning do something that does not occur during reasoning alone? These are some of the questions that we are currently investigating.

While I have reported our general findings on distributed reasoning elsewhere (Dunbar, 1997), here I will focus on some recent analyses of distributed reasoning that Mike Dama and I have conducted. What we did was to analyze sentence by sentence, the way that distributed reasoning occurs in labs. We used a new statistical technique that allows us to see how social processes such as question asking influence the ways that scientists represent their knowledge. What we found is that the major event that occurs during social interactions is generalization. The scientists in the lab often take a finding and attempt to generalize to other findings, both within and outside the lab. This process is important, as scientists frequently force each other to consider other representations and explanations for the data being considered or theory being proposed. What we see happening is that large changes in representation occur at these moments.

One place where distributed reasoning is particularly important is when a series of unexpected findings occur. When a single unexpected finding occurs, little distributed reasoning is present. In this situation, social interactions are usually a straightforward suggestion about method that does not necessitate any major change in the scientists' representation of the issue. But when a series of unexpected findings occur there is a large amount of distributed reasoning. Thus, it is not one scientist shouting "Eureka," but a number of scientists building a new model by adding different components to the model. As noted above, an important feature of this process is that often the members of the lab propose alternate models and explanations for the unexpected findings as well as resulting in the use of more distant analogies. By looking at many different laboratories we also have found that when groups of scientists reason, the diversity of the group is very important (Dunbar 1995). When all the scientists are from the same background it is difficult for them to generate multiple hypotheses, but when the scientists are from different backgrounds many different hypotheses can be generated. Clearly, these types of social-cognitive processes are at work in the peer review process underlying the publication of papers and evaluations of grant proposals.

Is distributed reasoning always beneficial? We believe that there are a number of important constraints on when and whether distributed reasoning will be successful. The first is the knowledge base that the group has. We have found that when the members of a group all have the same background, such as all having done graduate work on a particular organism, then the group performs no better than the individual. This can be seen in the analogies that the group makes. Groups from the same background will draw their analogies from that background. Groups with a varied background, but with common goals, will produce many different types of analogies and these analogies can be used to solve the problem that the lab is working on. Overall, our work on distributed reasoning indicates that it is an important component of contemporary science and is frequently an aspect of scientific discovery. Furthermore, our analysis reveals that distributed reasoning can help circumvent one of the main problems in human reasoning -generating different representations and understandings of both theory and data.

InVivo Science: What does it reveal?

Overall, the investigation of the cognitive mechanisms involved during “InVivo” scientific thinking and reasoning reveal a number of important mechanisms underlying model building in science. I have demonstrated that analogy, unexpected findings, experimental design, and Distributed reasoning all play a crucial role in the building of scientific models. Thus, no single cognitive process is involved in model building. Furthermore, while scientists can subjectively point to key moments in the changes in their models, an analysis of the changes in models over an extensive period of time indicates that new models arise through a series of small changes produced by a variety of different cognitive mechanisms. The goal of our current research is to determine how these different aspects of scientific model building are integrated to produce these new models.

Acknowledgments

This research was supported by grants from the Spencer Foundation, Office of Naval Research, NSERC, SSHRC, and FCAR. I would like to thank Laura-Ann Petitto for her comments on an earlier draft and Annie St. Laurent,Lisa Baker, and Mike Dama for their work on different aspects of this enormous project.

References

Baker, L. M., & Dunbar, K. (1996). Constraints on the experimental design process in real-world science. In G. W. Cottrell (Ed.), Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society (pp. 154-159). Mahwah, NJ: Lawrence Erlbaum Associates.

Baker, L.M., & Dunbar, K.(In press). Experimental design heuristics for scientific discovery: The use of baseline and known controls. *International Journal of Human Computer Studies.*

Dama, M., & Dunbar, K. (1996). Distributed reasoning. When social and cognitive worlds fuse. *In Proceedings of the Eighteenth Annual Meeting of the Cognitive Science Society*. 166-170.

Dama, M., & Dunbar, K. (1999).Distributed reasoning and representational change in real-world science laboratories. Manuscript submitted for publication

Dunbar, K. (1999) Beyond the myth of the unexpected: Are scientists the victims of chance? Manuscript submitted for publication

Dunbar, K. (1997). How scientists think: Online creativity and conceptual change in science. In T.B. Ward, S.M. Smith, & S.Vaid (Eds.) Conceptual structures and processes: Emergence, discovery and Change. APA Press. Washington DC.

Dunbar, K. (1995). How scientists really reason: Scientific reasoning in real-world laboratories. In R. J. Sternberg & J. E. Davidson (Eds.), The nature of insight (pp. 365-395). Cambridge, MA: MIT Press.

Dunbar, K. (1993). “In Vivo cognition: Knowledge representation and change in real-world scientific laboratories.” Paper presented at the Society for Research in Child Development. New Orleans, LA. March 26.

Forbus, K., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science,* 19, 141-205.

Friedland, P. E., & Iwasaki, Y. (1985). The concept and implementation of skeletal plans. Journal of Automated Reasoning, 1, 161-208.

Gallison, P. (1987). How Experiments End. Chicago, IL: Chicago University Press.

Hacking, I. (1983). Representing and Intervening. Cambridge: Cambridge University Press.

Holyoak, K.J., & Thagard, P. (1997). The analogical mind. *American Psychologist, 52,* 35-44.

Holyoak, K.J., & Thagard, P. (1995). Mental Leaps. Cambridge. MA: MIT Press

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

Kulkarni, D., & Simon, H. A. (1988). The processes of scientific discovery: The strategy of experimentation. Cognitive Science, 12, 139-175.

Thagard, P. (1997). Collaborative knowledge. Noûs, 31: 242-261.