Title
The Effects of Similarity and Individual Differences on Comparison and Transfer

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 32(32)

ISSN
1069-7977

Authors
Day, Samuel
Goldstone, Robert L.
Hills, Thomas

Publication Date
2010

Peer reviewed

The Effects of Similarity and Individual Differences on Comparison and Transfer

Samuel Day (day9@indiana.edu)
Dept. of Brain & Behavioral Sciences, 1001 E. 10th St.
Bloomington, IN 47405 USA

Robert L. Goldstone (rgoldsto@indiana.edu)
Dept. of Brain & Behavioral Sciences, 1001 E. 10th St.
Bloomington, IN 47405 USA

Thomas Hills (thomas.hills@unibas.ch)
Institut für Psychologie, Missionsstrasse 64A
4055 Basel, Switzerland

Abstract

Prior research has found that while people are generally quite poor at recognizing when a new situation is structurally similar to a known case, comparison of two analogous cases greatly improves the likelihood of achieving such recognition. Our study examines the effects of varying the similarity between these compared cases, both featurally and structurally. We find that between-case similarity has a significant impact on transfer, and that these effects interact with characteristics of the learner.

Introduction

Our minds are filled with valuable knowledge that we are unable to use. This is particularly true of what might be the most valuable knowledge of all: general principles that can be applied across a wide range of situations. Research in analogy has repeatedly found that principles that are learned in one context often fail to be retrieved when an individual is confronted with a deeply related situation that differs in concrete or contextual ways (e.g., Gentner, Rattermann & Forbus, 1983; Gick & Holyoak, 1983; Ross, 1984). For example, in Gick and Holyoak's classic (1980; 1983) analogy studies, individuals attempting to solve an insight problem routinely failed to recognize that the problem was analogous to one they had been taught earlier (unless given an explicit hint), and therefore failed to make use of their relevant knowledge. For both theoretical and practical reasons, researchers are keenly interested in finding ways to overcome this kind of impediment.

One approach that has shown great promise is simply asking learners to compare two different examples of a principle (e.g., Gick & Holyoak, 1983; Loewenstein et al, 2003; Gentner et al, 2003; Rittle-Johnson & Star, 2007). For example, Loewenstien and colleagues (2003) conducted research with MBA students enrolled in a course on negotiation. Some of the students compared two specific cases involving a "contingency contract," a useful but sometimes counterintuitive negotiation technique. Other students received the same two cases, but read and analyzed them separately, without any explicit comparison. The researchers found that students who had compared cases were nearly three times more likely to apply the relevant principle to a new case than those students who had analyzed the cases separately. Consistent with prior findings of poor analogical transfer in general, the students who had read but not compared cases performed no better on the transfer task than those who had received no training.

Results such as these point to the potential power of comparison. The most common explanation for these effects is that structural alignments generated when comparing two concrete examples serve to highlight meaningful structural commonalities between them, while simultaneously taking the focus away from elements that are extraneous or irrelevant (e.g., Markman & Gentner, 2000). This, in theory, allows a more explicit representation of the structure or principle itself, making it easier to recognize when it arises in new situations.

However, a great deal remains unknown about the factors that make comparison successful in transfer. Particularly, there is a surprising lack of research on how the relationship between the *compared* cases (such as their similarity) may influence the representations that are formed during comparison. Given that the similarities and differences between the cases are the basis for the knowledge that comparison is presumed to generate, this would seem to be a critical area for study.

For instance, will transfer to new situations be best when the features of the compared cases are relatively similar to one another, or when their content is more dissimilar? There are empirical reasons to predict either of these outcomes. Evidence for "conservative generalization" (Medin & Ross, 1989) suggests that the comparison of two examples that share significant surface commonalities may lead to a representation in which many of these irrelevant features are retained. If so, one of the primary assumed benefits of comparison—a more general representation—may be lost. Comparison of dissimilar cases may therefore lead to representations with broader generalizability. On the other hand, comparisons between overtly similar cases are likely to be less cognitively demanding, and may therefore help to "boot-strap" early learning processes. Consistent with this possibility, Kotovsky and Gentner (1996) found that young children were better able to perform matches on the basis of abstract structural commonalities after performing a similar task involving more perceptually similar stimuli.

A related issue is the effect of the similarity of the structures themselves. Most studies focusing on comparison and transfer have made use of cases with essentially identical relational structures. However, there are reasons to suggest that structural variation may be beneficial as well. For instance, some research has shown that comparing two "near-miss" cases (Winston, 1975), which are identical except for a crucial structural difference, may improve transfer (e.g., Gick & Paterson, 1992). This approach may be particularly effective when an individual needs to discriminate examples of a specific structure from other non-matching cases, as is generally the case in the real world.

The current study examines the impact of both featural and structural similarity in compared cases. Additionally, unlike previous studies, our design requires participants to discriminate different kinds of structures, which may be a more ecologically valid way of assessing the benefits of comparison. Finally, in contrast to previous research that has concentrated on analogical transfer in college-age students, our study uses 7th and 8th grade students in a science class. Children may be more prone to concrete interpretations of scenarios, and thereby miss connections between deeply related scenarios. Given the importance of students appreciating deep principles (e.g. diffusion, order from randomness, and our current topic of interest – feedback loops), it is particularly important to know how children's understanding of principles is influenced by superficial and deep similarities between scenarios.

Experiment

Participants

90 students from a public middle school participated in this study, as part of their regular class time in a General Science course. The group included both 7th- and 8th-grade students (n = 49 and 41, respectively) from six class periods. Roughly half of the students (n = 47) were part of the school's Accelerated Learning Program (ALPs), which is composed of students passing a science achievement test.

Materials and Design

We led the students' class sessions for two days. The first day involved general instruction on the concept of complex systems, including several real-world examples of such systems. This instruction did not include any specific discussion of feedback systems, our target principle. The experiment itself was conducted on the second day.

The overall design of the experiment was as follows: Brief instruction on feedback systems was followed by a pre-test, in which students classified specific scenarios as examples of positive or negative feedback, and answered inference questions about those cases. Students then interacted with two computer simulations, each of which could vary in terms of its content domain (biology or economics) and the type of feedback system it represented (positive or negative). These variations represented the experimental manipulation in the study. Afterwards, students explicitly compared and contrasted the simulations they had completed. Finally, the classification and inference task was administered a second time, as a post-test.

The initial instruction included brief descriptions of positive and negative feedback systems, and included an example of each. These definitions and examples were available to students throughout the experiment.

Pre-test and post-test The pre-test and post-test materials were designed to assess students' understanding of feedback systems, particularly the ability to discriminate positive and negative feedback systems. These materials included eight brief scenarios (averaging 50 words apiece), each describing a real-world phenomenon. Four of these scenarios represented positive feedback systems, and four represented negative feedback systems. For example, one scenario was the following:

> *The lynx is a natural predator of the hare. When lynx populations are small, hare populations increase rapidly. This makes the lynx population increase, since food is plentiful. However, a large lynx population reduces the number of hares, which ultimately brings the lynx population back down. This cycle repeats every ten years or so.*

After reading each scenario, participants classified it as an example of either a positive or negative feedback system by selecting a response from a 5-point scale: *Definitely negative, Probably negative, Don't know, Probably positive, Definitely positive.* They also answered one multiple-choice inference question about each scenario. For example:

> *As the lynx population decreases, the population of rabbits should: [Increase / Decrease]*

Identical items were given at pre-test and post-test. However, in order to minimize explicit memorization and reference to previous answers, students were not informed about the post-test until later in the experimental session.

Computer simulations. All students interacted with two computer simulations demonstrating feedback behavior. These were implemented in NetLogo (Wilensky, 1999), a software package for developing agent-based simulations. Each of the simulations depicted either a positive or a negative feedback system, and each instantiated one of two domains: biology (specifically, interacting slime mold cells) or economics (a simple stock market). This resulted in four relevant simulation types: Biology Positive, Biology Negative, Economics Positive, and Economics Negative. Two versions of each type were created, differing in cosmetic ways. This allowed some students to interact with two different versions of the same type without repeating an identical simulation. The main theoretical focus of our study was on the effects of the similarity between simulations; that is, whether the domain and/or feedback type were the same or different for each participant.

Each simulation began with a brief description of its behavior. For instance, the Economics Positive simulation presented the following introductory description:

> *"This simulation involves a small economic system. People in this system buy stocks, and they pay attention to what other people are doing. When they see someone else buying a stock, they are more likely to want to buy it themselves. When they see someone else* ***selling*** *a stock, they are more likely to sell it themselves. This creates a POSITIVE FEEDBACK LOOP. People buying the stock leads to even* ***more*** *people buying it. People selling the stock leads to even more people* ***selling*** *it."*

The presentation of the simulation was strictly guided although interactive, instructing students to perform specific actions and then to observe the resulting effects

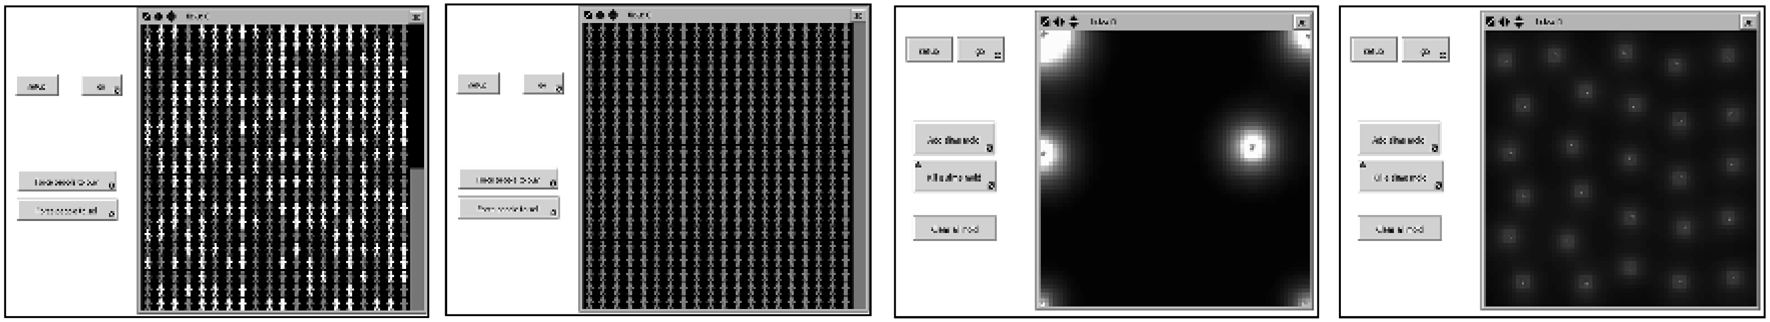

A. B. C. D.

Box 1: Simulations.

The **economic** (stock market) simulations begin (Frame A) with the 420 agents evenly divided between owning the stock (dark; red in the original simulation) and not owning the stock (white). A bar on the right side of the screen indicates the proportion of the agents currently owning the stock. While the simulation is running, each agent will buy or sell the stock with some specified probability. In the Positive Feedback version of the simulation, the probability of buying rather than selling is a positive function of the overall ownership of the stock. As more agents own the stock, the likelihood of new agents purchasing the stock increases. Conversely, as fewer agents own the stock, the likelihood of other agents selling the stock increases. Because of this, random initial fluctuations in stock ownership tend to be amplified over time, and the system quickly moves toward the extremes, resulting in either ownership by all agents or ownership by no agents (Frame B). In the Negative Feedback version, the probability of an agent buying the stock is a ***negative*** function of overall ownership. Therefore, increased overall ownership makes agents more likely to sell the stock, while decreased overall ownership makes agents more likely to buy. This tends to create homeostasis in the system. As the ownership of the stock begins to increase or decrease, the market quickly "corrects" itself and maintains an even proportion of owners and non-owners (as in Frame A).

In the course of the simulation, students are instructed to force a proportion of the agents to buy or sell the stock. This is accomplished by selecting the appropriate button on the left side of the screen, then clicking and dragging across the agents. These interactions serve to highlight the way that the system responds to small imbalances, by either amplifying them (positive feedback) or reducing them (negative feedback). Additionally, students are explicitly reminded at one point during the simulation that it is an example of a positive or negative feedback system. For instance, those in the Negative Feedback version were told: "Observe how this system is a negative feedback loop. People buying the stock leads to other people selling it, and people selling the stock leads to other people buying it. This tends to keep the system in balance, without allowing too many people to own or not own the stock at once."

The **biological** (slime mold) simulations begin with 27 agents (mold cells) randomly distributed on the screen. While the simulation is running, each cell moves about the screen probabilistically, and secretes a chemical that remains for a short period of time in its current location. In the Positive Feedback version, cells are attracted this chemical, and their likelihood of moving toward a location increases with the quantity of the chemical there. Over time, this results in the cells grouping into a small number of clusters (Frame C), since more cells in a given location leads to a greater amount of the chemical there, attracting even more individuals. (Chemical density is reflected by the brightness of a location). In the Negative Feedback version, cells tend to be repelled by the chemical, and are therefore more likely to move to locations where less of the substance is present. This results in the cells attempting to maintain a maximal distance from one another, leading to a relatively homogenous distribution across the field (Frame D).

During the simulation, users are instructed to add additional mold cells to the system, by selecting the "Add slime mold" button and clicking in the desired location on the screen. They are asked at various points to observe the relative effects of clustering these new cells close together versus spreading them out in the space. They are also reminded at one point that the simulation is an example of positive or negative feedback, and why. For example, users in the Positive Feedback version were told: "Observe how this system is a positive feedback loop. Cells produce the chemical in a certain location, which brings other cells to that location, which leads to even ***more*** of the chemical there. This tends to bring the cells together into large clusters.

on the system. For example, students in the Economics simulations were instructed at various times to force a proportion of the agents to buy or sell the stock and observe the results. At one point during each simulation, students were explicitly reminded of which type of feedback system the simulation portrayed (positive or negative), and specifically why this system's behavior reflected that feedback type. After being guided through several relevant actions, students were encouraged to interact freely with the system. Each simulation lasted approximately five minutes. Box 1 provides a detailed description of the simulations.

After completing both simulations, students were instructed: "Now we would like you to compare the two simulations that you just interacted with. Please write about the ways in which the two simulations were similar and different from each other, especially in terms of the way that they behaved." There was no time restriction on the comparison phase. After comparison, all students completed the classification and inference task again.

Predictions. The primary variable of interest is the change in performance between pre-test and post-test. There are several potential predictions about how this variable might be affected by the comparisons that students make. First, prior work on the effects of comparing analogous cases (e.g., Loewenstein et al, 2003) leads us to expect an overall improvement in classification and inference performance, reflecting generally stronger representations of the principles underlying feedback systems. Given that all students are explicitly comparing cases that share a feedback structure, it seems likely that their understanding of such structures should improve on average.

We also predict that the ***kinds*** of comparisons made may affect performance. Comparing two systems involving the same type of feedback (i.e., both positive or both negative) could lead to a bias in the interpretation of new cases. For instance, a student comparing two simulations involving negative feedback may be more likely to classify new cases as examples of negative feedback at post-test.

Another way in which the kind of comparison may matter is in whether it provides an appropriate balance between the *compatibility* (ease of alignment) and the *generalizability* of the two simulations. As discussed, the similarity of the compared cases may have two opposing influences on transfer. Cases that are more similar to one another may be easier to align, and may therefore provide a more straightforward basis for learning about their shared underlying structure. On the other hand, highly similar cases may artificially restrict students' representations of the relevant principles, leading them to only recognize the structure in new situations that are concretely similar to the learned cases. Less similar comparison cases may therefore lead to better generalization of the principles. We predict that learning will be optimal when dissimilarity on one dimension is "scaffolded" by relatively high similarity on another dimension. In the current context, we would predict relatively good performance from those comparing different feedback types in the same domain (e.g., Biology Positive and Biology Negative). In this case, the relevant differences in the positive and negative systems should be particularly highlighted because the concrete features of the simulations are otherwise highly similar. Likewise, strong performance is predicted for individuals comparing the same feedback type across different domains (e.g., Biology Positive and Economics Positive), since the same underlying principles can be observed across more diverse contexts, presumably supporting broader generalization.

We are also interested in potential effects of individual differences between students, and how these may interact with comparison. For instance, it is possible that students in accelerated classes will tend to focus more on the underlying principles of the simulations, and will therefore be less influenced by perceptual variation between them.

Results

Our data yielded several informative findings. Surprisingly, however, most of our initial predictions were not borne out. We first examined the overall improvement of the students between pre-test and post-test. Calculating improvement simply as post-test performance minus pre-test performance, there was no evidence of any improvement on average, either in classification (M = .03, $t(89)$ = 0.52, *n.s.*) or inference (M = .01, $t(89)$ = 0.78, *n.s.*).

Next, we examined possible bias effects in classifications. Specifically, we predicted that individuals who had compared two cases representing the same kind of feedback system (i.e., either two positive cases or two negative cases) would become more disposed to classify new cases as instances of that particular type. For each of these students (n = 43), we calculated bias as the shift toward whichever end of the classification scale matched the type of feedback cases that the student had compared. This measurement did not differ from zero (M = .01, $t(42)$ = 0.23, *n.s.*).

There was also no evidence for the kind of interaction between structural and featural similarity that we had predicted (analysis below). Neither of the conditions that included one similar dimension and one dissimilar dimension showed any improvement (see Figure 1). However, our analysis did reveal several significant results.

We conducted a 2 (Feedback similarity: Same v. Different) × 2 (Domain similarity: Same v. Different) × 2 (ALPs: Accelerated v. Regular classes) ANOVA on the improvement scores. The omnibus test indicated reliable differences between groups for the classification task, $F(7, 82) = 2.27$, $p < .05$. (No effects were found for the inference task on this or any other analysis discussed). Specifically, the test revealed main effects for both Feedback similarity ($F(1, 82) = 4.02$, $p < .05$) and Domain similarity ($F(1, 82) = 6.18$, $p < .05$). In both cases, improvement was greatest when dissimilar cases were compared. Interestingly, for both dimensions of similarity, performance actually decreased numerically at post-test when similar cases were compared (Feedback: similar = -.07, dissimilar = .13; Domain: similar = -.08, dissimilar = .16). This fact explains the absence of the predicted improvement in overall performance: increased scores associated with comparing dissimilar cases were largely offset by ***decreased*** scores resulting from the comparison of similar cases. As seen in Figure 1, the greatest improvement was seen in students who compared cases involving both different feedback

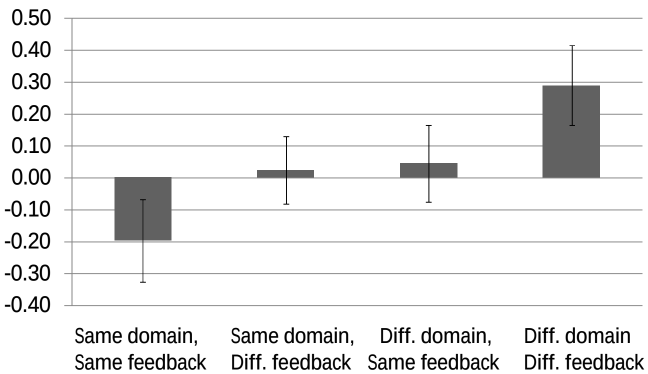

Figure 1: Post-test improvement, by condition

types ***and*** different domains, while the least improvement (actually negative) was seen in those whose comparisons involved the *same* domain and feedback type. Improvement by those in the Different-Different condition was reliably greater than zero (M = .26, $t(19) = 2.22$, $p < .05$). Those in the Same-Same condition were marginally less than zero (M = -.21, $t(21) = 1.79$, $p < .10$). No effect of membership in the accelerated class was observed ($F(1, 89) = 0.33$, *n.s.*).

The influences of structural and featural similarity therefore appear to reflect independent main effects. However, these two effects did not apply equally across all individuals. Interestingly, students in non-accelerated classrooms showed large effects of Domain similarity ($t(42) = 2.83$, $p < .01$), but no evidence of any influence from the similarity of the feedback types that were compared ($t(42) = 0.04$, *n.s.*; see Figure 2). In contrast, the ALPs students were influenced by Feedback similarity ($t(46) = 2.38$, $p < .05$) but not Domain similarity ($t(46) = 0.38$, *n.s*).

Discussion

Several conclusions are suggested by these data. First, the results are consistent with previous characterizations of explicit comparison as a powerful cognitive process that may have an important impact on the acquisition of generalizable principles. Under the right conditions, participants in our study improved reliably in their ability to classify new cases, even in very dissimilar domains. However, our data also suggest that the situation is more complex that is generally proposed, and that comparison is not uniformly beneficial. In fact, on average, explicit comparison by the students was not associated with any improvement at all at post-test. Under some circumstances, there were even trends suggesting that students might be negatively impacted by the comparison process (although these effects were not reliable, they were large enough to effectively offset any overall benefits of comparison). These results highlight the importance of exploring the comparison process more deeply, and attempting to establish the factors that influence comparison-based learning. The remainder of our findings begin to address these issues, exploring aspects of both the compared materials and the learners themselves.

Our study varied both the structural similarity (whether the systems involved the same or different feedback types) and the surface similarity (same versus different content domain) of the compared simulations. We predicted that learning would be optimal when dissimilarity along one dimension was “balanced” by higher similarity on another dimension, which we believed would facilitate alignment while still highlighting important structural features. This prediction was based in part on the approach that has generally been taken in the literature: either presenting the same underlying structure in dissimilar contexts (e.g., Loewenstein et al, 2003), or using “near-miss” cases involving the same content but slightly varying the relevant structure (e.g., Gick & Paterson, 1992). In contrast to our expectations, however, we found that post-test improvement was greatest when the cases were less similar to one another on both dimensions of similarity.

Of course, it is important not to over-interpret the results from one task and set of materials. Each dimension was only tested at two levels, one of which was very high similarity. It is possible (even likely) that these effects do not reflect a simple linear relationship between dissimilarity and transfer, but that there is in fact some optimal similarity level beyond which learning and transfer will decline. Regardless, our results do clearly indicate that the similarity of the compared cases—and not simply the similarity between the learning and transfer cases—is a critical factor influencing whether or not relevant knowledge will be successfully learned and applied. Furthermore, our results highlight the importance of using materials that will maximize the generalizability of the learned representations, and suggest that this factor may often be more important than attempting to facilitate alignment through high similarity.

Perhaps the most interesting—and challenging—finding from our study is the way in which properties of the comparison cases appear to interact with individual differences between learners. Transfer by the students in accelerated classes was influenced by the structural similarity between the cases, but not at all by the similarity of the domains involved. In contrast, structural similarity had no impact on students in regular classes, but learning in these individuals was significantly affected by domain similarity. This finding raises important issues about the effects of comparing cases.

The benefits of comparison are generally attributed to its ability to focus attention on relevant aspects of cases while

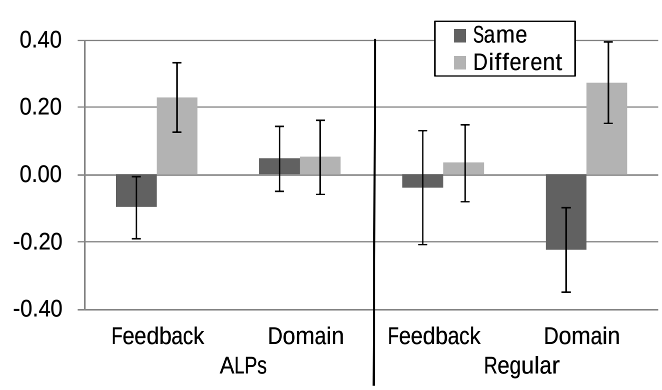

Figure 2: Post-test improvement for accelerated and regular classes.

backgrounding less relevant features. This is, in fact, a mechanism that likely frequently occurs. However, it is important to be mindful of the ways in which differences in individuals' representations of the cases will influence which aspects of the situations are highlighted, and to recognize that these do not always correspond with those that the experimenter may consider "relevant." While membership in the accelerated classes is certainly based on a number of interrelated factors—motivation, achievement, intelligence, ability to focus—it is clear that some difference between the groups is causing them to attend to different aspects of the simulations. These differences appear to have a stark impact on the effects of comparison.

Although more work will be necessary to establish the exact basis of these differences, it seems likely that the ALPs students are better able to look past the immediate surface features of a simulation, and to focus instead on its underlying structural relationships. There are many reasons that this might be the case. For instance, these individuals might be coming to the task with richer background knowledge about the systems that are being presented, and therefore have more cognitive resources available for learning. Consistent with this explanation, students in the accelerated classes had reliably greater performance at pre-test, prior to the primary instruction ($t(89) = 4.60$, $p < .001$). It is also possible that these students have adopted different learning strategies, and are more likely to view all instructional cases as examples of some relevant principle rather than simple facts to be learned independently. Bassok and Holyoak (1989, Experiment 3) found that individuals appeared to acquire the exact same material more concretely or more abstractly based on the specificity of the context in which it was presented. It is possible that successful students have learned to take advantage of this cognitive flexibility by deliberately treating new materials as instantiations of deeper principles, rather than ends in themselves. Previous research has found that experts tend to weigh structural similarities more than superficial similarities (Novick, 1988). The current results extend this finding; even non-experts that are generally high achieving in science show similar tendencies. As such, there appear to be domain-general individual differences in sensitivity to structure that go beyond expertise in a particular domain.

Future research will provide more insight into the exact processes underlying these differences, but our results make clear that characteristics of the learner must be considered when using comparison as an instructional tool. As our data show, cases that lead to reliable gains in one population may foster no improvement at all in another (even very similar) group.

Conclusions

Our knowledge is only valuable to the extent that we are able to make use of it. In previous research, the simple act of comparing two analogous situations has been shown to be extremely valuable in this regard, freeing up concepts that were otherwise bound to a specific context and allowing them to be employed in a much wider range of situations.

The current research shows, however, that these processes may interact in complex and unexpected ways with the features of the cases that are compared and with individual differences in the learner. Our results begin to establish some of the factors that influence the efficacy of comparison, and point the way to future research that may further help us take advantage of this powerful cognitive tool.

Acknowledgments

This work was supported by National Science Foundation REESE grant 0910218. We would like to thank Nancy Martin of Jackson Creek Middle School and Lisa Byrge for their help with our research.

References

Bassok, M., & Holyoak, K. J. (1989). Interdomain transfer between isomorphic topics in algebra and physics. *Journal of Experimental Psychology; Learning, Memory, and Cognition, 15,* 153-16.

Gentner, D., Loewenstein, J., Thompson, L. (2003). Learning and transfer: A general role for analogical encoding. *Journal of Educational Psychology, 95,* 393–408.

Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology, 25,* 524–575.

Gick, M. L., & Paterson, K. J. (1992). Do contrasting examples facilitate schema acquisition and analogical transfer? *Canadian Journal of Psychology, 46,* 539–550.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development, 67,* 2797–2822.

Loewenstein, J., Thompson, L. & and Gentner, D. (2003). Analogical learning in negotiation teams: Comparing cases promotes learning and transfer. *Academy of Management Learning and Education, 2*, 119–127.

Markman, A. B., & Gentner, D. (2000). Structure-mapping in the comparison process. *American Journal of Psychology, 113,* 501–538.

Medin, D. L., & Ross, B. H. (1989). The specific character of abstract thought: Categorization, problem-solving, and induction. In R. J.Sternberg (Ed.), *Advances in the psychology of human intelligence* (*Vol. 5*, pp. 189–223). Hillsdale, NJ: Erlbaum.

Novick, L. R. (1988a). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 510–520.

Rittle-Johnson, B., & Star, J. R. (2007). Does comparing solution methods facilitate conceptual and procedural knowledge? A study on learning to solve equations. *Journal of Educational Psychology, 99,* 561–574.

Ross, B.H. (1984). Remindings and their effects in learning a cognitive skill. *Cognitive Psychology, 16*, 371–416.

Winston, P.H. (1975). Learning structural descriptions from examples. In P.H. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill, pp. 157–210.

Wilensky, U. (1999). *NetLogo (and NetLogo User Manual).* http://ccl.northwestern.edu/netlogo/