The Impact of Abstract Ideas on Discovery and Comprehension in Scientific Domains

Shamus Regan (sregan@uic.edu)
Stellan Ohlsson (stellan@uic.edu)
Department of Psychology
The University of Illinois at Chicago
1007 West Harrison Street (M/C 285)
Chicago, IL 60607, U.S.A.

Abstract

The domain-specificity principle implies that domain-specific knowledge is the main determinant of scientific discovery. An alternative view is that scientists make discoveries by assembling and articulating abstract schemas. If so, prior activation of the relevant abstractions should facilitate discovery and comprehension. Two *in vitro* studies showed that abstract information can have as much or larger impact on scientific thinking as domain-specific information.

Abstraction and Discovery

The strongest determinant of a person's performance on a cognitive task is his or her task relevant knowledge. Recently, cognitive scientists have emphasized *domain-specific knowledge*, i.e., facts, principles and procedures that apply in one domain but are of limited usefulness in any other. What matters when playing chess is how much, and what, the player knows about the game of chess (Charness, 1989), but chess knowledge is not very useful outside that game; when diagnosing a patient, what matters is what, and how much, the physician knows about the relevant disease (Patel, Evans & Groen, 1989), but knowledge about diseases is not very helpful in non-medical tasks; and so on. The domain-specificity hypothesis implies that the factor that enables a scientist to make a discovery overlooked by others is his or her superior knowledge of the relevant facts.

The domain-specificity view contrasts with the long standing idea that human intelligence is a function of our ability to reason with *abstract knowledge*. Philosophers from Aristotle to Bertrand Russel have assumed that concepts like symmetry, subset, variation, probability, rate of change, complementarity, sequence, etc.,. confer cognitive power precisely because they apply across content domains (Ohlsson & Lehtinen, 1997). Outside cognitive science, it is still widely assumed that abstract thinking plays an essential role in complex cognitive tasks in general and scientific discovery in particular.

Most scholarly discussions of abstraction focus on the question of how abstractions are formed, but the question of interest here is how an abstraction, once formed, functions. One hypothesis is that an abstract concept, idea or schema can serve as a template for a domain-specific one (Ohlsson, 1993; Schank, 1986). By articulating a prior abstraction vis-à-vis a new situation or task, a person can generate a novel domain-specific structure (explanation, mental model, problem solution) for that situation or task (Ohlsson, 1992b).

This hypothesis is intuitively plausible but empirical research tends to support the domain-specificity view. There are numerous studies in which the participants apparently failed to apply a relevant abstraction in what we as observers think is the obvious way (e.g., Chen, Yanowitz & Daehler, 1995). A common weakness of such studies is that they do not compare the effect of abstract concepts, ideas or schemas with the effect of domain-specific ones. A second weakness from the present point of view is that the tasks used in those studies tend to be unrelated to the kinds of tasks involved in scientific reasoning and discovery.

In this paper we report two *in vitro* studies of the impact of abstractions on scientific thinking. First, we compared the impact of domain-specific and abstract information on students' performance on a laboratory version of the discovery of the structure of DNA. Second, we compared the effects of domain-specific and abstract information on students' ability to comprehend the theory of biological evolution.

Study 1: Discovering DNA

James Watson and Francis Crick discovered the structure of DNA in 1953, after several months of problem solving (Olby, 1974). Although prior research had identified the six types of molecules that made up DNA (sugar, phosphate and four nucleotide bases called adenine, thymine, guanine and cytosine), it was not known how those molecular building blocks fit together.

On their path to the correct structure, Watson and Crick identified the following eight properties of DNA: (a) The DNA molecule is an *elongated* structure. (b) The DNA molecule is *symmetric*. (c) The DNA molecule is held together by *two* (rather than one or three) sugar-phosphate strands, so-called backbones. (d) Due to asymmetries in the sugar molecule, each backbone has a direction; the two backbones in DNA are oriented in *inverse* directions. (e) The nucleotide bases are located *inside* the DNA molecule, with the two backbones twirled around them. (f) The backbones are connected via *pairs* of nucleotide bases, each pair analogous to the rung connecting the two sides of a ladder. (g) Every base pair consists of two *different* molecules. (h) More precisely, each base pair is *complementary* in that adenine is always paired with thymine and guanine with cytosine The first two properties were suggested by prior research by others. Properties c through h were *partial insights* (Ohlsson, 1992a) that Watson and Crick had to attain on the path to discovery.

Watson and Crick's main problem solving method was to build a physical model of the DNA molecule. They manufactured metal pieces that represented the various molecules and tried to put these together in a way that was consistent with quantitative measures (primarily X-ray data) of DNA as well as with its function as carrier of genetic information.

In searching the space of possible structures for DNA, Watson and Crick presumably drew upon their prior knowledge. However, Watson and Crick were not, in fact, experts in the relevant domain. Watson had recently been awarded his Ph.D. and Crick had yet to be awarded his; neither had ever resolved the structure of a large biological molecule. Furthermore, several of their insights had no basis in the chemistry of that time. In particular, there were no prior examples of double helixes, nor of molecules held together by complementary base pairs. As one would expect, the solution was unfamiliar; that is why its attainment was a major scientific advance. These observations suggest that the impact of domain-specific knowledge might have been limited.

A possible alternative is that Watson and Crick, while building their model of DNA, drew upon abstract concepts. After all, concepts like chain, symmetry, parallel sides, direction, inverse, pairing, unlike pairs and complementarity must have been in the conceptual repertoire of these two well-educated scientists. This view of the original discovery implies that prior activation of the relevant abstractions would have simplified the discovery. To evaluate this implication, we conducted what Dunbar (1995) calls an *in vitro* study of the discovery of DNA.

Materials In our laboratory version of the DNA problem, the six types of molecules (sugar, phosphate, four types of nucleotides) were modeled by foam board pieces of different shapes and colors. The pieces were shaped in such a way that they fit together in accordance with the possible chemical bonds between the corresponding molecules. There were 32 pieces. The task of putting this jigsaw puzzle together closely resembles the model-building task that occupied Watson and Crick in the summer of '53 (Olby, 1974).

To make the task performance more practical and easier to record, we reduced the target structure from a double helix to a ladder. This also simplified the problem so that the participants could solve it in a single sitting.(In a pilot study, 80% of the 30 participants solved it in 50 minutes and the median time to solution for those who succeeded was 26 mins. 8 secs.) Notice that all eight partial insights described above have to be attained to solve this two-dimensional version of the DNA problem.

Because many participants might have encountered the DNA molecule in chemistry courses, we removed the problem from its chemical context by describing the puzzle pieces as finds from an archaeological excavation of a fictitious pre-Egyptian civilization. The students were told that the pieces had been used to carry messages from the king to the villages throughout his land, and that they had originally been put together in such a way that the messages could be duplicated with high accuracy by blind priests. The students' task was to figure out the original arrangement of the pieces.

The participants were seated at a table. Their performances were video taped with a camera placed across and slightly above the table.

Design and procedure The 120 participants were randomly divided into seven groups. All groups solved the problem, but under different information conditions.

(a) *No information.* There were 30 participants in this group.

(b) *Domain-embedded information.* Three groups of 15 participants each were given hints formulated within the context of the archaeological cover story. The 2-hint group received hints about properties a and b only, the 5-hint group received hints about properties a through e and the 8-hint group about all eight properties of the target structure. A domain-specific hint presented one of the key concepts embedded within the archaeological cover story. For example, the idea of two parallel strands was presented by saying that ancient sources indicated that the structure "resembled the kingdom itself with its two long and straight eastern and western borders."

(c) *Abstract information.* Three groups of 15 participants each were given 2, 5 or 8 hints formulated abstractly. In a training phase, the experimenter gave the participants a statement of the abstraction, a single example that did not belong to either chemistry or to the archaeological cover story and asked them to generate an example of their own. For example, for the idea of two parallel strands the statement was "things with two parallel parts", the example was "a railroad track" and the instruction was to generate another example of something that has two parallel parts.

Results The main measure derived from the videotapes was the overall time to solve the problem. Figure 1 shows the result. Because the distribution of solution times was skewed, we report medians rather than means. The open points show the effect of the domain-specific hints. Participants who received five hints did better than the participants who received two hints, but the 8-hint group did not do better than the 5-hint group. The difference between two and eight domain-specific hints is 1357 - 1179 = 178 seconds, or 2 minutes 58 seconds. The solid points show the effect of the abstract hints. Although the 2-hint abstract group was slower than the corresponding domain-specific group, the decrease in solution time with number of hints is more systematic. The difference between two and eight abstract hints is 1677 - 1050 = 627 seconds, or 10 minutes and 27 seconds, a huge effect in terms of the reduction in the cognitive processing needed to reach the solution.

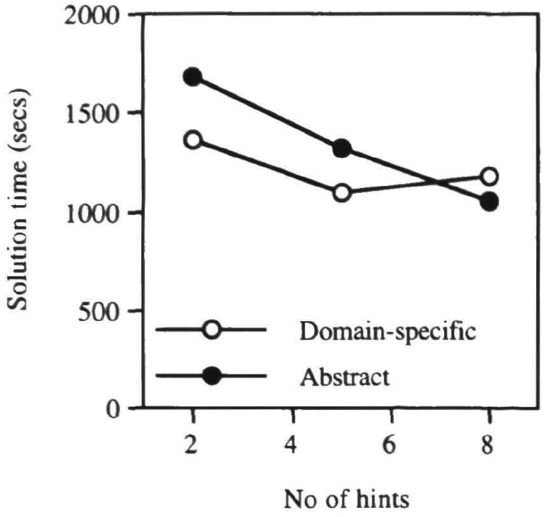

Figure 1: The median time to solution as a function of type of hint and number of hints.

The DNA molecule poses a complex and unfamiliar search problem. Figure 1 suggests that the eight partial insights that we extracted from the historical record functioned as heuristics that guided the search of our participants. The hints did not need to be domain-specific to be helpful. Indeed, the effect of the hints is larger in magnitude and more systematic when the hints were formulated abstractly. This result is consistent with our explanation for how Watson and Crick could solve a major scientific problem without superior domain-specific expertise: The problem is solvable on the basis of abstract concepts that Watson and Crick certainly acquired during their intellectual training.

Study 2: Comprehending Evolution

The elementary version of Darwin's theory of evolution explains particular biological adaptations with five component ideas: variation, inheritance, survival, reproduction rate and the accumulation of small changes over many generations. This biology-specific schema can be articulated vis-à-vis a particular adaptation. ("Polar bears vary in the thickness of their furs; thicker fur enables survival in a cold climate; ... ; over time, they evolved thicker fur.") Evolutionary theory has advanced beyond Darwin, but in this paper we consider his theory in its original version.

Less is known, at the level of day-to-day details, about how Darwin hit upon natural selection than how Watson and Crick discovered the structure of DNA, but Darwin's autobiography and his notebooks have been interpreted to imply that the idea of natural selection was discovered via an analogy with either the Malthussian theory of population growth or with artificial selection in animal breeding (Gruber, 1974; Holyoak & Thagard, 1995).

However, the five key concepts of the original theory were not themselves novel, and they would certainly have been known by Darwin in abstract form. Hence, an alternative hypothesis is that Darwin assembled these abstractions into a new abstract schema and constructed the biological theory by articulating that schema. This hypothesis implies that prior activation of those abstractions should have made that schema easier to construct.

To evaluate this implication, we conducted a second *in vitro* study in which students encountered biological evolution, not as a discovery problem but in the form of a textbook exposition that was 1-2 pages long and contained one illustrative example (the neck of the giraffe). Pilot studies (Ohlsson & Bee, 1991) had shown that constructing Darwinian explanations on the basis of such an exposition is a challenging intellectual task for young adults. We compared the effects of domain-specific information in the form of expert feedback on the participants' own explanations with the effects of abstract information in the form of training intended to establish an abstract variation-and-selection schema prior to reading.

Design and procedure The participants answered evolutionary explanation questions under three different information conditions.

(a) *No information.* Fifty undergraduate psychology students from the University of Pittsburgh and 95 from the University of Illinois at Chicago were asked questions of the form, "How did X evolve Y?" but they were not given any form of preparation, instruction or help. These two base line groups will be reported separately.

(b) *Domain-specific information.* The 20 participants in this group read an expository text that stated Darwin's theory and demonstrated how it can be applied to explain the long neck of the giraffe. They then answered a series of questions of the form, "How did X evolve Y?" (e.g., How did polar bears evolve thick fur?). They were given feedback in the form of expert answers. For each question, the participants first wrote down their own answer and then turned the page and read an expert Darwinian answer to that question. The data from the groups a and b are analyzed in more detail in Ohlsson and Hemmerich (1999).

(c) *Abstract information.* The 38 participants in this group were given pairs of descriptive texts from other domains than biology and were asked to state what the two texts in each pair had in common. Each text pair was designed to instantiate one of the five Darwinian ideas mentioned above or some combination of them. The last pair in the training sequence instantiated the complete Darwinian schema.

Results The participants' written explanations were coded in two ways: First, all explanations were coded with respect to which of the five Darwinian ideas that were expressed in the explanation. This *Darwin score* varied between 0 and 5. Second, all explanations were coded with respect to eight types of non-Darwinian explanation types that we had identified in a pilot study (Ohlsson & Bee, 1991); these explanation types are described elsewhere in this volume (Ohlsson & Hemmerich, 1999). This *misconception score* varied between 0 and 8. Both dimensions were scored by two independent coders.

Figure 2 shows the result in terms of a two-dimensional outcome space. The mean Darwin score is plotted on the vertical axis, the mean misconception score on the horizontal axis. The four data points represent the four groups of subjects. The two base line studies are shown separately, within the ring.

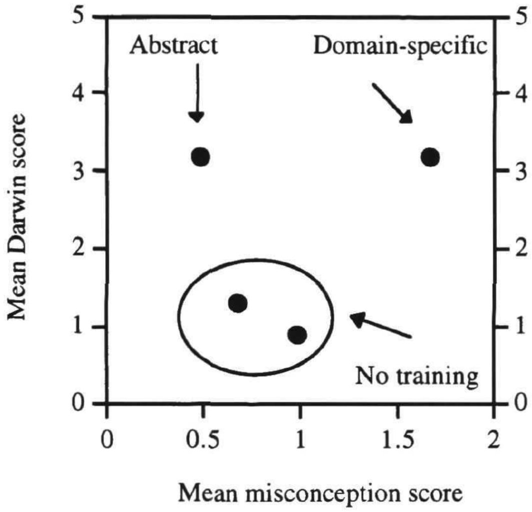

Figure 2: The results of four groups, displayed in a two-dimensional outcome space.

There are two main observations. First, both abstract and domain-specific information increased the use of the five Darwinian ideas in the participants' explanations, compared to the two no-information groups. The abstract information (which did not provide any knowledge about biology) had as strong an effect on the participants' tendency to use Darwinian ideas in their explanations as did the domain-specific feedback. Second, the two types of information had qualitatively different effects when viewed through the lens of the misconception scale. The domain-specific information *increased* the participants tendency to use the eight non-Darwinian ideas (e.g., crossbreeding) that we had identified in pilot work. The abstract information did not have this effect.

These results are consistent with the hypothesis that human beings, Mr. Darwin included, can construct the Darwinian theory by assembling a small group of abstract concepts into a new configuration and then articulating that configuration vis-à-vis biological evolution.

Discussion

The cognitive processes involved in discovery, problem solving, text comprehension, explanation and other cognitive performances must draw upon prior knowledge; there is no other resource for dealing with a situation or a task. The issue is how prior knowledge should be conceived.

There is no doubt that task relevant knowledge impacts on the performance of a task; this is mere common sense. However, the domain-specific knowledge principle might not be the whole story about how people accomplish unfamiliar tasks. In two studies, we showed that domain-specific information had a measurable impact on subsequent task performance, *but so did abstract information that had no relation to the target domain*. Information that aimed to activate certain abstract ideas in the participants' heads produced as large, or larger, effects on subsequent problem solving and comprehension as did domain-specific or domain-embedded information.

These findings are consistent with the idea that prior knowledge should be conceptualized as a repertoire of previously acquired structures (categories, schemas, strategies, rules, etc.) at different levels of abstraction. When a person encounters an unfamiliar task, domain-specific knowledge structures are unlikely to have a high level of fit (or else the task is not unfamiliar after all). In such cases, people respond to the task by articulating the best-fitting abstraction, or assembling a new abstraction which, in turn, is articulated (Ohlsson, 1993; Ohlsson & Lehtinen, 1997).

In both studies, the effects of both domain-specific and abstract information were small in magnitude. This is to be expected with short-term interventions. Future work should compare the effects of training that extends over longer periods of time. Another limitation is that although the quantitative measures verify that the information had an effect, they tell us little about how it worked. Future work will use more detailed data such as think-aloud protocols to elucidate the assembly and articulation of abstractions.

In summary, people in general and scientists in particular cannot operate on the basis of domain-specific knowledge vis-à-vis a task for which they lack such knowledge. The fact that we can solve an unfamiliar problem, understand an unfamiliar text, explain an unfamiliar phenomenon and make unexpected discoveries is due to the fact that we can operate with abstractions. Theories of human cognition that ignore this fact are incomplete.

References

Charness, N. (1989). Expertise in chess and bridge. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing* (pp. 183-208). Hillsdale, NJ: Erlbaum.

Chen, Z., Yanowitz, K.L., & Daehler, M.W. (1995). Constraints on accessing abstract source information: Instantiation of principles facilitates children's analogical transfer. *Journal of Educational Psychology*, 87(3), 445-454.

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

Gruber, H. E. (1974). *Darwin on man*. Chicago, IL: University of Chicago Press.

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

Ohlsson, S. (1992a) Information processing explanations of insight and related phenomena. In M. Keane and K. Gilhooly, (Eds.), *Advances in the Psychology of Thinking* (Vol. 1, pp. 1-44) London, UK: Harvester-Wheatsheaf.

Ohlsson, S. (1992b). The cognitive skill of theory articulation: A neglected aspect of science education? *Science & Education, 1*, 181-192.

Ohlsson, S. (1993). Abstract schemas. *Educational Psychologist, 28*, 51-66.

Ohlsson, S. & Bee, N. (1991). *Young adults' understanding of evolutionary explanations: Preliminary observations* (Technical report, November). Pittsburgh, PA: Learning Research and Development Center

Ohlsson, S., & Hemmerich, J. (1999). Articulating an explanation schema: A preliminary model and supporting data. *Proceedings of the 21st Annual Meeting of the Cognitive Science Society*, Vancouver, Canada, August 19-21, 1999.

Ohlsson, S. & Lehtinen, E. (1997). Abstraction and the acquisition of complex ideas. *International Journal of Educational Research, 27*, 37-48.

Olby, R. (1974). *The path to the double helix: The discovery of DNA*. Seattle, Washington: University of Washington Press.

Patel, V., Evans, D., & Groen, G. (1989). Biomedical knowledge and clinical reasoning. In D. Evans & V. Patel (Eds.), *Cognitive science in medicine* (pp. 57-112). Cambridge, MA: MIT Press.

Schank, R. C. (1986). *Explanation patterns*. Hillsdale, NJ: Erlbaum.