

The Temporal nature of Scientific Discovery: The roles of Priming and Analogy

Kevin Dunbar & Christian D. Schunn
Department of Psychology, McGill University

Abstract: One of the most frequently mentioned sources of scientific hypotheses is analogy. Despite the attractiveness of this mechanism of discovery, there has been but a small success in demonstrating that people actually use analogies while solving problems. The study reported below attempted to foster analogical transfer in a scientific discovery task. Subjects worked on two problems that had the same type of underlying mechanism. On day 1, subjects discovered the mechanism that controls virus reproduction. On day 2, subjects returned to work on a problem in molecular genetics that had a similar underlying mechanism. The results showed that experience at discovering the virus mechanism did facilitate performance on the molecular genetics task. However, the verbal protocols do not indicate that subjects analogically mapped knowledge from the virus to the genetics domain. Rather, experience with the virus problem appeared to prime memory. It is argued that analogical mapping can be used flexibly in scientific discovery contexts and that primed knowledge structures can also provide access to relevant information when analogical mapping fails.

The origins of scientific hypotheses and theories have often been considered to be outside the purview of scientific investigation. Recently, however, a number of accounts of scientific discovery have appeared that provide detailed mechanisms of discovery (e.g., Holland, Holyoak, Nisbett and Thagard, 1986; Klahr & Dunbar, 1988; Langley, Bradshaw, Simon & Zytkow, 1987). One such account is the Scientific Discovery as Dual Search framework of Klahr and Dunbar (Klahr & Dunbar, 1988). Scientific Discovery as Dual Search (SDDS) is based upon the assumption that scientific reasoning consists of a search in two main problem spaces --an hypothesis and an experiment space. Using this framework, it has been possible to explore the manner in which the current hypothesis guides experiment space search, and how experimental results in turn determine search of the hypothesis space. For example, Klahr & Dunbar (1988) have discovered that when subjects fail to find an hypothesis by searching the hypothesis space, they switch to searching the experiment space. Dunbar (1989) and Klahr, Dunbar, & Fay (1990) have discovered

a number of heuristics for experiment space search.

While we are now acquiring a more detailed picture of the heuristics that guide experiment space search, relatively little is known about the heuristics governing the search of the hypothesis space. Klahr and Dunbar (1988) have shown that when subjects attempt to make a discovery, they initially search memory for an hypothesis. When an hypothesis has been found, the subjects then conduct experiments to see if the hypothesis is correct. If the hypothesis is incorrect, the subjects then either modify the hypothesis, search memory for a new hypothesis, or switch to a search of the experiment space. However the actual mechanisms underlying this search of the hypothesis space are unclear. The purpose of the research reported here was to provide a more detailed account of Hypothesis space search.

Research in other laboratories suggests two main mechanisms for hypothesis space search; analogical mapping (Gentner, 1983; Holland, Holyoak, Nisbett & Thagard 1986), and Reminders (Ross; 1989). Holland, et al. (1986) have emphasized the central role of

analogical mapping as a source of initial scientific hypotheses. Surprisingly, empirical research has demonstrated that subjects often fail to notice an analogical mapping when one is present. Most research that has demonstrated an inability of subjects to make use of analogies has been conducted using problems that demand little prior knowledge, and the subjects do not work on the problems for extensive amounts of time (e.g., Gick & Holyoak, 1983). When subjects are given more extensive training on a source domain there is a greater probability of analogical mapping occurring (e.g., Bassock & Holyoak, 1989). This suggests that the type of representation that subjects have of both the source and the target domain determines whether analogical mapping will occur. Gentner's (1983) structure mapping theory suggests that subjects will map from one domain to another only when they can abstract the relational structure of the source domain and the target domains. Both Gentner's theory and the goal oriented approach of Holyoak and Thagard (1989) suggest that subjects must have a detailed representation of the both the source and the target problems to make analogies.

Another mechanism that could be involved in mapping previously acquired knowledge on to a new problem is reminders; the use of an earlier example that the problem solver is reminded of while solving a problem. Reminders must be part of the analogical mapping process or could occur without being used in analogical mapping. Ross (1989) has proposed that when solving a problem there is a search of memory for related information. If the search is successful, the subject will be able to use the earlier example to solve the problem. If the earlier example has the same structural characteristics as the current problem the subject may be able to analogically map the earlier solution on to the current problem. If there is no structural overlap the subject could still use the reminding by fitting the previous solution to the current problem.

When applied to a scientific reasoning context, these views suggest that when a scientist has knowledge of a source domain,

analogical mapping will occur only when a representation of the target domain has the same structure or the same subgoals as the source problem. When working on a problem, the scientist will initially retrieve some knowledge thought to be relevant to the problem. Experiments are then conducted. The experiments lead to a more detailed representation. When a more detailed representation of the problem is acquired, a search of memory for a structure or solution that will map onto the target domain will occur. While working on a problem a scientist may retrieve many different analogies to formulate hypotheses until a satisfactory solution to the problem is reached.

In sum, previous research suggests that if an attempt to map from the base to the target domain is made when subjects begin working on a problem, subjects will have superficial information about the target domain. Subjects should use their superficial knowledge of that domain to search memory for hypotheses relevant to the problem. As a result of experimentation, subjects should learn that their initial hypotheses do not hold. Then, as they learn more about the relations among the elements of the target problem --through experimentation-- they should develop a more detailed representation of the problem. This should make it possible to search memory for a structurally similar solution, or solutions that have been used to solve similar problems.

The Molecular Biology domain

To investigate the temporal course of the development of hypotheses a task that fosters the development of changing representations of a problem is needed. Rather than invent an arbitrary problem, two problems from molecular biology that involve similar underlying mechanisms were used. The source problem was one of discovering why viruses are sometimes dormant (do not reproduce), and other times are active (reproduce rapidly). Molecular biologists have discovered that dormant viruses secrete enzymes that inhibit virus reproduction. This is a form of negative regulation. The target problem also involved

negative regulation. The target problem was to discover how genes are controlled by other genes. The specific question addressed was why genes secrete beta-galactosidase only when there is lactose present. Again the mechanism underlying this is negative regulation (Jacob & Monod, 1961). Monod and Jacob were awarded the Nobel prize for discovering that the control mechanism in viruses and in lactose is negative regulation.

To investigate the time course of analogical mapping a paradigm developed by Dunbar (1989) was used. In the Dunbar (1989) study subjects were taught some elementary knowledge about molecular genetics. Then, subjects were taught how to conduct molecular genetics experiments on the computer. Then they were given the problem of discovering the way that genes control other genes. Subjects tended to use their knowledge of the genetics domain to formulate their initial hypotheses. Subsequent hypotheses were formed by inducing the concept of negative regulation from patterns of data or memory search.

In the current study, subjects were taught about viruses and genetics and were shown how to conduct simulated experiments on the computer. For both the virus and the genetics problem subjects were asked to discover the control mechanism. The tasks that the subjects were given and the potential analogical mappings will now be described in detail.

The Experiment

Subjects. Sixty McGill undergraduates were paid to participate in the experiment. All subjects had taken an introductory biology course and they were not familiar with gene regulation in bacteria or viruses. Their knowledge of molecular genetics consisted of knowing that DNA and RNA exist.

Procedure. There were three conditions. In the *No-Virus* condition, subjects participated only in the gene regulation problem. There were two experimental conditions -- *Negative-Virus* and *Correlated-Virus*. In these conditions subjects attempted to discover why viruses are sometimes active

and sometimes dormant. The virus problem was presented on day 1, and the genetics problem was presented on day 2. Subjects conducted their experiments on a MacIntosh II computer. The display was highly interactive: subjects conducted experiments by using menus and selecting various options. Once the experiment was designed subjects could conduct the experiment and monitor the results. A permanent record of all experiments and results was available after each experiment was conducted.

Day 1. Virus Problem. The virus problem was carried out in three phases. First, the subjects were taught some basic facts about viruses and biochemistry, and were shown two methods of determining whether a chemical is involved in virus reproduction. One method was to view the amount of a chemical that is present in an active or dormant virus. By doing this the subject could see whether there was a relationship between the amount of a chemical and the number of viruses present. The second method was to add a chemical to an active or dormant virus and monitor its effect on the number of viruses present.


Second, the subjects were instructed on how to give a verbal protocol. Third, the virus problem was given to the subjects. They were told that viruses were sometimes active and sometimes dormant, and that biochemists thought that one of three chemicals that the virus can secrete was responsible for making the virus active or dormant. Subjects were asked to discover which of the three chemicals causes the virus to be active or dormant and what the control mechanism was.

For the *Negative-Virus* group the mechanism causing the virus to become dormant was negative regulation: The dormant viruses secrete a chemical that prevents generation of new viruses. For the *Correlated-Virus* group, there was a positive correlation between the amount of Q that was present and the number of viruses. In this case, Q was a by-product of the number of viruses present. There was no causal mechanism for this group; this was a control condition. Subjects worked on this problem until they felt that they had

discovered the correct answer. Subjects were not told if they had discovered the correct mechanism. All subjects were asked to return for another problem the next day.


There were 12 possible experiments that could be conducted. in the virus problem. The small number of experiments made it extremely likely that subjects would conduct all possible experiments. In particular, in the *Negative-Virus* condition, subjects should conduct an experiment where Q is added to an active virus and the number of viruses decreases, (see Figure 1). This would suggest that Q inhibits virus reproduction. Due to the small size of the experiment space it was expected that most subjects would discover the correct mechanism.

FIGURE 1: NEGATIVE-VIRUS CONDITION



Day 2. Genetics problem. All three groups of subjects were given the genetics problem. There were three phases in the genetics problem. First, the subjects were taught some basic facts about molecular biology. They were told that certain genes control the activities of other genes by switching them on when there is a nutrient present. This is an example of positive regulation --the regulator gene senses that there is a nutrient present and then releases substances that instruct other genes to secrete enzymes that can utilize the nutrient. The example of positive regulation given to the subjects can be seen in Figure 2. Here, the A gene switches on the enzyme producing genes. When the A gene is absent and fructose is present no enzyme is produced

(FIG 2A). When the A gene is present in a diploid cell the cell secretes an enzyme that breaks down the fructose (FIG 2B).



In the second phase of the genetics problem, subjects were shown how to give a verbal protocol. In the third phase, subjects were presented with two findings about a set of genes labelled I, O, and P: The enzyme producing genes secrete an enzyme when there is lactose present, and do not secrete the enzyme when there is no lactose. Subjects were asked to discover how the enzyme producing genes are controlled. The presentation of the I,P,O genes was identical to the A,B,C genes. The differences were (i) that there were different letters on the genes (ii) the nutrient was lactose rather than fructose, and (iii) the enzyme secreted was betagalactosidase rather than delta. Finally, unbeknownst to the subject, there was a different underlying mechanism --negative regulation rather than positive regulation. The subjects have to discover that the I and O genes negatively regulate (i.e., inhibit) the activity of the enzyme producing genes until a nutrient is present. Subjects must also learn that the I gene can be present on either the male or female chromosome, but that the O gene can only be on the O gene.

There are 120 possible experiments that can be conducted (6 amounts of nutrient x 20 genetic combinations). For a normal Ecoli, the amount of enzyme produced is half whatever the amount of nutrient is. The I gene mutants (designated I- mutants) produce an output of 876, and O- mutants produce an output of 527. The I- and O- mutants produce this amount of enzyme regardless of the amount of nutrient administered. Thus, I- and O- nutrients will produce enzymes even when

no nutrient is administered. This is a strong clue for negative regulation being involved. In this study, the P gene plays no role at all.

Subjects were expected to make a number of different mappings from both the virus and the genetics domains during the course of the experiment. First, there should be a mapping from the positive regulation example given at the start of the genetics problem. The positive regulation example had a virtually identical structure to the target problem. There were three regulator genes, three enzyme producing genes, nutrients with similar names (lactose and fructose), and enzyme products that also have similar names (beta and delta). Thus, as in the Dunbar (1989) study, it was expected that initial hypotheses would be of positive regulation.

Once subjects discover that none of the genes work by positive regulation there should be a search of the hypothesis space for another mechanism. For the *Negative-Virus* group, it was expected that the presence of large amounts of enzyme when either the I or O genes were mutants would be equivalent to the situation of large numbers of virus being present when Q was absent. At this point, subjects should mention that the genetics problem is like the virus problem. The subjects should then map over the concept of negative regulation from the virus to the genetics domain.

If the virus problem is used to help solve the genetics problem, then it could occur in a number of stages. First, subjects would have to extract a representation such as *X inhibits Y* from *Q inhibits virus production*. Then they should apply this relation to the genetics problem. Second, even with this mapping, the problem is not solved; the *X inhibits Y* relation must be modified to include two genes (I and O), and the that the I and O genes inhibit in different ways. Thus, subjects should retrieve the *X inhibits Y* relation, and then modify it (by experimentation) to fit the current problem. The genetics problem not only involves a memory search for hypotheses, but also a modification of hypotheses to solve the problem.

Results

Day 1 (virus problem). All of the 20 subjects in the *Correlated-Virus* condition discovered that the positive non-causal relationship between the amount of Q and the number of viruses. Sixteen of the 20 subjects in the *Negative-Virus* group discovered that the dormant virus secretes a chemical that inhibits virus reproduction. The four subjects that did not discover negative regulation concluded that there was a complex interaction of chemicals that switch on reproduction.

Day 2 (genetics problem). As shown in Table 1, for the *No-virus* group, 7 out of 20 subjects discovered that the I and O genes inhibit and that the P gene plays no role. For the *correlated-virus* condition, the results were almost identical: 7 subjects discovered the roles of I, O, and P. This indicates that the experience with conducting virus experiments on day 1 had no beneficial effect on performance. Results for the *Negative-Virus* group were very different; 12 of the subjects discovered the mechanism of genetic control. This indicates that discovery of the negative regulation of viruses on day 1 did indeed benefit performance on the genetics problem.

Table 1
Subjects discovering genetics mechanism

Condition	No-Virus	7 (35%)
	Correlated virus	7 (35%)
	Negative virus	12 (60%)

As can be seen from Table 2, 12 of the subjects who discovered negative regulation in the virus problem also discovered negative regulation on the genetics problem. Four of the subjects who discovered negative regulation in the virus problem did not discover negative regulation in the genetics problem. These subjects did mention inhibition, but also proposed a positive regulation mechanism.

Table 2
Effects of discovering virus mechanism on discovering genetics mechanism

	discover genetics	miss genetics
discover virus	12	4
miss virus	0	4

Initial hypotheses. It is possible to determine whether subjects immediately saw if the mechanism that they learned in the virus problem was applicable to the genetics problem by examining initial hypotheses. None of the subjects in the *No-Virus* or the *Correlated-Virus* conditions initially proposed negative regulation as a potential mechanism. Only one of the subjects in the *Negative-Virus* condition proposed that negative regulation as an initial hypothesis. This indicates that subjects initial mappings were not from the virus domain. Instead, initial hypotheses were taken from the example demonstrating positive regulation at the beginning of the genetics problem.

Subsequent hypotheses. As mentioned above, previous research using this paradigm (Dunbar, 1989) has shown that subjects discover negative regulation using two main strategies. One is by searching memory for an appropriate mechanism and the other is by inducing the mechanism from a series of experimental results. If subjects were using the knowledge obtained from the virus experiment, then subjects should discover the concept of negative regulation by a memory search, rather than by inducing negative regulation from a series of experimental results

We were able to categorize performance as memory search by one of two criteria. One criterion was whether subjects stated that they were searching memory for hypotheses or that the genetics problem reminded them of the virus problem. None of the subjects in either the *Correlated-Virus* or the *Negative-Virus* groups mentioned the virus experiment while they were working on the genetics problem. The other criterion was if on the first experiment that was inconsistent with positive regulation, subjects proposed negative regulation. For example, when an I-experiment is conducted there is a very large output of betagalactosidase. This result is inconsistent with positive regulation. All three groups of subjects conducted an I- or an O- experiment as one of their first four experiments. In the *No-virus* condition, 4 subjects suggested negative regulation after their first I-

experiment. In the *Correlated Virus* condition 4 subjects proposed negative regulation. By contrast, 11 subjects in the *Negative-virus* condition proposed negative regulation .

Incorrect final hypotheses. It is also possible that the virus problem influenced the types of incorrect hypotheses proposed. In the *No-Virus* condition only 5 of the 13 subjects mentioned negative regulation. In the *Correlated-Virus* condition 4 of the 13 subjects mentioned negative regulation. In the *Negative-Virus* condition 8 subjects failed to propose the correct mechanism. The 4 subjects who missed both the virus and the genetics problem never mentioned inhibition. Three of the 4 subjects who solved the virus problem but did not solve the genetics problem mentioned negative regulation. Thus, all but one of the subjects who solved the virus problem proposed some form of negative regulation in the genetics problem.

One possible explanation for the better performance of the *Negative-Virus* group could be that these subjects learned various experimental skills while working on the virus problem. The performance of the *Correlated-Virus* group does not support this explanation. If subjects acquired knowledge of how to conduct experiments during the virus experiment, then both the *Correlated-Virus* and *Negative-virus* groups should have benefitted. As only the *Negative-virus* group showed improved performance, this explanation seems unlikely.

Discussion

The results of this study illustrate the dynamic nature of the scientific discovery process. Initial hypotheses were formed by analogically mapping from the same domain. When initial hypotheses were disconfirmed there was a search of memory for an hypothesis that could account for the discrepant results. This search was not based on superficial features of the problem. Subjects searched for a mechanism that could account for the current results. Once the mechanism of negative regulation had been retrieved it had to be modified to fit the current context.

Subjects in the *Negative-Virus* condition did not appear to discover negative regulation by analogical mapping from the virus domain. If subjects were mapping knowledge over from the virus domain, then there should have been a reference in the verbal protocols to the virus domain prior to proposing an inhibitory mechanism. None of the subjects made any reference to the virus domain. Furthermore, at the end of the experiment all subjects were asked whether they had thought of the virus experiment while working on the genetics problem. All the subjects claimed that they had not thought of the virus experiment. When asked if they thought that the virus problem had helped them on the genetics problem almost all subjects stated that the experience with the genetics problem helped them design better experiments. However, the fact that the performance of the subjects *No-Virus* condition was virtually identical to that of the subjects in the *Correlated-Virus* condition indicates that experience with the virus problem *per se*. had no effect on experimental strategies. In sum, these results appear to indicate that there is no direct mapping from the virus to the genetics problem.

The fact that subjects made no explicit mapping from the virus to the genetics problem was surprising. Most theories of analogical reasoning assume that when mapping occurs, it is a conscious activity that occurs by searching memory for a match between the target problem and other possibly relevant mappings. The results of this study suggest a different mechanism: The virus problem may have primed the concept of negative regulation. When the subjects discovered that the mechanism was not one of positive regulation, they engaged in a memory search. The memory search finds negative regulation because it has been primed by the virus problem. This priming mechanism would account for the fact that (a) subjects initial hypotheses were not concerned with negative regulation, and (b) no evidence of an explicit mapping from the virus to the genetics domain appeared in the protocols.

A number of mechanisms could have produced the results obtained in this study.

One is priming of a pre-existing knowledge structure. Cheng and Holyoak (1986), and Fong and Nisbett (1989) have proposed that there are a number of well ingrained concepts in memory that they have labelled 'pragmatic reasoning schemas.' A concept such as negative regulation could be one of these types of schemas. In the current study, solving the virus problem may have primed the concept of negative regulation, making it more likely to be retrieved while solving the genetics problem.

If the prior learning episode primes a concept that is already present in memory, then it suggests that this knowledge structure priming will occur only for concepts that are represented in memory prior to participation in the experiment. However, studies of implicit learning (cf. Reber, 1989) have discovered that when subjects observe a number of instances that follow a rule, they form an abstract representation of the rule that is not tied to the particular context under which the 'rule' was learned. A similar type of learning may have occurred in this study. By learning through experimentation, subjects may have acquired an abstract concept that is missing the contextual information about how the concept was acquired. When working on the genetics problem this abstract representation may have been accessed and used to solve the problem. As the concept had no contextual information associated with it, no reminders occurred. We are currently investigating this hypothesis..

Finally, we can now provide a more dynamic account of the heuristics governing hypothesis space search in SDDS. First there is a search for an hypothesis that can account for an instance that has similar features to the target problem. These initial hypotheses can be based upon the superficial features of the target problem and also upon what little is known about the deep structure features of the problem (cf. Faries & Reiser, 1988). Second, if this search is successful, then the hypothesis is tested. Third, if the hypothesis fails the test, then there is a search for other potential hypotheses, if none are found there is an attempt to modify the initial hypothesis to accommodate the data. This later search is

qualitatively different from the first search; subjects now have more knowledge of the target problem and can search for an underlying mechanism rather than a match for superficial features. When a mechanism is being searched for, the current state of activation of all mechanisms in memory will determine which mechanisms will be retrieved. If, as in the current study, a mechanism has been primed by a prior problem, then it will be retrieved.

References

- Bassock, M., & Holyoak, K.J. (1989). Interdomain transfer between isomorphic topics in algebra and physics. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 153-166.
- Cheng, P., & Holyoak, K. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, 17, 391-416.
- Dunbar, K. (1989). Scientific reasoning strategies in a simulated Molecular genetics environment. In *the proceedings of the 11th annual meeting of the Cognitive Science society*, MI Ann Arbor, 426-433.
- Dunbar, K., & Klahr, D. (1989). Developmental differences in Scientific Discovery Strategies. In D. Klahr, & K. Kotovsky (Eds.). *Simon and Cognition: Proceedings of the 21st Carnegie-Mellon Symposium on Cognition*. Erlbaum: Hillsdale, New Jersey.
- Faries, J.M., & Reiser, B.J. (1988) Access and use of previous solutions in a problem solving situation. In *the proceedings of the tenth annual meeting of the Cognitive Science Society*. Montreal, Quebec.
- Fong, G.T., & Nisbett (1990). Immediate and delayed transfer of training effects in statistical reasoning. Manuscript submitted for publication.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou, & A. Ortony (Eds.), *Similarity and Analogical reasoning*. Cambridge, Cambridge University Press.
- Gick, M.L., & Holyoak, K.J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.
- Holland, J., Holyoak, K., Nisbett, R.E., & Thagard, P. (1986). *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press.
- Holyoak, K. J. & Thagard, P.J. (1989). A computational model of analogical problem solving. In S. Vosniadou, & A. Ortony (Eds.), *Similarity and Analogical reasoning*. Cambridge, Cambridge University Press.
- Jacob, F., & Monod, J. (1961). Genetic regulatory mechanisms in the synthesis of proteins. *Journal of Molecular Biology*, 3, 318-356.
- Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science*. 12, 1-48.
- Klahr, D., Dunbar, K., & Fay. (1990). Designing good experiments to test "bad" hypotheses. In J. Shrager, & P. Langley (Eds.), *Computational models of discovery and theory formation*. Ann Arbor, MI: Morgan Kaufmann.
- Langley, P, Simon, H.A., Bradshaw, G.L., & Zytkow, J. (1987). *Scientific Discovery: Computational explorations of the creative processes*. Cambridge, MA: MIT press.
- Reber, A. S. (1989) Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, 118, 219-235.
- Ross, B.H. (1989). Reminders in learning and instruction. In S. Vosniadou, & A. Ortony (Eds.), *Similarity and Analogical reasoning*. Cambridge, Cambridge University Press.

Notes

This research was supported by grant number OGP0037356 from the National Sciences and Engineering Council Canada to Kevin Dunbar and a summer NSERC research award to Christian D. Schunn.

Correspondence can be addressed to: Kevin Dunbar, Department of Psychology, McGill University, Montreal, Quebec, Canada H3A 1B1
Email: DUNBAR@HEBB.PSYCH.MCGILL.CA