

Priming, Analogy, and Awareness in Complex Reasoning

Christian D. Schunn

Carnegie Mellon University

Kevin Dunbar

McGill University

AS of 2001

Now at LRDC University of Pittsburgh

Now at Dartmouth College NH

Running Head: PRIMING, ANALOGY, AND AWARENESS IN REASONING

Address all correspondence to:

Kevin Dunbar, Dept. of Psychological & Brain Sciences, Dartmouth College
Hanover, NH 03755

Abstract

The mechanisms by which a concept used in solving one complex task can influence performance on another complex task were investigated. We tested the hypothesis that even when subjects do not spontaneously make an analogy between two domains, knowledge of one domain can still spontaneously influence reasoning about the other domain via the mechanism of priming. Four groups of subjects (two experimental and two control) were given a simulated biochemistry problem on Day 1 and a simulated molecular genetics problem on Day 2. For the two experimental groups, the solution to the biochemistry problem involved inhibition. For the two control groups, the solution did not involve inhibition. On Day 2, all subjects received the same version of the molecular genetics problem in which the solution involved the concept of inhibition. Subjects in the experimental conditions were more likely to attain the correct answer, to propose inhibition, and to propose inhibition early in the problem solving session, than subjects in the control conditions. However, subjects in the experimental conditions made no reference to the biochemistry problem either in their verbal protocols, or a post-task questionnaire. The results are interpreted as demonstrating that an implicit process –priming– can make old knowledge available for current problem solving.

Priming, Analogy and Awareness in Complex Reasoning

Much research on higher level cognition, such as that involved in reasoning and problem solving, has focused on the use of highly structured knowledge and how this knowledge is generated, retrieved, and used. This research has produced rich theories of reasoning, problem solving, and analogy (e.g., Holland, Holyoak, Thagard & Nisbett, 1986; Newell & Simon, 1972; Gentner, 1989). However, the way in which lower-level cognitive processes (e.g., priming, memory activation, attention) are used during higher-level cognition has received little attention in recent research. While some of the cognitive mechanisms that are a cornerstone in lower-level cognitive research have been incorporated into models of thinking and reasoning (e.g., spreading activation: Anderson, 1982, 1993; Holyoak & Thagard, 1989) we still know relatively little about the role that lower-level cognitive processes may play in higher level thinking. A fundamental goal of the research reported in this paper is to explore the link between lower-level cognitive processes and higher-level reasoning. We will address this issue by focusing on the lower-level cognitive mechanisms that enable a concept used in solving one problem to influence performance on a new problem in a different domain.

The issue of how old knowledge can be applied to new situations has been one of the most intensely investigated issues in higher-level cognitive research over the past 15 years. Much of this research has focused on analogy. There is now a considerable literature on how knowledge of one domain, the “source”, can be used—by drawing an analogy—to help solve a problem in a new domain, the “target.” For example, recent discussions of building a better Internet have drawn analogies from the highway building program of the 1950s (the source) to the current problem of rebuilding the Internet (the target).

Researchers have argued that there are two major phases involved in making an analogy: access and use (cf. Ross, 1989a; Gentner & Forbus, 1991). In the access phase, a source analog or abstract concept is selected or accessed. One class of access mechanisms are classic memory retrieval mechanisms in which subjects actively search memory for a source that has similar sets of properties and/or goals to those of the current problem (e.g., Gentner & Forbus, 1991; Gick & Holyoak, 1980; Holyoak & Koh, 1987; Ross, 1984, 1989a). In the use phase, the selected source analog or abstract concept is applied to the current context. Holyoak and his colleagues have argued that subjects map over problems that have similar goals and underlying structures (Holyoak, 1985; Thagard & Holyoak, 1989). Gentner (1983, 1989) has argued that subjects map over entire relational structures. Overall, this research has generated finely detailed accounts of the higher level mechanisms underlying the use of analogies and has led to the development of a number of computational models of analogical reasoning that can predict performance of subjects under a wide variety of conditions (e.g., Gentner, 1989; Gentner & Forbus, 1991; Halford, 1992, Holyoak & Thagard, 1989).

While such models of analogical reasoning have successfully accounted for a wide variety of performance by subjects, both the authors themselves, and various commentators on analogy have noted that analogical reasoning has two main drawbacks as an all-purpose mechanism by which old knowledge can affect the acquisition of new knowledge. First, even when good analogies are available subjects frequently fail to use analogies (e.g., Gentner & Landers, 1985; Gick & Holyoak, 1980). Second, analogy is computationally expensive, demanding much working memory capacity and extensive mental operations (cf. Halford, 1992; Wharton, Holyoak, Downing, Lange, Wickens, & Melz, 1994). The former point has been the object of intense research over the last

decade and researchers have identified a large number of constraints on when and how analogical reasoning will occur. Briefly, this research has shown that subjects will only try to make analogies when certain relationships exist between the source and the target. For analogical reasoning to be attempted, the source and the target must share one or more of the following characteristics; similar surface features, similar goals, and the same underlying systems of relations (e.g., Catrambone & Holyoak, 1989; Gentner, 1989; Gentner, 1989; Holyoak & Koh, 1987; Ross, 1989a; Wharton et al. 1994). What is puzzling is that analogy appears to be used under very limited circumstances, yet people nonetheless benefit from having experience with similar knowledge.

In order to address this issue, we will distinguish between explicit and implicit mechanisms of accessing source analogs: Explicit access processes are those in which subjects actively search through previous problem solving episodes or concepts that can be used to solve the current problem. Implicit processes are those in which subjects are unaware of the mechanisms that lead to a particular concept or analog becoming available—often because such processes are automatic. This distinction is important, because it suggests that there are mechanisms—other than explicit analogy—by which source knowledge can become available in complex reasoning. In this paper, we will focus on priming as an implicit mechanism for making source knowledge available. We investigate the hypothesis that while it may be difficult for subjects to explicitly access information needed to generate an analogy, many of the implicit processes that are involved in analogical reasoning can still work even if an explicit analogy is not made.

One implicit process that has the potential to have an effect on reasoning and transfer is priming. When subjects are re-exposed to a pre-existing piece of information, be it a word or a concept, this word or concept becomes primed. Priming makes items more readily available and faster to respond to (e.g., Graf, Shimamura, & Squire, 1985; Neely, 1979; Tulving & Schacter, 1990). It is thought that primed items have lower thresholds than unprimed items (cf. Anderson, 1983; 1993, Neely 1979; Roediger & McDermott, 1993). Previous research has distinguished among various forms of priming (Roediger & Challis, 1992; Blaxton, 1989; Flores d'Arcais, Schreuder, & Glazenborg, 1985). We focus on priming of concepts, in which some re-exposure to an existing concept, in the form of the concept itself, a synonym, conceptual coordinate, or otherwise conceptually related item makes the concept more readily accessible. Priming of concepts, while less commonly studied than more perceptual repetition priming effects, has been demonstrated in many domains, including political reasoning (Young, Thomsen, Borgida, Sullivan, & Aldrich, 1991), generating anagram solutions (White, 1988), trait judgments (Smith & Branscombe, 1988), category production (Rappold & Hashtroudi, 1991), and free recall (Roediger & Challis, 1992). Furthermore, priming of concepts has been found to last over extended periods of time (Becker, Moscovitch, Behrmann, 1994; Rappold & Hashtroudi, 1991; Srull & Wyer, 1979). Hamann (1990) found that priming of category exemplars decayed much less rapidly when study involved a semantic task than when study involved a perceptual task.

The priming of concepts (henceforth called priming) has been used to investigate text comprehension (e.g., Seifert, McKoon, Abelson, & Ratcliff, 1986; Wharton & Lange, 1994), but so far, priming has not been investigated in complex reasoning and problem solving. Our goal is to determine whether priming has effects on the higher level reasoning processes involved in complex problem solving. We focus on priming because it is regarded as a fundamental mechanism involved in cognition, and it helps solve the two above-mentioned problems with analogy—it is computationally inexpensive and

does not necessarily require a complex mental structure to occur.¹ The argument presented here is that a concept could be primed in one problem solving context, and hence make that concept more available in another context. A concept that is primed should be more readily accessed and available for use in later problem solving than an unprimed concept. The prediction that we are making is that even though the cognitive system may not have produced an analogy, priming of concepts can still occur and can have an impact on performance in another task.

To investigate these issues, a two day experiment was constructed in which subjects worked on two related tasks on separate days. A context was chosen whereby a concept could be primed. The concept that we used was inhibition. Subjects had been previously taught the concept of inhibition in four of their compulsory psychology courses, thus it was expected that this concept would be one that could be primed by solving a problem that involved inhibition. On the first day, subjects were given a very simple source problem (the virus problem). In this problem, subjects conducted simulated biochemical experiments on a computer. Using the simulated laboratory, subjects discovered that viruses are dormant because they are inhibited. It was expected that this manipulation would prime the concept of inhibition. The next day, subjects returned for another, ostensibly unrelated, task (the genetics problem). In the genetics problem subjects had to discover how a set of genes are controlled. None of the surface features of the source or target problems were the same. The two tasks had different interfaces, and were presented as unrelated tasks. Thus, no explicit mappings between the two problems were expected.² However, the mechanism of control for both the virus and the genetics task involved the same fundamental concept — inhibition. Thus, it was predicted that subjects who worked on the virus problem would be more likely to discover the concept of inhibition in the genetics problem than subjects in control conditions who were not exposed to the concept of inhibition on the previous day. As the two tasks were very different, subjects were not expected to make an explicit analogy between the two tasks. We hypothesized that if transfer occurred it would occur via the implicit mechanism of priming.

An important question that our hypothesis raises is how would one know whether priming is at work? There are three criteria that we see as being important. The first is that of awareness. Priming is generally regarded as an implicit process that occurs outside awareness, and a vast literature has repeatedly demonstrated that subjects are usually unaware that priming has occurred (See Roediger & McDermott, 1993 for a recent review). Second, reaction time studies have demonstrated that primed concepts are responded to faster than unprimed concepts (e.g., Becker, Moscovitch, Behrmann, 1994; White, 1988). The third criterion is one of retrieval of episodic or contextual information. If a concept is primed, rather than a specific instantiation of that concept retrieved, then we would expect subjects not to recall contextual information. However, if subjects are retrieving specific instances of a concept, then we would expect them to recall contextual information about the concept. Thus, if priming is to occur in problem solving, then subjects should be unaware of the priming, should retrieve concepts more quickly when they are primed, and not retrieve contextual information about the last use of the concept.

The way that we investigated awareness was by collecting verbal protocols from subjects while they performed a task. Verbal protocols consist of subjects continuously saying what they are thinking while they work on a problem. It is thought that when generating verbal protocols, subjects are verbally stating the contents of working memory (cf. Ericsson & Simon, 1980; 1993). Thus, verbal protocols are thought to provide a

running record of the contents of working memory. If subjects are explicitly reminded of, and are hence aware of, one problem while working on another, this should show up in the verbal protocols since explicit reminding would involve working memory. If there is no reference to the source problem in the verbal protocols, but there is improved performance on the target problem, this would suggest that an implicit mechanism, such as priming, is at work.

An alternative mechanism to priming that might produce similar results is learning. We distinguish three forms of learning: learning the new concept of inhibition, learning new facts about the inhibition concept, and learning new procedures useful to discovering inhibitory regulatory mechanisms. Learning the concept of inhibition is an unlikely candidate, since subjects already have the concept of inhibition. However, this potential mechanism can be tested by studying the way in which subjects discover inhibition in the day 1 (virus) task. If the concept of inhibition is being learned *de novo*, then subjects should have to struggle with the task, and propose the inhibition concept after an extended period of time. If, however, the concept of inhibition is already known to the subjects, then they should be able to propose inhibition immediately after uncovering evidence in favor of such a regulatory mechanism.

The second type of learning that could occur is learning new facts about a pre-existing concept. Given that subjects do have the concept of inhibition before they participate in the day 1 (virus) task, part of what they learn in solving the virus task is to modify their concept of inhibition for this current problem. Thus, in addition to retrieving the concept of inhibition, they instantiate this concept in a new context, the virus task. Specifically, they are learning new facts involving the inhibition concept—enzymes can inhibit virus reproduction. If subjects are retrieving these new facts that they have learned, then it should be possible to distinguish between priming and new fact learning using verbal protocols. If, on day 1, subjects learned to modify the concept of inhibition, then we would expect references to the day 1 (virus) problem to occur in the day 2 protocols. Subjects should be retrieving specific information about the virus problem to help solve the genetics problem, since the new facts would be specifically about the virus task. If however, priming of the concept of inhibition occurred in day 1, then subjects need not retrieve contextual information while solving the day 2 genetics problem. Thus, we expect that no references to the day 1 virus problem would occur in the day 2 genetics problem.

The third form of learning that could occur on the day 1 virus problem is procedural learning: Subjects may learn new procedures in solving the virus task that can be used to discover inhibition. Since this knowledge is procedural in form, it need not appear in the verbal protocols (cf. Ericsson & Simon, 1980, 1993; Singley & Anderson, 1987), and so is not distinguishable from priming using verbal protocols. However, the presence of procedural learning can be studied by looking at the changes in the procedures that subjects use when they are working on the day 2 genetics task. In particular, if procedural learning is a dominant factor in day 1, then subjects should learn specific experimental procedures useful for discovering inhibition, and should use these procedures on day 2. In contrast, priming would not be expected to produce any such procedural changes. This hypothesis can be investigated further by having an additional control condition of subjects that are told the solution to the virus problem, rather than having to discover inhibition through experimentation. If procedural learning is important, then only the group that discovers the solution through experimentation would benefit from working on the virus task. If, however, priming is the main mechanism of

transfer, then both groups should benefit from knowledge of the virus task. Such a manipulation was included in the experiment.

In summary, this experiment involves subjects working on the virus and genetics problems on two consecutive days. Both problems involved the same underlying mechanism—inhibition. It was expected that solving the virus problem would have an effect on solving the subsequent genetics problem. If this happens, then the question is: at what level does the transfer occur? In particular, will subjects discover inhibition through priming, learning, or through explicit analogical reasoning to the virus problem?³

METHOD

Subjects. Eighty McGill undergraduates were paid to participate in the experiment—twenty in each of four groups. All subjects had taken an introductory biology course. None of the subjects were familiar with the mechanisms of gene regulation or virus reproduction, as these topics were not covered in their compulsory biology course. Since inhibition is a very basic concept, it is likely that all the subjects already had knowledge of inhibition from other sources. However, all subjects had taken psychology courses that had introduced them to the concept of inhibition, providing a recent, concrete example of this knowledge. Specifically, all subjects took introductory courses in: (1) Learning, where they were given lectures on Pavlovian inhibition, inhibitory stimulus generalization, and the general concept of inhibition; (2) Motivation, where they encounter inhibitory neural mechanisms; (3) Perception, where they encountered inhibition of neural circuits, lateral inhibition, inhibition in hearing, and on the skin; (4) Cognition, where they learned of inhibitory connections in connectionism, inhibitory priming, and proactive interference. Thus, using subjects from this pool ensured that they all had recent multiple exposures to the concept of inhibition and had approximately the same knowledge of the concept of inhibition.

Apparatus. Subjects conducted experiments on a Macintosh II computer. The experimental interfaces were written in cT (Sherwood & Sherwood, 1988).

Procedure. Day 1, the virus task. There were two phases of the virus task; an Introductory phase and a Discovery phase. In the Introductory phase, subjects were taught some basic facts about biochemistry. Then they were shown that viruses can be made active or dormant through controlling enzymes. They were introduced to the methods that they would use for investigating the virus problem. In the Discovery phase, subjects were told that sometimes viruses are dormant and do not reproduce upon entering a cell. Subjects were also told that scientists suspected that a certain enzyme controls the activity of the viruses. Three different enzymes, labeled G, K, and Q, were given as possible causes of viruses becoming active or dormant. Subjects were then given the task of discovering how the enzymes controlled the virus activity.

Subjects designed experiments by using a program that simulated the action of enzymes and the number of viruses present. Subjects chose whether to have active or inactive viruses. They then chose whether to view enzyme activity in the cell or to examine the effects of adding an enzyme to the cell. Next, they chose which enzyme to investigate (i.e., which enzyme to view or add depending on the type of experiment). If subjects chose "VIEW," then they saw a graph over time of the number of viruses that were present as well as the amount of enzyme selected. For example, in Figure 1a, the subjects saw that as the number of viruses increased the amount of enzyme G also increased. If, subjects chose "ADD", they again saw the number of viruses present over time, but with an enzyme added. For example, in figure 1b, the subjects saw that when

enzyme Q was added the number of viruses decreased. Subjects could discover which enzyme controlled virus activity by comparing the effects of adding the different enzymes to active and inactive viruses.

Insert Figure 1 about here

The number of distinct experiments that subjects could conduct was very small. The three experimental choices; active/inactive, view/add, and enzyme (G, K, or Q) produced twelve possible experiments. This set was purposely made very small for two reasons: First, most subjects would be able to discover the solution. Second, the knowledge acquired by all subjects would be fairly uniform.

Conditions. We used a between subjects design in which subjects were randomly assigned to one of four conditions. In all four conditions two of the enzymes played no causal role (enzymes G and K); it was the role of the third enzyme (enzyme Q) that was manipulated. Subjects had to discover different mechanisms in the different conditions. Conditions 1 and 2 were control conditions where the subjects did not discover inhibition. Conditions 3 and 4 were experimental conditions where subjects discovered inhibition.

1. In the Correlated condition, enzyme Q was correlated with virus growth. Subjects had to discover that Q is not causally related to virus reproduction.

2. In the Positive regulation condition, enzyme Q switched on or activated virus growth. Subjects had to discover this mechanism.

3. In the Negative regulation condition, enzyme Q inhibited virus reproduction. Subjects had to conduct experiments to discover this mechanism.

4. In the Story condition, enzyme Q inhibited virus reproduction. In this condition subjects did not conduct experiments; they were told that inhibition was the mechanism of control. This condition was used as a further control condition to investigate the hypothesis that procedural learning of strategies for discovering inhibition might be at work in the Negative regulation condition. It may be the case that subjects in the Negative regulation condition acquire procedural knowledge specific to discovering inhibition. In the Story condition, subjects who are merely told the solution to the problem would not discover this procedural knowledge. By comparing the performance of the subjects in the Negative and Story conditions, we were able to investigate this hypothesis.

Prior to beginning the Discovery phase, subjects were instructed on how to give a verbal protocol. We used the methods of Ericsson and Simon (1982; 1993) for obtaining verbal protocols: Subjects were told to say everything that they were thinking while performing the task. Subjects gave concurrent verbal protocols throughout the discovery phase. Subjects conducted experiments until they thought they had discovered how the enzymes controlled the virus activity. No feedback on the correctness of their final solutions was given. At the end of Day 1 subjects were asked to state the role of each of the three enzymes.

Day 2, the genetics task. The same form of the genetics task was used in all conditions. This task is the same task used by Dunbar (1993) and a more extensive account of this task can be found there.⁴ The genetics task consisted of two main phases. In the Introductory phase, subjects were taught some basic facts about molecular biology and were introduced to methods used for investigating genetic regulation using a

simulated molecular genetics laboratory on a computer. In the Discovery phase, subjects were told that E. coli need glucose to live, and that the E. coli obtain glucose by breaking down lactose into glucose. They were told that a set of genes on the E. coli chromosome secretes a substance called betagalactosidase (henceforth called beta-gal) that breaks down lactose. Then they were given the following problem: how is the production of beta-gal controlled? The subjects were given two findings. First, the beta-gal producing genes secrete beta-gal when there is lactose present. Second, they do not secrete beta-gal when there is no lactose. They were told that genes I, P, and O were thought to be involved in controlling the production of betagalactosidase.

The subjects conducted simulated experiments on a set of genes found in E. coli. As can be seen in Figure 2a, the E. coli were simulated in the following way: The cell wall of E. coli was represented as large ovals, and the genes as a series of boxes (Note that the computer interface for the simulated genetics experiment was totally different from that in the virus problem and subjects should not be reminded of the virus problem by the genetics interface). There were two kinds of genes: producer genes (unlabeled boxes) that produce the beta-gal which breaks down Lactose, and controller genes (labeled boxes) that control the producer genes. The lactose was represented with smaller squares, and the beta-gal with black dots (see Figure 2b). The I, P, O genes controlled the beta-gal genes such that beta-gal was produced only when Lactose was present. The subjects' task was to determine how the I, P, and O genes were involved in controlling the beta-gal genes.

Insert Figure 2 about here

Two main experimental methods were available to the subjects. The first method was mutating one of the controller genes and observing the effect on the output of the beta-gal genes. If the mutation changes beta-gal output (e.g., by shutting output off), then that controller gene is implicated. This kind of experiment is called a haploid experiment. The second method involves using diploid cells in which there are two sets of genes: one set of genes that have the producer genes and the controller genes, and the other set of genes that just have the controller genes. By producing different combinations of mutations on the two sets of genes subjects could determine whether, and how, a particular gene was involved in controlling the output.

The methods were described to the subjects in the context of an example regulation problem: Subjects were shown how certain genes (labeled A, B, & C) control the activities of other genes by switching them on when there is a nutrient present. This is an example of positive regulation—the "A" gene is switched on when a nutrient is present. The "A" gene then releases substances that instruct other genes to secrete enzymes that will breakdown the nutrient. The number of genes and the experimental interface used in the example are identical to that in the Discovery phase. The only difference is that the genes have different labels in the Discovery phase (A, B, & C rather than I, P, & O) and the underlying mechanism is one of Activation — rather than Inhibition. Following the results of Dunbar (1993) it was expected that subjects would map the A, B, C example onto the genetics problem and that subjects' initial hypotheses would be that the I, P, O genes turn on beta-gal production. This manipulation was used

to ensure that subjects would all begin the genetics problem with an incorrect activation hypothesis.

The A,B,C example had two main functions. First, we wished to present an example of how to conduct experiments in this domain. Second, we wished to make the task suitably difficult such that subjects would not be performing at ceiling in the base conditions. This was achieved by biasing subjects toward activation hypotheses. From previous research (Dunbar, 1993), we knew that subjects would all start with activation hypotheses and only begin to consider alternate hypotheses, such as inhibition, after much experimentation. It is the effect of priming of the concept of inhibition on the search for alternate hypotheses that was the focus of this experiment.

Subjects conducted experiments until they thought they had discovered how the I, P, O genes controlled the producer genes. The genetics problem had a number of levels to the answer. First subjects had to discover that the I and O genes were inhibitors. A more difficult challenge was to discover how the I and O genes interact and that they work in different ways. The O gene physically blocks beta-gal production and must be adjacent to the beta genes to inhibit them. The I gene does not have to be on the same chromosome as the beta-gal genes because it sends out an inhibitor that binds to the O gene. Subjects can discover this through the use of diploid experiments. Based on Dunbar's (1993) results, we expected few subjects to discover these subtle aspects of genetic control, and hence the analyses of results will not focus on this level. Experimentation consisted of three components: designing an experiment, watching the experiment run, and viewing a summary table of past experiments. There were 120 possible experiments that could be conducted: subjects could choose from six different amounts of nutrient. Only one gene could be mutated in each chromosome at a time.⁵ Subjects had to discover that the I and O genes negatively regulate (i.e., inhibit) the activity of the enzyme producing genes, whereas the P gene plays no role.

The virus (Day 1) and genetics (Day 2) tasks were different along many dimensions. At the level of domain, the virus task was presented as a biochemistry task, whereas the genetics task was presented as a molecular genetics task. At the level of data presentation, the virus task involved graphs and the genetics task involved tables. At the level of the interface, the subjects selected options with buttons in the virus task and pull-down menus in the genetics task. Finally, at the level of to-be-explained phenomena, the virus task involved a decrease in output with the addition of an object, whereas the genetics task involved an increase in output with the removal of an object. However, the underlying essential concept in both tasks is inhibition. In the virus task, the Q enzyme inhibits virus growth. In the genetics task, the I and O genes each inhibit beta-gal output.

Post-genetics task questionnaire. In addition to collecting verbal protocols, we also questioned subjects to further assess whether they had engaged in explicit use of the virus task while working on the genetics problem. Subjects were given the following three questions: "Did you think about yesterday's experiment at any point today?", "Do you see any similarities between yesterday's and today's experiments?" and "Did you think yesterday's experiment helped you today?"

RESULTS AND DISCUSSION

Virus Task Results

As stated in the introduction, the virus task was expected to be relatively simple. This was the case: In the Correlated condition, all 20 (100%) subjects discovered that there was no causal link between Q and virus growth, in the Positive condition, 19 (95%) subjects discovered that Q was an activator, and in the Negative condition, 18 (90%)

subjects discovered that Q was an inhibitor. Further, none of the subjects in the Positive and Correlated conditions proposed inhibition in the virus task. Thus, the virus manipulations were successful.

To investigate whether subjects already had the concept of inhibition (i.e., that it was a pre-existing concept that could be primed), the verbal protocols of the Negative condition were analyzed for first and subsequent mentions of the concept of inhibition. In the Negative virus task, there was one crucial experiment that was strongly suggestive of inhibition (adding enzyme Q to an active virus). Fourteen (70%) subjects mentioned inhibition immediately after conducting this crucial experiment. Three (15%) subjects mentioned inhibition before conducting this crucial experiment (i.e., without any evidence for inhibition). Two (10%) subjects mentioned inhibition within two experiments beyond this crucial experiment. The remaining subject never conducted this crucial experiment, and never mentioned inhibition. The finding that subjects were able to propose the concept of inhibition so quickly indicates that they already had the concept of inhibition.

Similar analyses were conducted for the concept of activation in the Positive condition. In the Positive the virus task, there were two critical experiments each of which demonstrated the presence of a positive regulation mechanism (adding enzyme Q to either virus type, producing greater virus growth). The verbal protocols were coded for first mention of activation. Two (10%) subjects mentioned activation prior to conducting a crucial experiment. Sixteen (80%) subjects mentioned activation immediately after conducting a crucial experiment. One (5%) subject mentioned activation in the experiment following their first crucial experiment. The remaining subject never mentioned activation. Thus, for the majority of the subjects, activation was a readily available concept. Overall, these results indicate that subjects did indeed possess the concepts of inhibition and activation prior to participating in the virus (day 1) experiment.

We can now turn to the effect of the virus problem on subjects' performance in the genetics task. Only subjects who correctly discovered the role of the Q enzyme in the virus condition were included in the analyses of the genetics task. Thus, analyses were conducted on 20 subjects in the Correlated, 18 in the Positive, and 19 in the Negative conditions.

Genetics Task Results

To determine the effect of knowledge acquired in the virus task on performance in the genetics task, five types of analyses were conducted. First, the analyses of overall success of subjects in each condition were conducted to determine whether there were any differences in performance between subjects in the four groups. Second, to assess whether there were any procedural learning effects, an analysis was conducted on the kinds of experiments performed by subjects. Third, an analysis of subjects' verbal protocols was conducted to determine whether subjects were explicitly mapping from the virus problem to the genetics problem. Fourth, the results of the post-genetics task questionnaire were analyzed to provide converging evidence for the verbal protocol findings. Fifth, the point at which subjects first proposed inhibition was analyzed to determine further whether the concept of inhibition was primed by the virus task.

1. Overall Performance on the Genetics Task

A primary measure of how experience with the virus problem affected performance on the genetics task was the number of subjects that solved the genetics task. To assess performance on the genetics task, the subjects' statement of their final

hypotheses were coded for the role that each gene played. The subjects' final hypotheses were coded by two independent coders, producing an inter-coder reliability of 97%. Differences were resolved through discussion. A subject's hypothesis was coded as correct if the roles of all three genes were correct (P=No role, I=Inhibitor, O=Inhibitor). Subjects in the Negative and Story conditions were more likely to reach the correct solution (8 and 8 subjects respectively) than in the Correlated, and Positive conditions (3 and 5 subjects respectively) $\chi^2(1) = 3.6, p < .05$. As another control, the performance in the genetics task in this study can be compared to the results found by Dunbar (1993, Experiment 1), in which subjects from the same subject pool were given the same genetics task without being exposed to the virus task (see Figure 3). Dunbar found that 20% (4 of 20) subjects reached the correct solution. This comparison suggests that the virus task had a facilitatory effect in the Negative and Story conditions (mean of 41% correct), and no effect (of facilitation or interference) in the Correlated and Positive conditions (mean of 21% correct).

Insert Figure 3 about here

The primary focus of the experiment was on whether discovering inhibition in the virus task led to a greater likelihood of proposing inhibition in the genetics task. To test for such an effect, a further analysis was conducted on the frequency of inhibitory hypotheses in each of the four conditions. The analysis was conducted on the number of inhibitory hypotheses in the subjects' final hypotheses for the I, O and P genes (see Figure 4). There was a main effect for condition ($F(3,73) = 3.3, p < .03$). Post-hoc analyses revealed that there were more inhibitory hypotheses in the two inhibitory conditions (Negative and Story) than in the non-inhibitory conditions (Scheffé $F(3,73) = 8.4, p < .05$). Thus, experience with inhibition in the virus task had a significant effect on proposing inhibition in the genetics task. A similar analysis was conducted on the number activation hypotheses in the subjects' final hypotheses for the I, O, and P genes (see Figure 4). The effect of condition was not significant ($F(3,73) < 1, p > .7$).⁹

Insert Figure 4 about here

2. Evidence of Procedural Transfer

To assess whether the manipulations produced any changes in the way in which subjects conducted experiments in the genetics task, nine measures of the experimentation process were analyzed. There were no effects of condition on number of experiments conducted ($F(3,73) = 1.3, p > .25$), time on task ($F(3,73) = 1.8, p > .15$), or time spent per experiment ($F(3,73) < 1$). A further three of these measures focused on the particular kinds of experiments which are especially important to the genetics task (see Table 1). First, an ANOVA was conducted on the proportion of diploid experiments with a normal female gene and a mutant male gene (Normal/mutant). Such experiments cannot discriminate between any mechanisms, and so generally indicate confusion in the subject. The effect of condition on proportion of Normal/mutant experiments was nonsignificant

($F(3,73) = 1.6, p > .2$). Second, an ANOVA conducted on the proportion of zero lactose experiments, which are very useful for identifying negative regulation,⁷ found no effects of condition ($F(3,73) = 1.7, p > .15$). Third, an ANOVA was conducted on the proportion of crucial experiments conducted by each subject. To arrive at the correct hypothesis with sufficient evidence, subjects had to conduct five kinds of experiments, called crucial experiments. These crucial experiments were the three different haploid mutants (I-, P-, and O-) and the diploid mutants involving the two causally implicated genes (I-/Normal & O-/Normal).⁸ The effect of condition on this measure was nonsignificant ($F(3,73) < 1$).

Insert Table 1 about here

The final three measures focused on the quality and breadth of the experiments conducted. First, an ANOVA was conducted on the mean number of features varied from one experiment to the next. Feature variation has often been used as a measure of strategy use in scientific experimentation (e.g., Bruner, Goodnow, & Austin, 1956; Kuhn & Phelps, 1982; Tukey, 1986). Varying multiple features between experiments confounds sources of differences, whereas varying only one feature at a time provides unambiguous tests of the effects of particular features. The effect of condition on the mean number of features varied was nonsignificant ($F(3,73) < 1$). Second, an ANOVA was conducted on an index of the degree to which subjects explored each of the factors in the task (i.e., types of haploid mutations, types of male diploid mutations, types of female diploid mutations, and amount of nutrient). This measure, called the dimension search index, was the sum of the number of levels of each factor that the subject explored, and ranged from 4 to 18. There were no effects of condition on this measure ($F(3,73) < 1$).⁹ Finally, an ANOVA was conducted on the proportion of the 20 possible gene combinations that were tried by each subject, another measure of breadth of experimentation (c.f. Dunbar, 1993; Klahr & Dunbar, 1988; Klahr, Fay, & Dunbar, 1993). There were no effects of condition on this measure either ($F(3,73) < 1$).

In sum, there were no indications that there were differences in the experimentation process across conditions. Thus, it is unlikely that the subjects obtained any special new knowledge in the virus task that influenced the way in which they conducted experiments. We now turn to an analysis of the verbal protocols to determine whether subjects generated inhibition hypotheses as a result of explicit reminding and analogical mapping.

3. References to the Virus Problem in the Verbal Protocols

The verbal protocols from the genetics task were coded for any references that subjects made to other sources of knowledge while they were working on the genetics problem. If subjects were using explicit analogy, then references to the virus task should occur in their verbal protocols. We divided references to prior sources of knowledge into three types (a) references to the ABC example that the subjects were shown at the beginning of the genetics task, (b) references to the virus problem, and (c) references to any other knowledge.

Insert Table 2 about here

As can be seen from Table 2, a total of 17 references were made to other sources of knowledge. However, only two subjects made references to the virus task, and these subjects were not in the Negative or Story conditions. This indicates that the reason that more subjects proposed inhibition in the Negative and Story conditions was not because these subjects were explicitly mapping from the virus problem to the genetics problem, or that the genetics task reminded them of the virus problem. Below, we give two examples of each of the three types of reference that subjects made while working on the genetics task.

References to the ABC example. References to the ABC problem were the most explicit use of mapping from one source of knowledge to the genetics tasks. Subjects' references to the ABC example were where the subjects mapped A, B, C over to the I, O, P genes. For example, subject AD5 in the negative condition said:

“The thing is I have to find... what is the controller gene... And the case- remember when we had A missing- OK, let me just... OK- B, C, we added A, B, C- and when we- and there was a reaction. Said that A was the controlling gene. So I want to find the case- It seems like, all three are reacting...”

Subject LK15 in the Story condition said:

“The absence of P is the same as the presence of A... in the.. in the previous experiment...”

References to the virus task. No subject made any explicit reference to any of the details of, or mechanisms involved in, the virus problem. The two examples provided below were the only references to the virus problem in 80 protocols. Even these references were very general and not very helpful for the subjects. Subject LM5 stated that;

“Oh, gee... and I was thinking of taking biology yesterday. Now today I don't think I will...”

The other subject (NA12) stated that

“Ha, ha, this is a tougher one than yesterday...”

References to other knowledge. Like references to the virus problem, references to other sources of knowledge were very general and no explicit mappings were made to a source problem that could be mapped over to the genetics problem. Subject EF14 in the positive condition said;

“Its just like ‘mastermind’ I have to figure out what each gene's doing?”

Subject AK1 in the correlated condition stated

“This is worse than a final”

Overall, the examination of the protocols provides little evidence for any form of explicit analogical mapping or reminding from the virus problem to the genetics problem. The finding that all but two of the subjects failed to mention the virus task suggests that mechanisms other than explicit analogy were responsible for the better performance of the inhibitory condition subjects over the non-inhibitory condition subjects.

4. Post-Genetics Task Questions

To further address the issue of how performance on the genetics task was influenced by the virus task, three questions were given to subjects on completion of the genetics task. The goal of asking the questions was to uncover the relationships that the subjects saw between the virus and genetics tasks. These questions were open-ended and allowed the subjects room to state any relationships or similarity that they saw between the two tasks.

A. “During the experiment did you at any time think of yesterdays experiment?”¹⁰

There was a 92% agreement between coders for this question. Only 37% of subjects reported that they had thought about the previous day’s experiment at any time. The percentage of subjects reporting that they had thought of yesterday’s experiment were 35, 41, 57, and 20 in the Correlated, Positive, Negative, and Story conditions, respectively.

The interaction with condition was not significant ($\chi^2(3) = 5.1, p > .15$). Furthermore, 13% of the subjects spontaneously responded that they had no memory for the details of the previous day’s experiment (There was a 92% intercoder agreement for this response). Surprisingly, in the Story and Negative conditions, these subjects were more likely to have solved the genetics task ($r = .3, \chi^2(1) = 6.9, p < .01; r = .26, \chi^2(1) = 2.6, p > .1$).¹¹

Sixty-one percent of the subjects responding that they had thought of the virus task during the genetics task (N=14) specified a particular point in time at which they had thought of the virus task. All of these subjects mentioned the same point in time: the introduction phase to the genetics task. No other particular point in time was mentioned. Thus, one explanation for why some subjects reported thinking about the virus problem and yet no references appeared in the verbal protocols in the genetics task is that subjects only thought of the virus task during the introduction to the genetics task. This explanation could also reconcile the lack of explicit analogy with the apparent demand characteristics of the study to make some link between the two sessions. Furthermore, this result explains the lack of an association between claiming to have thought about the virus problem and solving the genetics task: a few subjects did think back to the virus task very early in the genetics task, but then discarded it immediately since it did not seem relevant.

B. “Do you see any similarities between yesterday’s and today’s experiments and what are they?” This question was not asked of subjects in the Negative condition, but was asked of subjects in the other three conditions. Responses to this question were coded into five categories: Methods (used elimination, added/subtracted things, etc.), Goals for experiments (goal was to find the causal agent, trying to see what causes things, etc.), Mechanism (same underlying mechanism, e.g., inhibition), Elements (mention elements from either domain as being similar), and None (believes there are no similarities). There was 98% agreement among the two coders on this measure.

Insert Table 3 about here

There was no similarity in mechanism for the Correlated and Positive conditions, and subjects did not see any (see Table 3). The two subjects in the Story condition only noticed the similarity when they were asked about similarity, rather than when working on the problem. These subjects claimed not to have thought about the virus task during the genetics task, and also believed that working on the virus task had not helped working

on the genetics task. Thus, it appears unlikely that subjects in the Story condition explicitly mapped knowledge from the virus task to the genetics task --even the subjects solving both tasks seemed generally unaware of the crucial similarity between the two tasks when asked to explicitly compare the two domains. With regards to elements, subjects in the Story condition noticed similarities of elements (e.g., that Q corresponds to I). When more general similarities were mentioned, subjects in the Positive and Correlated conditions were more likely to state that the general goals and methods used in the virus and genetics tasks were the same.

C. “Do you think that yesterday's experiment helped you in any way?” Intercoder reliability was 96% for this question. While 37% of the subjects thought that experience with the virus task helped during the genetics task, the responses did not interact significantly with condition ($\chi^2(3) = 3.1, p > .3$). Furthermore, there was no association between believing that the virus task had helped and actually solving the genetics task ($\chi^2(1) < 1; r = -.19, \chi^2(1) = 1.3, p > .2$ for the Story and Negative conditions respectively).

Overall, the results of the protocol analyses and post-genetics interview reveal that subjects were not aware that the genetics problem did, or could have, helped them discover the solution to the genetics problem. This indicates that the subjects were not using strategies of explicitly mapping the solution of the virus problem to that of the genetics problem.

5. Point at which Subjects Propose Inhibition

As proposed in the introduction, experience with the virus problem may have primed subjects' existing concepts of inhibition. If this were the case, then subjects in the inhibitory conditions should propose inhibition earlier than subjects in the non-inhibitory conditions. To assess this hypothesis, the verbal protocols from the genetics task were coded for the point at which subjects first mentioned inhibition. The reliability for coding when subjects first mentioned inhibition was 98%. Two categories were used: Immediately (no more than one experiment beyond the first piece of evidence that could indicate inhibition), and Later.¹² As can be seen from Table 4, subjects in the inhibitory conditions were more likely to propose inhibition earlier than subjects in the non-inhibitory conditions ($\chi^2(1) = 3.8, p < .05$). These results are consistent with the hypothesis that the concept of inhibition was primed by the virus task.

Insert Table 4 about here

GENERAL DISCUSSION

The results of this experiment indicate that it is possible for experience with one problem to facilitate performance on another problem without subjects being aware that one problem has helped them solve the other problem. In this experiment, subjects were given the virus problem on Day 1 and the genetics problem on Day 2. For two groups of subjects, the solution to the virus problem involved inhibition. When these subjects returned for the genetics problem on Day 2, more subjects attained the correct answer and more subjects proposed inhibition than in either of the two control conditions. These subjects made no reference to the virus problem either in their verbal protocols, or in a questionnaire administered after the genetics problem was completed. Thus, subjects had

no awareness that experience with the virus problem had an effect on their performance in the genetics problem.

The finding that subjects were not aware that the virus problem helped them indicates that subjects were not using explicit analogical mapping from the virus problem to the genetics problem. If they had been making explicit analogies, then this would have shown up in the verbal protocols and post genetics questionnaire. It should be noted that analogies and reminders were present in the verbal protocols, but the analogies and reminders were not to the virus problem. The existence of analogies and reminders in the protocols indicates that verbal protocols do pick up analogies. Furthermore, verbal protocols have been found to be good indicators of analogical transfer (Ross, 1984; 1989b). Thus, the lack of references to the virus problem in the verbal protocols is not an artifact of using verbal protocols. While it could be argued that lack of a mention of the virus task in the genetics problem is not strong evidence that the virus task was not thought of, we obtained similar results in the post-problem solving questionnaire. This converging evidence, using different dependent measures, leads us to believe that subjects really did not explicitly access the virus problem while working on the genetics problem.

The use of verbal protocols was a critical feature in this experiment. Other researchers have tended to use post-task interviews and questionnaires to assess whether subjects were reminded of a prior problem. Protocols provide a much stronger measure of reminding, as they reflect the on-line contents of working memory. Thus, if we had used only a post-task questionnaire and subjects did not mention the virus task, it could have been the case that subjects were reminded during the task, but while answering the questionnaire failed to remember that they had used the virus problem to help solve the genetics problem. Given that we did collect on-line protocols, this potential ambiguity was avoided; we had access to any potential analogies or reminders during the task. Thus, the use of on-line verbal protocols makes it possible to rule out explicit analogy as the cause of transfer in this task.

Why did subject not engage in explicit analogy? Comparing the genetics task and inhibitory versions of the virus task, the important similarity is at the level of inhibition, a single relation rather than a complete system of relations. At the level of relational systems underlying the solutions, the two tasks are, in fact, very different. The solution to the virus task involves simple inhibition (e.g., enzyme inhibits virus growth). In contrast, the solution to the genetics task involves inhibition of inhibition (e.g., nutrient inhibits the inhibition of enzyme output). Given that similarity was only at one very specific level, the lack of any use of explicit analogy is consistent with the findings of previous research where subjects fail to draw analogies when there are neither surface features, or systems of relations common to the source and target domains (e.g., Clement & Gentner, 1991; Wharton et al, 1994; Halford, 1992). The novel finding presented here is that despite the lack of use of explicit analogy, there is positive transfer from the virus problem to the genetics problem.

The results of this experiment are consistent with the view that the concept of inhibition was primed by the virus task, and hence was more available for subjects when they worked on the genetics problem. In the introduction, we argued that priming was a potentially important mechanism particularly when subjects do not make use of an explicit analogy. A primed concept would be more available for problem solving than unprimed concepts. Given that priming is regarded as an implicit process, priming should occur without awareness. Thus the transfer mechanism of priming is both consistent with

the hypothesis that there was not enough information available for analogical mapping to occur and the finding that subjects had no awareness of the effects of the virus problem on the genetics problem. The priming hypothesis that we have proposed is consistent with previous work on the priming of concepts and extends the range of situations over which priming can occur.

We are calling the positive transfer that we obtained in our experiment priming. However, another way of interpreting the findings is as a form of implicit analogy. That is, some implicit process similar to analogy yet different from priming may have been responsible for the transfer that we found. This may be the case, however, there is currently no existent account of what the mechanisms underlying implicit analogy might be. Regardless of whether the phenomenon is called priming, or implicit analogy, it is clear that the observed effects are different from those typically thought to occur in problem solving transfer studies: explicit analogical transfer.

A potential alternative explanation for the results of this experiment is that the findings are due to learning effects rather than priming. However, there are a number of reasons that lead us to believe that learning is not the cause of the findings: First, subjects already had the concept of inhibition, as demonstrated by the ease with which they discovered inhibition in the virus task. Second, subjects did not mention the virus task during the genetics task. This suggests that subjects were not making use of new facts about inhibition that they learned in the virus task. Third, there were no differences across conditions in the kinds of experiments conducted by subjects, suggesting that the transfer was not caused by learning of new procedures for experimentation. Finally, the subjects in the Story condition also showed transfer, providing further evidence that the source of transfer was not learning particular methods of discovering inhibition.

The issue of learning versus priming is necessarily complex. We have treated learning mechanisms as being distinct from priming mechanisms. Treating priming and learning as different processes makes it possible to derive clear predictions about the types of data that would be expected under either explanation, which we summarized in the previous paragraph. However, it may be the case that priming and learning exist along a continuum much like automatic and controlled processes have been argued to be on a continuum (cf. Cohen, Dunbar & McClelland 1990; MacLeod & Dunbar, 1988). Thus, future experiments might show that concepts can be changed via learning with subjects being unaware of the change. This is a logically possible mechanism that potentially could be used to explain the data in the current experiment. However for such an account to be true, it would be necessary to show that there is a change in subjects concept as a result of the initial experience with the virus task. All the analyses that we have conducted indicate that there was no change in the concept of inhibition in the subjects as a result of their experience with the virus problem. Thus, priming appears to be the most parsimonious explanation for our data.

Another relevant issue to our findings is that subjects in all conditions were biased towards activation hypotheses at the beginning of the genetics task. Thus, what may be happening is that this recent experience with activation may interfere with discovering inhibition. Dunbar (1993) has shown that this is indeed the case: Subjects are biased toward activation and set goals of finding activation and distort evidence to fit activation hypotheses. In the experiment reported here, we were interested in whether subjects could overcome this bias by having the concept of inhibition primed by experience with a prior problem. The finding that subjects in the inhibitory conditions performed better than subjects in the other conditions demonstrates that priming of another hypothesis can

make it more likely for subjects to overcome a bias to a particular hypothesis. This type of finding suggests that when subjects search for new hypotheses, a recently primed concept will be more readily available and will be proposed, even when subjects are biased to another hypothesis. Thus, priming of hypotheses may be an important mechanism underlying hypothesis generation.

The approach taken in this paper makes it possible to begin to elaborate the links between lower-level and higher-level cognitive processes. We have argued that under certain circumstances it is difficult for the cognitive system to engage in analogical reasoning, yet knowledge of a source can have an effect on reasoning about a target. This research raises a number of issues concerning when subjects will be aware of using knowledge from one domain to solve a problem in a related domain. In traditional studies of transfer, subjects acquire a new concept during the course of the experiment and then must transfer this new concept to another problem. In the current experiment, subjects already had the concept before entering the experiment. It may be the case that subjects are least likely to be aware of a source when the source primes a pre-existing general concept such as inhibition. General concepts such as inhibition, in addition to being associated with many specific instantiations, are likely to exist independently of these specific instantiations. Thus, when a primed general concept memory is accessed it will not result in the retrieval of specific instantiations. However, as in the traditional transfer studies, when the primed concept is a novel concept uniquely tied to a particular instantiation then it is likely that the details of the instantiation regarding when and how the concept was acquired will be retrieved and the subject will make an explicit analogical mapping from the source to the target.

The results and theoretical approach taken in this paper make it possible to understand a number of puzzling findings recently obtained in the literature. For example, Mandler and Orlich (1993), and Lovett and Anderson (1994) have found that performance on a target problem was affected by a source problem, but subjects were not aware of using the source problem while working on the target. The approach that we have taken in this paper suggests that the reason they obtained little awareness of the relationship between a source and a target was that the source problem primed a concept and that this primed concept was then available for solving another problem. Thus, it may be the case that priming, rather than analogical mapping, occurred in their experiments. What we take both their results and our results to mean is that when subjects do not have enough information available to make an explicit analogical mapping, source knowledge can still have an effect through lower-level cognitive processes such as priming.¹³

The findings of this experiment unmask some of the problems associated with using retrospective reports to discover the origins of hypotheses. Many psychologists have used retrospective accounts of a scientific discovery to propose theories of how scientific discoveries were made (e.g., Boden, 1993). However, the results of this experiment demonstrate that even with a relatively straightforward scientific reasoning task, subjects have great difficulty in remembering the sets of factors that contributed to their generation of a new hypothesis. Results like the ones obtained here and that obtained from studies of scientists' on-line problem solving (cf. Dunbar, 1994; Dunbar and Baker, 1994) indicate that scientists are often unaware of the origins of their hypotheses. Given the constructive nature of human memory (e.g., Bartlett, 1932; Owens, Bower, & Black, 1979) it is highly likely that many of the reports of the origins of specific hypotheses in science are incorrect: the scientists may not have a memory of

where their hypotheses came from because their hypotheses were the result of priming, and they must construct a non-veridical story that gives the origin of their hypotheses.

The goals of the research reported in this paper were two-fold. First, we wanted to investigate the roles of lower-level cognitive processes involved in higher-level cognition. We investigated one potential mechanism – priming – and found evidence that priming is at work in complex problem solving. Second, we wished to investigate the paradox regarding the importance of analogy and the empirical findings that subjects have great difficulty in using analogy. We argued that one lower-level process that is important in analogy – priming – could occur, and have an effect on performance, even when subjects are unable to make an explicit analogy. Thus even, when analogy fails, old knowledge can have an effect on the acquisition of new knowledge. Overall, our results demonstrate that it is possible to explore the role of lower-level cognitive processes in higher-level cognition and suggest new avenues for uncovering these mechanisms.

REFERENCES

- Anderson, J. R. (1982). The Architecture of Cognition. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1993). Rules of the mind. Hillsdale: NJ, LEA.
- Bartlett, F. C. (1932). Remembering: A study in experimental and social psychology. London: Cambridge University Press.
- Becker, S., Moscovitch, M., & Behrmann, B. (1994). Long-term semantic priming: A computational account and empirical evidence. Manuscript submitted for publication.
- Blaxton, T. A. (1989). Investigating dissociations among memory measures: Support for a transfer-appropriate processing framework. Journal of Experimental Psychology: Learning, Memory, and Cognition, *15*, 657-668.
- Boden, M. (1993). The creative mind: Myths & mechanisms. New York, NY: BasicBooks, Inc.
- Bruner, J. S., Goodnow, J. J. & Austin, G. A. (1956). A study of thinking. New York: Wiley.
- Catrambone, R., Holyoak, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. Journal of Experimental Psychology: Learning, Memory, and Cognition, *15*, 1147-1156.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. Psychological Review, *97*(3), 332-361.
- Dunbar, K. (1994). How scientists really reason. Scientific discovery in real-world laboratories. In R.J. Sternberg, & J. Davidson (Eds.) Mechanisms of Insight. Cambridge, MA: MIT Press.
- Dunbar, K. (1993). Concept discovery in a scientific domain. Cognitive Science, *17*, 397-434.
- Dunbar, K. & Baker, L. (in press). Goals, analogy and the social constraints of scientific discovery. Behavioral and Brain Sciences.
- Duncker, K. (1945). On problem solving. Psychological Monographs, *58*.
- Ericsson, K. A., & Simon H. A. (1980). Verbal reports as data. Psychological Review, *87*, 215-251.
- Ericsson, K. A., & Simon H. A. (1993). Protocol analysis: Verbal reports as data (Revised Edition). Cambridge, MA: MIT Press.
- Flores d'Arcais, G. B., Shreuder, R., & Glazenborg, G. (1985). Semantic activation during recognition of referential words. Psychological Research, *47*, 39-49.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou and A. Ortony, (Eds.), Similarity and analogical reasoning. New York, NY: Cambridge University Press.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. Cognitive Science, *7*, 155-170.
- Gentner, D., & Forbus, K. (1991). MAC/FAC: A model of similarity-based retrieval. Proceedings of the International Conference on Systems, Man, and Cybernetics (pp. 504-509). Hillsdale, NJ: Erlbaum.
- Gentner, D. & Landers, R. (1985, November). Analogical reminding: A good match is hard to find. Paper prepared for the panel on Commonsense Reasoning at the International Conference on Systems, Man & Cybernetics, Tucson.

- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou, & A. Ortony (Eds.), Similarity and analogical reasoning (pp. 199-241). Cambridge University Press.
- Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. Cognitive Psychology, *12*, 306-355.
- Gick, M. L. & Holyoak, K. J. (1983). Schema induction and analogical transfer. Cognitive Psychology, *14*, 1-38.
- Graf, P., Shimamura, A. P., & Squire, L. R. (1985). Priming across modalities and priming across category levels: Extending the domain of preserved function in amnesia. Journal of Experimental Psychology: Learning, Memory, and Cognition, *11*, 386-396
- Halford, G. S. (1992). Analogical reasoning and conceptual complexity in cognitive development. Human Development, *35*, 193-217.
- Hamann, S. B. (1990). Level-of-processing effects in conceptually driven implicit tasks. Journal of Experimental Psychology: Learning, Memory, and Cognition, *16*, 970-977.
- Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), The psychology of learning and motivation (Vol. 19). New York: Academic Press.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. (1986). Induction: Processes of inference, learning, and discovery. Cambridge, MA: MIT Press.
- Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. Memory & Cognition, *15*, 332-340.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. Cognitive Science, *13*(3), 295-355
- 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., Fay, A. L., & Dunbar, K. (1993). Heuristics for scientific experimentation: A developmental study. Cognitive Psychology, *25*(1), 111-146
- Kuhn, D., & Phelps, E. (1982). The development of problem-solving strategies. In H. Reese (Ed.), Advances in child development and behavior (Vol. 17, pp. 1-44). New York: Academic Press.
- Lovett, M. C., & Anderson, J. R. (1994). Effects of solving related proofs on memory and transfer in geometry problem solving. Journal of Experimental Psychology: Learning, Memory, and Cognition, *20*, 366-378.
- MacLeod, C. M., & Dunbar, K. (1988). Training and Stroop-like interference: Evidence for a continuum of automaticity. Journal of Experimental Psychology: Learning, Memory, & Cognition, *14*(1), 126-135.
- Mandler, J. M., & Orlich, F. (1993). Analogical transfer: The roles of schema abstraction and awareness. Bulletin of the Psychonomic Society, *31*(5), 485-487.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory. Roles of inhibitionless spreading activation and limited-capacity attention. Journal of Experimental Psychology: General, *106*, 226-254.
- Newell, A. & Simon, H. A. (1972). Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.
- Okada, T. (1994). Collaborative concept discovery in a scientific domain. Unpublished doctoral dissertation, Carnegie Mellon University.

- Owens, J., Bower, G. H., & Black, J. B. (1979). The “soap opera” effect in story recall. Memory & Cognition, *7*, 185-191.
- Rappold, V. A., & Hashtroudi, S. (1991). Does organization improve priming? Journal of Experimental Psychology: Learning, Memory, and Cognition, *17*, 103-114.
- Roediger, H. L. III & Challis, B. H. (1992). Effects of exact repetition and conceptual repetition on free recall and primed word-fragment completion. Journal of Experimental Psychology: Learning, Memory, and Cognition, *18*, 3-14.
- Roediger, H.L. III & McDermott (1993). Implicit memory in normal human subjects. In F. Boller & J. Grafman (Eds.), Handbook of Neuropsychology. (Vol. 8, pp. 63-131). Amsterdam: Elsevier.
- Ross, B. H. (1984). Reminders and their effects in learning a cognitive skill. Cognitive Psychology, *16*, 371-416.
- Ross, B. H. (1989a). Distinguishing types of superficial similarities: different effects on the access and use of earlier problems. Journal of Experimental Psychology: Learning, Memory, and Cognition, *15*, 456-468.
- Ross, B. H. (1989b). Reminders in learning and instruction. In S. Vosniadou, & A. Ortony (Eds.), Similarity and analogical reasoning (pp. 438-469). Cambridge University Press.
- Seifert, C. M., McKoon, G., Abelson, R. P., & Ratcliff, R. (1986). Memory connections between thematically similar episodes. Journal of Experimental Psychology: Learning, Memory, and Cognition, *12*, 220-231.
- Sherwood, B. A., & Sherwood, J. N. (1988). The cT language. Champaign, IL: Stipes Publishing.
- Tukey, D. D. (1986). A philosophical and empirical analysis of subjects' modes of inquiry in Wason's 2-4-6 task. Quarterly Journal of Experimental Psychology, *32*, 5-33.
- Tulving, E. & Schacter, D. L. (1990). Priming and human memory systems. Science, *24*, 301-306.
- Smith, E. R. & Branscombe, N. R. (1988). Category accessibility as implicit memory. Journal of Experimental Social Psychology, *24*, 490-504.
- Strull, T. K. & Wyer, R. S. Jr. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. Journal of Personality and Social Psychology, *37*, 1660-1672.
- Wharton, C. M., Holyoak, K. J., Downing, P. E., Lange, T. E. (1994). Below the surface: Analogical similarity and retrieval competition in reminding. Cognitive Psychology, *26*(1), 64-101.
- Wharton, C. & Lange, T.E. (1994). Analogical transfer through comprehension and Priming. In Proceedings of the sixteenth annual conference of the Cognitive Science Society. Atlanta, GA. Hillsdale, NJ: Erlbaum.
- White, H. (1988). Semantic priming of anagram solutions. American Journal of psychology, *101*, 383-399.
- Young, J., Thomsen, C J., Borgida, E., Sullivan, J. L., & Aldrich, J. H. (1991). When self-interest makes a difference: The role of construct accessibility in political reasoning. Journal of Experimental Social Psychology, *27*, 271-296.

Author note

Christian Schunn, Department of Psychology; Kevin Dunbar, Department of Psychology.

This research was supported by a grant to Kevin Dunbar from the National Sciences and Engineering Council of Canada (OGP0037356), by fellowships from FCAR and NSERC to Christian Schunn, and also from NICHD grant R02-HD25211 to David Klahr. We would like to thank Laura Ann Petitto, David Klahr, Marsha Lovett, Brian Ross, Colleen Seifert and two anonymous reviewers for comments on earlier drafts of this paper. Parts of this research were first presented at the Annual meeting of the Cognitive Science Society, Cambridge, MA, July 26, 1990.

Correspondence regarding this manuscript should be addressed to Kevin Dunbar, Department of Psychology, McGill University, Montreal, Quebec, Canada H3A 1B1, or to Christian Schunn, Department of Psychology, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213. Electronic mail: dunbar@ego.psych.mcgill.ca or schunn+@cmu.edu.

Footnotes

¹ Part of the computational simplicity is caused by the lack of complex mental structures. That is, since the priming that has been observed occurs for single relations/concepts rather than for complex systems of relations, the computationally expensive variable-binding process is greatly simplified.

² In addition, there were no common systems of relations involved in both problems. A number of authors (e.g., Clement & Gentner, 1991; Wharton et al, 1994; Halford, 1992) have noted that when there are neither surface features, nor systems of relations similar in both the source and target, subjects will not be reminded of the source and analogical mapping will not occur.

³ It should be noted that priming and analogy are not mutually exclusive since the effects of priming can be orthogonal to the occurrence of analogical reasoning. That is, both priming and analogical reasoning can occur simultaneously. However, the effects of priming can be seen most easily in the absence of analogical reasoning, and therefore that case is emphasized here.

⁴ This task is based upon the experiments that Jacob and Monod (1961) used to discover the lac operon. Jacob and Monod won the Nobel prize for this work. The reasoning used in this task is representative of the types of reasoning that scientists engage in while conducting their research. Thus, it is a useful task to investigate complex reasoning.

⁵ This restriction made it fairly difficult to isolate the effects of a particular controller gene. However, the state of technology at the time that Jacob and Monod made their discovery was similarly restricted.

⁶ The lack of priming effects due to day 1 experience on the concept of activation are not surprising—the more recent experience with activation in the introductory phase of the genetics task most likely overwhelmed the older experience.

⁷ As a test of internal validity, the proportion of zero lactose experiments correlated significantly with the presence of negative regulation in subjects' final hypotheses ($r = .21$, $p < .04$).

⁸ As an indicator of internal validity, failing to conduct a crucial experiment led to not solving the task 90% (18/20) of the time.

⁹ This measure was used by Okada (1994), with the genetics task, as a measure of quality of search.

¹⁰ The post genetics task questionnaire was developed during the course of conducting the experiment. Not all subjects in all conditions were asked all three questions. In particular, none of the subjects in the Negative condition were asked, "Do you see any similarities between yesterday's and today's experiments and what are they?", and only thirteen of the subjects in the Negative condition were given the remaining two questions. Subjects in the other conditions were given all the questions.

¹¹ In the Story condition, 1 of 8 of the solvers had "No memory" of the virus task, in contrast to 0 of 12 non-solvers. The corresponding n 's for the Negative condition were 2 of 5 solvers and 1 of 8 non-solvers.

¹² The important feature influencing when subjects should propose inhibition was whether subjects had conducted an experiment that could suggest inhibition. Given that there was considerable variability in when subjects conducted this type of experiment, it was important to develop a measure of the speed with which

subjects proposed inhibition that removed this incidental source of variance.

Therefore, we used the simple categorical measure of whether subjects proposed inhibition immediately following the first inhibition evidence.

¹³ Priming may also be responsible for the performance of subjects in other experiments. For example, in Gick and Holyoak (1980) 20% of subjects proposed the desired solution without being given a hint to use the prior problem. Some of these subjects may have had the solution primed by the first problem. Use of verbal protocols would make it possible to determine whether this was indeed the case. We would like to thank Colleen Siefert for proposing this possibility.

Table 1

Global Measures of Performance on the Genetics Task for Each Condition (with St. Dev.).

	Correlated		Positive		Negative		Story	
	M	SD	M	SD	M	SD	M	SD
N of experiments	15.5	5.2	13.2	4.4	13.4	3.9	13.1	4.3
Time on problem (in minutes)	43.0	13.6	36.6	14.9	36.1	20.4	38.5	13.8
Time/experiment	2.9	0.6	2.9	1.1	2.6	0.8	3.0	1.1
% Normal/X experiments	9.3	9.3	13.5	12.3	9.0	10.3	15.1	10.8
% zero lactose experiments	6.9	12.0	11.7	16.8	16.5	20.9	6.4	12.1
% of crucial experiments	99.0	4.5	94.4	16.6	96.8	7.5	97.0	9.8
Features varied	1.7	0.2	1.8	0.2	1.7	0.2	1.8	0.2
Dimension search index	13.0	2.1	12.6	2.5	13.1	2.0	12.1	1.8
% of gene combinations	54.8	17.0	50.0	19.1	50.3	12.5	49.8	15.0

Table 2

References to Other Sources of Knowledge in the Verbal Protocols of Subjects in the Genetics Task for Each Condition.

Mentioned Object	Correlated	Positive	Negative	Story
ABC	0	2	4	3
Virus	0	2	0	0
Other	2	1	2	1

Table 3

Numbers of Kinds of Responses in Each Condition to the Question, “Do You See Any Similarities Between Yesterday's and Today's Experiments?”.

	Correlated	Positive	Story
Methods	8	8	3
Goal	7	1	1
Mechanism	0	0	2
Elements	4	4	9
None	3	5	8

Table 4

When Subjects First Mention Inhibition in Each Condition.

	Correlated	Positive	Negative	Story
Immediate	6	7	10	12
Non-immediate	14	11	9	8

Figure captions


Figure 1. Graphs that subjects saw while performing the virus task. 1a: Graph resulting from viewing enzyme G in a cell with an active virus. This shows that as the number of viruses increases, so does the amount of enzyme G. 1b: Graph resulting from adding enzyme Q to a cell with an active virus (in the Negative condition). This shows that the number of viruses decreases when enzyme Q is added.

Figure 2. The computer interface used in the genetics task. Figure 2a: E. coli with controller and producer genes. Figure 2b: E. coli breaking down Lactose with beta-gal.


Figure 3. The proportion of solvers in each condition, with the expected baseline solution rate from Dunbar (1993).

Figure 4. Mean number of inhibition and activation hypotheses in each condition.

1a)



LEGEND




EXPERIMENT


active virus
view G

Continue

1b)



LEGEND




EXPERIMENT

active virus
add Q

Continue

2a



2b

