Topics in Cognitive Science 2 (2010) 265–284

ISSN: 1756-8757 print / 1756-8765 online
DOI: 10.1111/j.1756-8765.2009.01055.x

The Education of Perception

Robert L. Goldstone,[a] David H. Landy,[b] Ji Y. Son[c]

[a]*Psychological and Brain Sciences, Indiana University*
[b]*Department of Psychology, University of Richmond*
[c]*Department of Psychology, California State University, Los Angeles*

Received 6 April 2009; received in revised form 20 August 2009; accepted 24 August 2009

Abstract

Although the field of perceptual learning has mostly been concerned with low- to middle-level changes to perceptual systems due to experience, we consider high-level perceptual changes that accompany learning in science and mathematics. In science, we explore the transfer of a scientific principle (competitive specialization) across superficially dissimilar pedagogical simulations. We argue that transfer occurs when students develop perceptual interpretations of an initial simulation and simply continue to use the same interpretational bias when interacting with a second simulation. In arithmetic and algebraic reasoning, we find that proficiency in mathematics involves executing spatially explicit transformations to notational elements. People learn to attend mathematical operations in the order in which they should be executed, and the extent to which students employ their perceptual attention in this manner is positively correlated with their mathematical experience. For both science and mathematics, relatively sophisticated performance is achieved not by ignoring perceptual features in favor of deep conceptual features, but rather by adapting perceptual processing so as to conform with and support formally sanctioned responses. These ''rigged-up perceptual systems'' offer a promising approach to educational reform.

Keywords: Perceptual learning; Education; Mathematical reasoning; Complex systems; Scientific reasoning

1. Introduction

Understanding science and mathematics, on first sight, relies on the highest of high-level cognition. Scientific reasoning depends on analytic thought, making novel and creative associations between dissimilar domains, and developing deep construals of phenomena that run counter to untutored perceptions. In fact, Quine (1977) considered a hallmark of advanced

Correspondence should be sent to Robert L. Goldstone, Psychological and Brain Sciences, Indiana University, Bloomington, IN 47408. E-mail: rgoldsto@indiana.edu

scientific thought to be that it no longer requires notions of overall perceptual similarity as the basis for its categories (see also Goodman, 1972). The rationale for this claim is that unanalyzed perceptual similarities may lead one astray. For example, marsupial wolves may closely resemble the placental wolves customary in the Northern Hemisphere, but they are evolutionarily rather distant cousins. The generalization of this example is that as a scientist develops more complete knowledge about the reasons why an object has a property, then overall perceptual similarity becomes decreasingly relevant to generalizations. The knowledge itself guides whether the generalization is appropriate. Historical, as well as lifelong (Carey, 2009; Chi, Feltovich, & Glaser, 1981), maturation in scientific and formal reasoning can be viewed as increasing reliance on deep principles and decreasing reliance on potentially misleading perceptual resemblances.

Although there is certainly justification for this opposition between superficial perception and principled understanding (Sloman, 1996), we will be advocating the converse strategy of trying to ground scientific and mathematical reasoning in perceptual processing. In particular, we pursue the agenda of co-opting natural perceptual processes for tasks requiring abstract or analytic reasoning (see also Barsalou, 2005; Goldstone & Barsalou, 1998). A first reason for pursuing this agenda is that our visual and auditory perceptual systems are relatively extensive neuroanatomically speaking, and are phylogenetically early. They are, accordingly, good candidates for reuse by later developing, high-level cognitive processes.

Second, because perceptual systems encode aspects of external objects in a relatively direct, ''raw'' fashion, they can implicitly represent certain aspects of those objects without explicit machinery to do so (Palmer, 1978). Representations that intrinsically preserve physical properties are often more efficient than purely symbolic representations because they do not require external constraints to assure proper inferences. Relatively raw representations are particularly useful when one does not know what properties will be needed at a later point, or explicitly how to compute the needed properties. Surprisingly then, perceptual representations are often times most useful for more complex cognitive processes—those without simple definitions or rules.

Third, many properties of abstract cognition, when explored from the perspective of processes that could furnish them, are also found in perceptual systems. Rule use in high-level cognition is paralleled in perception by selective attention. Both involve highly efficient selection of relevant, and inhibition of irrelevant, attributes. Schizophrenics have difficulty inhibiting both inappropriate thoughts and irrelevant perceptions (Beech, Powell, McWilliam, & Claridge, 1989). Conversely, many of the perceptual and cognitive symptoms of childhood autism, including hypersensitivity to sensory stimulation, abnormally narrow generalizations from training, and lack of productive language, may be traced to an overly selective attentional process (Lovaas, Koegel, & Schreibman, 1979). Other examples of parallel processes between perception and high-level cognition include structural binding (of features into objects for perception, or fillers into roles for cognition), simplification (through visual blurring or strategic cognitive filtering), and cross-domain matching (synesthesia in perception, or analogical reasoning in cognition). Considerations of individual differences, task manipulations, and neuropsychological data provide enough evidence for correlations between perceptual and conceptual tasks to encourage pursuit of the possibility that they are linked by overlapping processes rather than by mere analogy (Barsalou, 2008; Goldstone & Barsalou, 1998).

2. Adapting perception to fit cognition

An important reason why perception and high-level cognition are more closely related than might appear at first sight is that perception is not limited to what we see at first sight. We adapt our perceptual systems to fit our higher-level cognitive needs. Experts in many domains, including radiologists, wine tasters, and Olympic judges, develop specialized perceptual tools for analyzing the objects in their domains of expertise (Gauthier, Tarr, & Bubb, 2009). Much of training and expertise involves not only developing a database of cases or explicit strategies for dealing with the world but also tailoring perceptual processes to more efficiently gather information from the world (Gibson, 1991). Even expertise in mathematics and science often involves perceptual learning. Biology students learn to identify cell structures, geology students learn to identify rock samples, chemistry students learn to recognize chemical compounds by their molecular structures, and students of mathematics rely on recognizing regularities in notation (Cajori, 1928).

Research in perceptual learning indicates influences of cognitive tasks on perceptual systems that are surprisingly early in the information processing sequence. For example, practice in discriminating small motions in different directions significantly alters electrical brain potentials that occur within 100 ms of the stimulus onset (Fahle & Morgan, 1996). Prolonged practice with a subtle visual categorization results in much improved discrimination, but the improvements are highly specific to the trained orientation of grating patterns (Notman, Sowden, & Özgen, 2005). This profile of high specificity of training is usually associated with changes to early visual cortex (Fahle & Poggio, 2002).

Perceptual change can also be functional to naturally occurring tasks of an individual. Expertise can lead to improvements in the discrimination of low-level, simple features, as with the documented sensitivity advantage that radiologists have over novices in detecting low-contrast dots in X-rays (Sowden, Davies, & Roling, 2000). Expertise for visual stimuli as eclectic as butterflies, cars, chess positions, dogs, and birds has been associated with an area of the temporal lobe known as the fusiform face area (Bukach, Gauthier, & Tarr, 2006). In general, perception is adapted to promote the categories or responses required for performing a task, and these adaptations often occur at an early stage of processing.

As a field, perceptual learning has been principally concerned with low- to middle-level changes to perceptual processing (Fine & Jacobs, 2002). However, in our two case studies from science and mathematics, we will be considering perceptual learning in fairly late processes that are implicated in object and event interpretation because the payoffs for perceptual flexibility, at all levels, are too enticing to forego. They allow an organism to respond quickly, efficiently, and effectively to stimuli without dedicating online executive control resources. Instead of strategically determining how to use unbiased perception to fit one's needs, it is often easier to rig up a perceptual system to produce task-relevant representations, and then simply leave this rigging in place without strategic control. This ‘‘rigged-up perceptual systems'' (RUPS) hypothesis can be stated explicitly as follows: An important way to efficiently perform sophisticated cognitive tasks is to convert originally demanding, strategically controlled operations into learned, automatically executed perceptual processes. These tasks can be understood as on par with the kinds of ''visual routines'' proposed by Shimon Ullman (1984) to account for how people extract information from a visual scene using processes such as shifting attentional focus, indexing items, tracing boundaries, and spreading activation from a point to the boundary of an area. We will consider this RUPS with respect to two high-level cognitive tasks with educational relevance—one from scientific reasoning and one from mathematics.

3. Tuning perception to scientific principles

In Section 1, we described the marsupial versus placental wolf case as an example of superficially similar objects with deeper, biological differences. An example of the converse comparison, superficially dissimilar objects with deep commonalities would be whales and hippopotamuses, whose current dissimilarities belie their relatively recent shared ancestor. Cases like these are pedagogically important, because it has proven difficult for students (and other people) to transcend superficial appearances and appreciate hidden, deep principles (Carraher & Schliemann, 2002). Learners in many domains do not spontaneously transfer what they have learned to superficially dissimilar domains (Detterman, 1993; Gick & Holyoak, 1980, 1983). Physics professors who have taught students to find the time required for a ball to fall on the ground from a 200-ft tower have been shocked when their students fail to see the applicability of the same equations for finding the time required for a ball to fall to the bottom of a 200-ft well (David Perkins, personal communication, 1998). Although the teacher thinks in terms of general gravitational acceleration, the students fail to make the leap from towers to wells.

To understand transfer, instead of examining these all-too-common failures, we consider from our classroom and laboratory observations (Goldstone & Sakamoto, 2003; Goldstone & Son, 2005; Goldstone & Wilensky, 2008; Son & Goldstone, 2009a) a situation where spontaneous transfer occurs, and analyze its perceptual basis. This situation involves the scientific principle of ''competitive specialization.'' In competitive specialization scenarios (Rumelhart & Zipser, 1985), a good solution is found if every region has an agent reasonably close to it. For example, an oil company may desire to place oil drills such that they are well spaced and cover their territory. If the oil drills are too close, they will redundantly access the same oil deposit. If the oil drills do not efficiently cover the entire territory, then some oil reserves will not be accessed.

3.1. Ants and Food

The first example of competitive specialization involves ants foraging food resources drawn by a user. The ants follow exactly the following three rules of competitive specialization. At each time step, (a) a piece (pixel) of food is randomly selected from all of the food drawn by a student, (b) the ant closest to the piece moves with one rate, and (c) all of the other ants move with another rate. In interacting with the simulation, a learner can reset the ants' positions, clear the screen of food, draw new food, place new ants, move ants, start/stop the ants' movements, and set a number of simulation parameters. The two most critical user-controlled parameters determine the movement speed for the ant that is closest to the selected food (called ''closest rate'' in Fig. 1) and the movement speed for all other ants (''not closest rate''). Starting with the initial configuration of three ants and three food piles shown in Fig. 1, several important types of final configuration are possible and are shown in Fig. 3. If only the closest ant moves toward a selected piece of food, then this ant will be the closest ant to *every* patch of food. This ant will continually move to new locations on every time step as different patches are sampled, but it will tend to hover around the center of mass of the food patches. The other two ants will never move at all because they will never be the closest ant to any food patch. This configuration is suboptimal because the average distance between a food patch and the closest ant (a quantity that is continually graphed) is not as small as it would be if each of the ants specialized for a different food pile. On the other hand, if all of the ants move equally quickly, then they will quickly converge to the

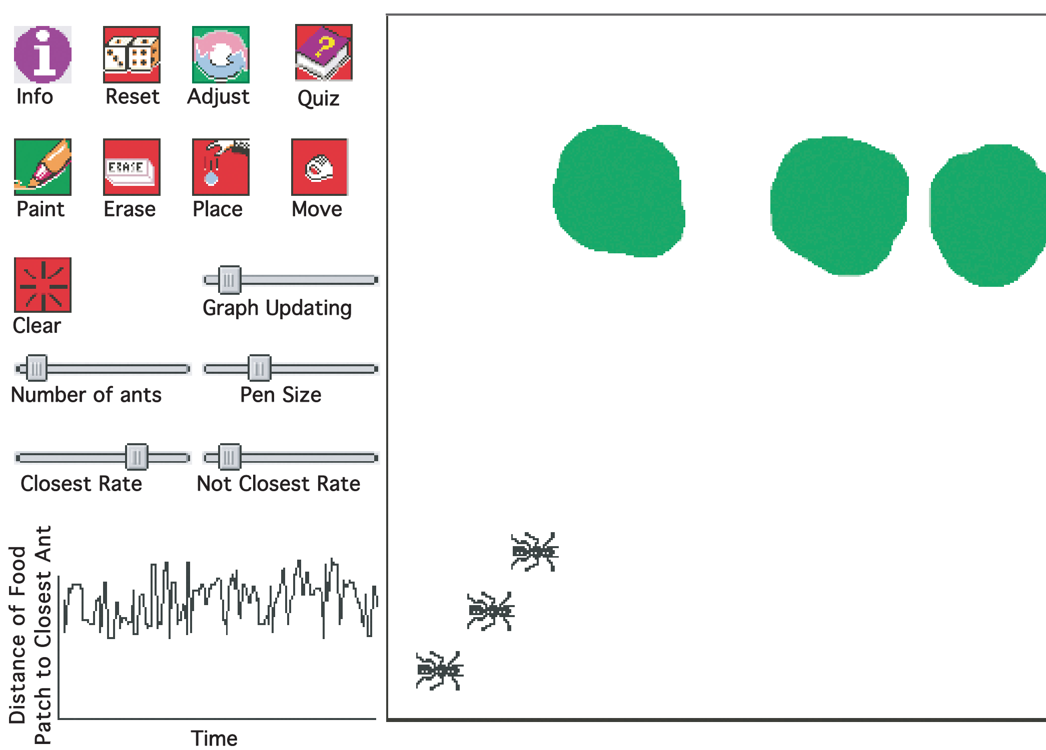

Fig. 1. A screen dump of an initial configuration for the ''Ants and Food'' computer simulation. At each time step, a patch of food is randomly selected, and the ant closest to the patch moves toward the patch with one speed (specified by the slider ''Closest Rate'') and the other ants move toward the patch with another speed (''Not Closest Rate'').

same screen location. This also results in a suboptimal solution because the ants do not cover the entire set of resources well—there will still be patches that do no have an ant nearby. Finally, if the closest ant moves more quickly than the other ants but the other ants move too, then an approximately optimal configuration is achieved. Although one ant will initially move more quickly toward all selected food patches than the other ants, eventually one of the other ants will move so as to be the closest to another patch, thereby allowing for locational specificity for both ants.

An important, subtle aspect of this simulation is that poor patterns of resource covering are self-correcting so the ants will almost always self-organize themselves in a 1-to-1 relationship to the resources regardless of the lopsidedness of their original arrangement if good parameter values are used.

3.2. Pattern Learning

The second example of competitive specialization, shown in Fig. 2, involves sensors responding to patterns drawn by the user. Just like the Ants and Food scenario, the Pattern

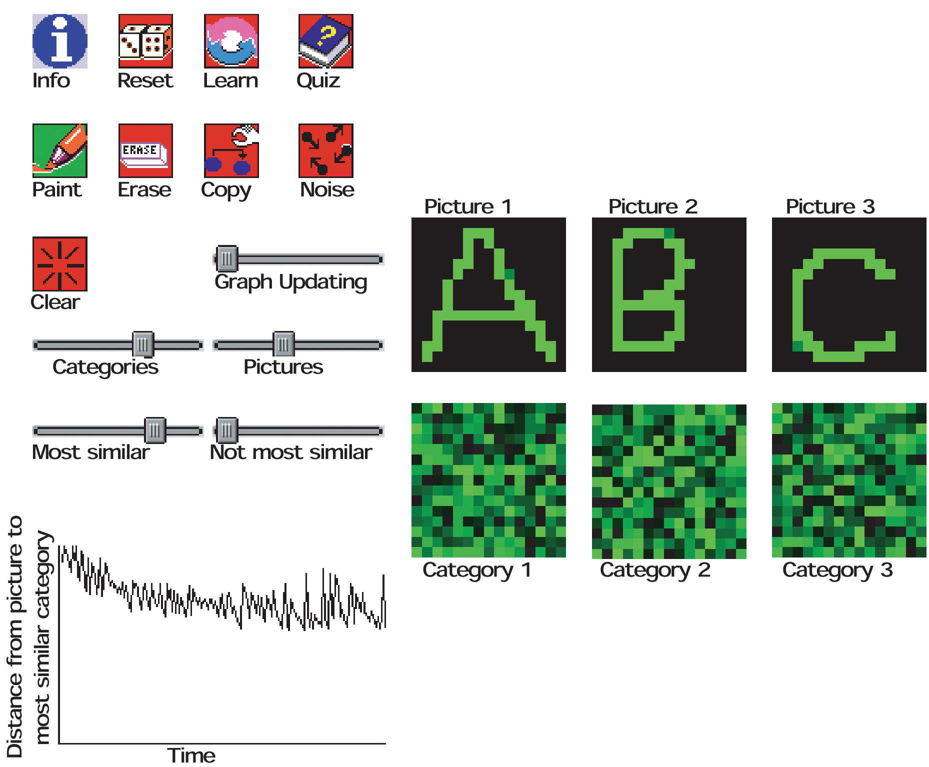

Fig. 2. A screen dump for the simulation ''Pattern Learning.'' Users draw pictures, and prior to learning, a set of categories are given random appearances. During learning, a picture is selected at random, and the most similar category to the picture adapts its appearance toward the picture at one rate (specified by the slider ''most similar'') while the other categories adapt toward the picture at another rate (''not most similar'').

Learning simulation also follows the three rules of competitive specialization. At the beginning of the simulation, the sensors respond to random noise. But at each time step, a pattern is randomly selected from all of the patterns drawn by a user, and the sensor most similar to that pattern adapts to become more similar to that pattern at a particular rate. All of the other sensors adapt toward the selected pattern at another rate. Users can reset sensors to become random again, draw patterns, erase patterns, copy patterns, add noise, start/stop pattern learning, and change a number of parameters. The most important parameters are the rates of adaptation for the most similar and not most similar sensors.

3.3. Transfer by primed perceptual interpretations

As may be clear to the reader, these two situations are, at their heart, the same system. Pattern Learning is a high-dimensional generalization of the two-dimensional Ants and Food scenario. Some of the functional equivalents of the scenarios are shown in Fig. 3.

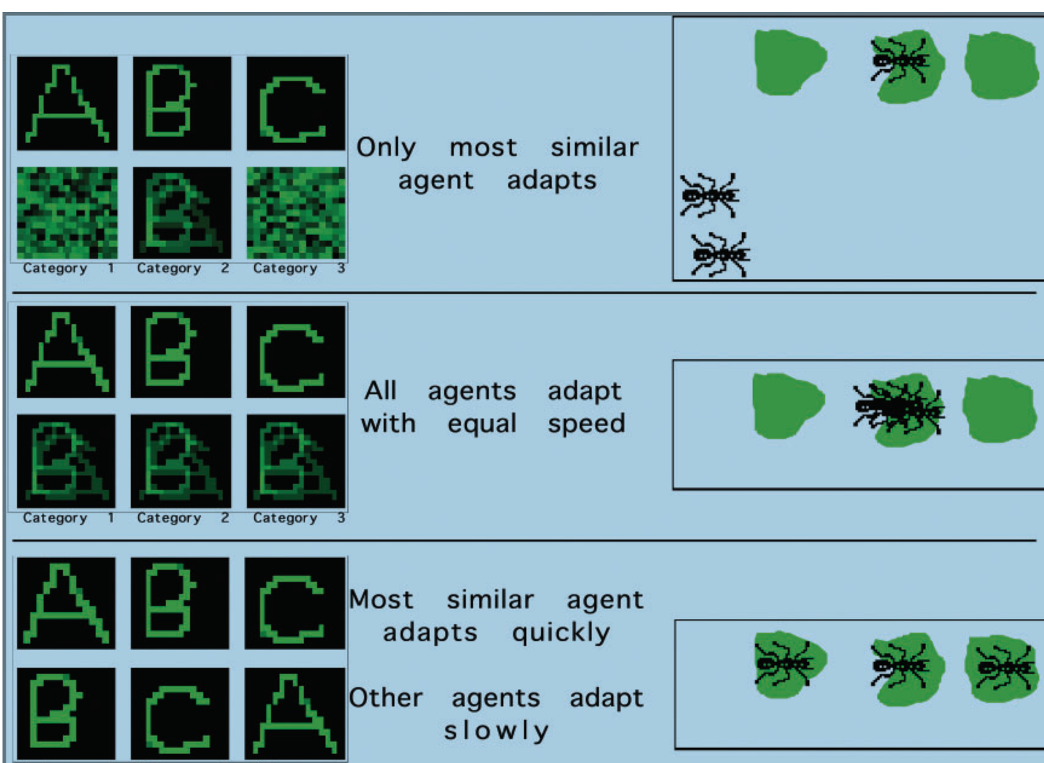

Fig. 3. The basis for the isomorphism between the Ants and Food and Pattern Learning simulations. If only the most similar agent to a resource adapts, then often a single agent will move toward the average of all of the resources. If all agents adapt equally quickly, then they will all move toward the average position. If the agent closest to a resource patch moves much faster than the other agents but all agents move a bit, then each of the agents typically becomes specialized for one resource type.

Surprisingly, our laboratory and classroom investigations with these two demonstrations of competitive specialization have shown that students can, under some circumstances, transfer what they learn from one simulation to another (Goldstone & Sakamoto, 2003; Goldstone & Son, 2005). Even though there is a lack of superficial perceptual features shared by Figs. 1 and 2, we nonetheless believe that the observed transfer is due to perceptual learning. In particular, we believe that students who interact with the Ants and Food simulation actively interpret the presented perceptual patterns, and carry over these interpretations to the more complex Pattern Learning situation. We see this process of carrying over interpretations as akin to Leeper's (1935) classic finding that an ambiguous man/rat drawing is automatically interpreted as a man when preceded by a drawing of a man and as a rat when preceded by a rat. This phenomenon is not ordinarily thought of as transfer, but it is an example of a powerful influence of prior experiences on high-level perception. Originally dissimilar events can come to be seen as similar because the perceptual interpretations may be highly selective, perspective-dependent, and idealized.

Two alternative characterizations of this RUGS effect are usefully delineated because they bear generally on mechanisms of perceptual learning. One classic dispute, as characterized by Gibson and Gibson (1955, p. 34), concerns the question ''Is learning a matter of enriching previously meager sensations or is it a matter of differentiating previously vague impressions?'' According to the enrichment view, perceptions change as sensory information becomes associated with and enriched by accompanying information such as labels, outcomes, or contexts (Postman, 1955). These enrichments can be said to bias one's perception of an event. According to the Gibsons' differentiation view, perceptions change not by becoming connected to learned associations, but by becoming more connected to the external world and its properties. Thus, it is assumed that learning involves responding to previously ignored sensory information. The Gibsons typically interpret this as learned perceptual selection of critical event information. Our account of the psychological change associated with exploring the Ants and Food situation is more consistent with the Gibsons' account in that our students do not seem to be simply biased to interpret all situations via a competitive specialization lens. Only situations that truly instantiate the principle are interpreted according to the principle (Goldstone & Sakamoto, 2003; Son & Goldstone, 2009a). However, rather than characterizing the perceptual change as improved selection of task-relevant *features* as the Gibsons do, we emphasize that the entire process of interpreting a situation is altered.

One source of evidence that students are indeed basing their interpretation of Pattern Learning on their interpretation of the preceding simulation is that they can solve problems in the Pattern Learning scenario more effectively when it has been preceded by Ants and Food than a scenario governed by a different principle (Goldstone & Son, 2005; Son & Goldstone, 2009a). One prime facie reason to believe that the beneficial transfer is perceptual in nature is the students' spontaneous visualizations, examples of which are shown in Fig. 4. A student was asked to visually describe what would happen when there are two categories and four input pictures that fell into two clusters: variants of As and variants of Bs. The student drew the illustration in Fig. 4A. In this illustration, adaptation and similarity are both represented in terms of space. The two categories (''Cat1'' and ''Cat2'') are depicted

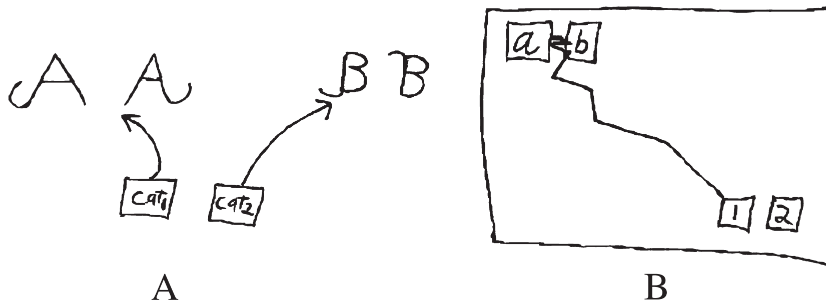

Fig. 4. Typical visualizations of students expressing their knowledge of Pattern Learning. In (A), the student represents the adaptation of categories by moving them through space toward two clusters of stimuli. The similarity of the two variants of ''A'' is represented by their spatial proximity. In (B), a student was asked to illustrate a problem that arises when the closest category to a pattern adapts, but the others do not adapt at all: a single category is shown oscillating between two patterns. Note the similarity to the top panel of Fig. 3 with the Ants and Food simulation.

as moving spatially toward the spatially defined clusters of ''A''s and ''B''s, and the ''A'' and ''B'' pictures are represented as spatially separated. The context for Fig. 4B was a student who was asked what problem might occur if the most similar detector to a selected picture moved quickly to the picture, whereas the other detectors did not move at all. The student showed a single category (shown as the box with a ''1'') moving toward, and eventually oscillating between, the two pictures. Again, similarity is represented by proximity and adaptation by movement.

Students' verbal descriptions also give evidence of the spatial diagrams that they use to explicate Pattern Learning. Students frequently talk about a category ''*moving over* to a clump of similar pictures.'' Another student responded that ''this category is being *pulled in two directions—toward* each of these pictures.'' A third student also uses spatial language when describing, ''These two squares are *close* to each other, so they will tend to *attract* the same category to them.'' All three of these reports show that students are construing visual similarity in terms of spatial proximity. Our students find making the connections between adaptation and motion, and between similarity and proximity to be natural, but only after they have had experience with Ants and Food. Specifically, 30% of participants used spatial terminology when describing Pattern Learning when it was preceded by Ants and Food, whereas only 8% of participants did so when Ants and Food was presented first. Moreover, participants' use of spatial terminology was significantly correlated with their success at solving problems posed in the Pattern Learning scenario, such as developing an automated procedure for developing three categories that each becomes uniquely specialized toward one of three input pictures. Pattern Learning problems were solved an average of 40 s sooner when students used spatial terminology to describe the scenario.

Our proposal is that students are using the literal, spatial models that they learned while exploring Ants and Food to understand and predict behavior in Pattern Learning. If students connect visual similarity to spatial proximity, they conduct the same kinds of mental simulations in Pattern Learning that they perform when predicting what will happen in new Ants and Food situations. So we do not believe that students abstract a formal structural description that unifies the two simulations. Instead, they simply apply to a new domain the same perceptual routines that they have previously acquired. To the experienced eye, identical perceptual configurations are visible across the two simulations. The property of ''one thing trying to cover everything'' is seen in both simulations on the top row of Fig. 3, and ''Everybody is doing the same thing'' is seen in both simulations in the middle row. The critical point is that these paired simulations are obviously not perceptually similar under just any construal. It is only to the student who has understood the principle of competitive specialization as applied to Ants and Food that the two situations appear perceptually similar.

There is additional evidence from students' interpretations that their perception of cross-scenario similarity is driving transfer. Many interpretations of the Ants and Food scenario would not be expected to increase perceived similarity to Pattern Learning because they were domain-specific, ''arthropocentric'' interpretations of the ants' behavior. Participants frequently described ants as scaring each other away, avoiding crowds, being tempted by food, or being tired. A judge tallied the number of domain-specific interpretations over all of the descriptions, both unique and shared. For an interpretation to count as domain-specific, it needed to be (a) applicable to sentient agents like ants but not nonsentient agents like the categories of the second simulation, (b) described in terms of ''ants'' or ''food'' rather than more abstract language such as ''forager'' or ''groups,'' and (c) not simply a specific instantiation using ants and food of the abstract rules that in fact governed the ants' behavior. Participants who formed these kinds of domain-specific interpretations showed no beneficial transfer to Pattern Learning.

This result would also be consistent with students' conceptual understandings driving their perceptual interpretations rather than our thesis—that the perceptual interpretations are driving transfer. However, by manipulating the likelihood of producing a domain-specific perceptual interpretation, we affect transfer. In particular, when the Ants and Food are both given concrete manifestations in which they are clearly identifiable by their graphical simulation elements as ants and food, then participants are much more likely to give domain-specific interpretations such as ''this ant is tired'' or ''this ant is scared of this other ant.'' These domain-specific interpretations give rise to worse transfer to Pattern Learning (see also Goldstone & Sakamoto, 2003; see also Son & Goldstone, 2009b). In contrast, when the ants and food are depicted by simple dots for the ants and blobs for the food, then fewer domain-specific interpretations are given and better transfer is achieved. Given this result, we believe that the perceptual encodings of the simulation elements drive transfer, rather than a priori individual differences in understanding driving perceptual interpretations.

The benefit of idealized over concrete depictions for transfer also undercuts treating ''perceptual'' as equivalent to ''superficial.'' The facilitation in understanding Pattern Learning when it has been preceded by Ants and Food is explainable by the deployment of a dynamic, spatial model of agents moving toward resources with differential rates, and each agent tending to move quickly only toward resources that are close. Perceptual aspects that are not relevant to this model do not promote beneficial transfer. This is consistent with other results. Mathematical systems are more readily transferred when they are conveyed with generic symbolic forms rather than more concrete graphical representations even when the latter have features that intrinsically conform to the underlying formalism (Kaminski, Sloutsky, and Heckler, 2008).

4. Adapting perception and action to algebraic reasoning

Our second case study of the application of RUPS is in the domain of algebraic reasoning. Even more so than scientific reasoning, mathematical reasoning is traditionally assumed to involve formal operations. Algebraic symbolic reasoning is assumed to depend on internal structural rules. Algebra is often considered to be the best example of widespread computational cognition—the application of laws to structured strings, where those laws generalize on the basis of the identity of the symbols, and that identity is arbitrarily related to the symbols' content (Fodor, 1992). In contrast, our prior research (Landy & Goldstone, 2007a,b) has indicated that our participants are heavily influenced by groupings based on perceptual properties when performing both algebra and arithmetic. Despite being reminded of, and verbally subscribing to, standard order of precedence rules according to which multiplications are performed before additions, our college student participants are much more likely to calculate an incorrect solution value of 25 for ''2+3 * 5 = ?'' than ''2 + 3*5 = ?.'' The speed and accuracy of calculation suffers when the physical spacing around an operation is inconsistent with its order of precedence, for example, when there is less space around the ''+'' sign than ''*'' sign. In general, participants seem to create perceptual groups of notational elements, and use these groups, rather than just calculation rules, to perform mathematics.

From these results, it might be concluded that there is an inherent and perpetual conflict between rule-based and perceptual processes. In contrast, the RUPS position is that we adapt our perceptual processing so as to produce results that are consistent with formally sanctioned mathematical rules. We look for ways of easing the load on our executive control system—the system required to prioritize, execute, and monitor rules. One excellent way to ease this load is to have our perceptual systems naturally do the formally right thing. We briefly consider three results from algebraic reasoning that suggest RUPS.

4.1. Eye movements and order of precedence

Rather than memorize the precedence rule ''multiplication and division before addition and subtraction,'' we may learn to move our eyes in a way that has the effect of honoring this explicit rule automatically. In one experiment supporting this conjecture, participants' eye movements were measured when they were asked to solve arithmetic problems like ''2 × 3 + 4'' (called a ''multiplication–addition'' problem because the multiplication occurs first) and ''2 + 3 × 4'' (an ''addition–multiplication'' problem) (Landy, Jones, & Goldstone, 2008). Fig. 5 shows that participants tended to look at the multiplication portion of the expression during the early stages of the trial. In addition, their very first eye movements tended to be toward the multiplication, and gazes to multiplications lasted longer than

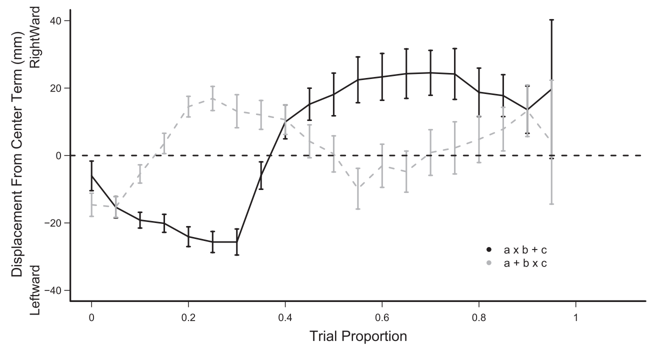

Fig. 5. Mean gaze position (rightward gaze is plotted as positive) for addition–multiplication and multiplication–addition problems across the duration of a trial (expressed as the proportion of the trial). Participants tend to look at the multiplication sign before the addition sign.

gazes to additions. Similar results are found whether multiplication is denoted by an ‘‘×’’ or dot.

4.2. Influence of order of precedence on attention tasks not explicitly involving mathematics

We interpret the previous results as suggesting that people’s perceptual systems become rigged up over practice with mathematics to automatically gravitate toward exactly the region of an equation that they should first process according to the dictates of mathematics. An alternative account is that attention is quickly but strategically being allocated exactly to where it ought to be for the mathematical task at hand. To address this counterproposal, we have also explored the allocation of attention in tasks that use mathematical notation but do not require a mathematical response.

In one such task, we simply asked participants which side of a multiplication–addition or addition–multiplication expression such as ‘‘2 × 3 + 4’’ or ‘‘2 + 3 × 4’’ contained the target operation. The target operation was switched from ‘‘+’’ to ‘‘×’’ every 20 trials. Participants were faster and more accurate when the target was ‘‘×’’ rather than ‘‘+,’’ and this advantage was even found when a centered dot was used to denote multiplication. In fact, this detection advantage is even found with novel, counterbalanced operators that have an order of precedence that is learned during an initial phase of the experiment. Thus, participants show a detection advantage for the notational symbol that has a higher order of precedence in the mathematic system that is learned, even when the assignment of the symbol to order of precedence is randomized.

As a third test of the tendency of notational elements with high order of precedence to attract attention, we employed a Flanker Task (Eriksen & Eriksen, 1974). In our instantiation of this task, participants saw expressions like ‘‘4 × 5 + 6 × 7’’ and simply had to respond as quickly as possible as to what the center operator was, pressing one key for ‘‘+’’ and another key for ‘‘×.’’ As Fig. 6 shows, there is a modest but statistically significant influence on accuracy of the congruence of flanking operators, but only when they had a

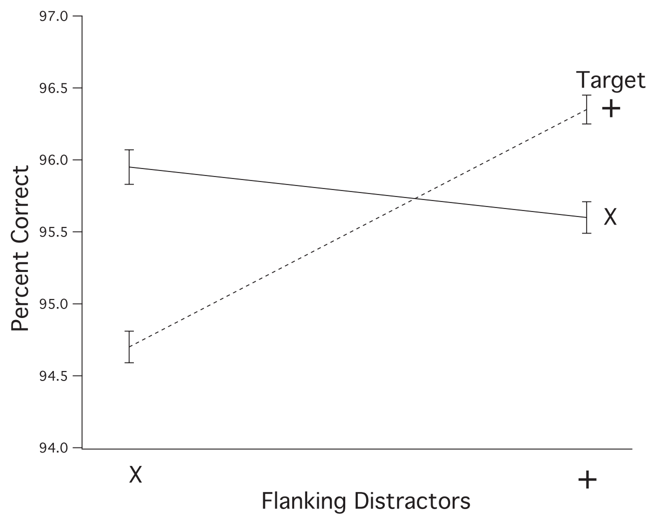

Fig. 6. Results from a Flanker Task by D. Landy and R. L. Goldstone (unpublished data). Making a response to a central operator is only appreciably hindered by incongruent peripheral operators when the central operator has a lower order of precedence.

higher order of precedence than the central operator. That is, participants are significantly less likely to give a correct response of ''+'' in the above incongruent expression than they are for the expression ''4 + 5 + 6 + 7'' in which the central and flanking elements are congruent in pointing toward ''+.'' This strong congruence effect was not found when ''×,'' the notation element corresponding to the operation with a high order of precedence, was in the center. Overall, response times were faster with ''+'' than ''×'' targets, arguing that this effect is not simply due to faster processing of the ''×.'' Apparently the ''×'' sign attracts more attention than the ''+'' sign in the context of mathematical expressions, even when no mathematical calculations are required.

4.3. Learning to visualize transposition

One final source of evidence in favor of RUPS for algebraic reasoning concerns reasoners' deployment of imagined movement events to solve math problems. When students first learn how to solve for a variable in a problem like ''$X + 4 = 9$'' they first learn that making the same change to both sides of an equality preserves that equality. Thus, it is formally sanctioned to subtract 4 from both sides of the equality, producing ''$X + 4 - 4 = 9 - 4$,'' which simplifies to ''$X = 5$.'' However, many students report that, after having gained experience with this principle, they learn to shorten solution time by simply moving the 4 from the side of the equation with the ''X'' to the opposite side, changing the sign of the 4 in the process. Although this transposition operation is highly intuitive, it is noteworthy that this kind of spatial transformation does not appear in most leading models of algebra (e.g., Anderson, 2005).

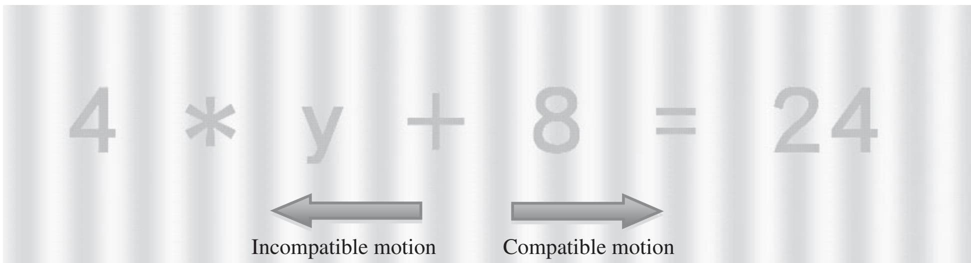

Fig. 7. As participants solved for the variable in equations like the above, a vertically oriented grating continuously moved either to the left or to the right. Although irrelevant for the task, when the movement of the grating was compatible with the movements of the numbers required by transposition, participants were more accurate.

To test whether the intuitively plausible spatial transposition strategy is naturally adopted by our participants, we presented displays like that shown in Fig. 7 to participants and had them solve for the variable. The equation was superimposed on top of a vertically oriented grating that continuously moved to either the left or right. The movement of the grating was either compatible or incompatible with the movement of numbers implicated by a transposition strategy. For the equation ''$4 * Y + 8 = 24$'' shown in Fig. 7, a rightward motion of the grating would be compatible with transposition because, in order to isolate Y on the left side, the 4 and 8 must be moved to the right. However, for the equation ''$24 = 4 * Y + 8$,'' a rightward motion would be incompatible. Participants solved the equations more accurately when the grating motion was compatible with transposition. Accuracy on incongruent and congruent motion trials were 95.2% versus 96.3%, respectively, a numerically small but significant difference, paired $t(56) = 2.5$, $p < .05$. No difference was found for response times.

The accuracy difference is consistent with a ''visual routines'' (Ullman, 1984) approach to mathematical cognition, according to which people engage in dynamic, visual–spatial routines to perform perceptual computations. Of particular relevance to the perceptual learning aspect of this transposition routine, we also found that participants who have taken advanced mathematics courses are more affected by the compatibility of the background motion than students with less experience. Participants reporting having taken calculus were substantially more accurate overall, mean error rate = 0.1, than non-calculus-takers, mean error rate = 0.27; $F(1, 56) = 7.9$, $p < .01$. Experienced participants were also more affected by the compatibility of the transposition motion and background motion than were non-calculus-takers, $F(2, 114) = 4.71$, $p < .05$. Accordingly, we conclude that the imagined motion strategy is a smart strategy that students come to adopt through experience with formal notations, rather than a strategy that students initially use while learning, and then abandon as their sophistication increases. Learned perceptual routines are not at odds with strong mathematical reasoning; they are likely the means by which strong mathematical reasoning is possible.

4.4. Recapitulation of RUPS for math

The theoretical upshot of this work is to question the assumption that mathematical cognition generally operates like formal systems of algebra or mathematical logic. The laws of algebra are formal in the sense of operating on the basis of axioms that are independent of visual form. Multiplication is commutative no matter what terms are multiplied and how they look. Mathematical cognition *could* have worked like this too. If it had, it would have provided a convenient explanation of how people perform algebra (Anderson, 2007). However, the conclusion from our experiments is that seeing ''$X + Y * A$'' cannot trivially be translated into the symbolic mentalese code * (+ (G001, G002), G003). Although it would be awfully convenient if a computational model of algebraic reasoning could aptly assume the transduction from visual symbols to mental symbols, taking this for granted would leave all of our current experimental results unexplained. We should resist the temptation to posit mental representations with forms that match our intellectualized understanding of mathematics. A more apt input representation to give a computational model (D. Landy, unpublished data) would be a visual–spatial depiction of notational elements that includes their absolute positions, spacings, sizes, and accompanying nonmathematical pictorial elements (such as the moving grating in the transposition experiment). Furthermore, much, perhaps all, of the processing of these mathematical expressions occurs in the same visual–spatial medium that houses this input representation. Processes like marking off elements as handled, combining terms to form simpler expressions, and moving, distributing, and factoring terms can all take place within this space with no translation to a purely symbolic form. The power of maintaining both the inputs and processes in this visual–spatial format is substantially amplified because perceptual processes can be customized by the reasoner to save labor and the need for executive control.

Our primary claim has been that people's mathematical operations that need not logically involve space or spatial transformations nevertheless do involve them. The operations most likely involve spatial processes that are not restricted to vision, but it is also clear that ongoing visual processing can facilitate or impede these spatial processes. Our results could be interpreted as arguing against the use of symbolic representations of mathematics, but others have argued that it is precisely perceptual processes that establish symbolic descriptions (Pylyshyn, 2000) and mathematical notations are, after all, symbolic expressions. Accordingly, rather than opposing symbolic processing of mathematics, we interpret our results as challenging conceptions of symbols as divorced from analog, symbolic information. In this respect, we offer a reinterpretation of Newell and Simon's (1963, 1976) influential ''Physical Symbol System Hypothesis.'' Their hypothesis was that physical symbol systems had the necessary and sufficient means for producing intelligent action. A symbol system includes both physical symbols such as marks on paper or punches on a computer tape, and the explicit rules for manipulating these tokens. In practice, all of their physical symbols were distantly related to their worldly referents, and were digital and discrete entities such as the strings ''P Q'' and ''GOAL 7 TRANSFORM L3 INTO LO.'' The arbitrary nature of these entities was by design because they wanted symbols to be able to designate any expression whatsoever without any a priori prescriptions or limitations. We concur with Newell and Simon's emphasis on physical symbols and believe in paying even more attention to symbols' physical attributes involving space, shape, and perceptual grouping. Accordingly, our revised physical symbol systems hypothesis is that symbols are not arbitrary, unconstrained tokens, but rather are organized into perceptual groups and processed using space. This conception of physical symbols makes them far more constrained than those underlying Newell and Simon's General Problem Solver, but these constraints are not only limiters, but permitters as well. For Specific Problem Solvers that are humans, it is a good policy to design symbols that can be processed efficiently given what we know about perceptual and cognitive mechanisms.

5. Conclusions

There is a striking difference between the sluggish rate with which human biology changes and the dizzying pace of progress in science and mathematics. In advocating an account of math and science that is grounded in perception and action, we must confront the fact that we are using the essentially same kinds of brains to understand advanced modern formalisms that have been used for millennia. Perceptual learning is at the core of our resolution to the discrepancy between scientific and neuro-evolutionary progress and a grounded account of the former. Our biology adapts at a plodding rate, but one thing that evolution has built into this biology is the ability to learn rapidly during an organism's life (Sterelny, 2003). We humans are prodigiously adept at programming ourselves to fit tasks at hand. As Clark (2009, p. 59) writes,

> We do not just self-engineer better worlds to think in. We self-engineer ourselves to think and perform better in the worlds we find ourselves in. We self-engineer worlds in which to build better worlds to think in. We build better tools to think with and use these very tools to discover still betters tools to think with.

A credible and worthy hope for education is to teach students to take the natural affordances of our long-tuned perceptual systems, which are at their core spatial and dynamic, and retask them for new purposes.

The end result of rigged-up perceptual systems is to have automatically deployed attentional and interpretational processes (see also Shiffrin & Schneider, 1977). This leaves open the question of automaticity of the *development* of these RUPS. The results from the transfer of the competitive specialization principle are suggestive that RUPS can be automatically developed as well as deployed; our students did not know that the competitive specialization principle would be relevant for a second task and yet they still created perceptual interpretations capable of transfer. Still, strategic control over the development of RUPS seems not only possible but pedagogically important. When a student knows that a particular interpretation or perceptual process supports a formal principle, then he or she can actively provide inputs and reinforcement to aid the development of the perceptual skill.

A RUPS perspective on pedagogical innovations potentially offers a new way of understanding why and how attempts to make difficult materials perceptually concrete may either promote or interfere with their understanding. The existing literature is mixed, sometimes suggesting that concrete or manipulative materials are useful for grounding abstractions, co-opting real-world knowledge, and motivating students (Carraher, Carraher, & Schliemann, 1985; Glenberg, Gutierrez, Levin, Japuntich, & Kaschak, 2004; Verschaffel, Greer, & De Corte, 2000), but other times suggesting that concrete materials limit generalization, distract students from essential principles, and interfere with symbolic interpretations (DeLoache, 1995; McNeil, Uttal, Jarvin, & Sternberg, 2009; Son & Goldstone, 2009b; Uttal, Liu, & DeLoache, 1999). If we think of education as training perception, then using perceptually rich objects is beneficial, but it is critical to develop perceptual routines that operate over these objects in effective ways. Typically, extraneous details of pedagogically motivated objects should be eliminated lest students develop routines that perseverate on these elements (McNeil & Jarvin, 2007). However, objects that invoke perceptual processes of selection, grouping, scanning, focus, and binding—processes that bridge to exportable principles—are likely to confer benefits. Diagrams are perceptual objects that are designed exactly with an eye toward these kinds of cognitive benefits. Representational techniques such as Euler Circles for logic, Cayley diagrams for group theory, and Feynman diagrams for quantum field theory are widespread in science and mathematics education at all levels. Often times, these techniques are construed simply as tools for translating abstract contents into concrete and intuitive formats. However, the RUPS interpretation is that they are also methods for changing perceptions so that students become sensitive to otherwise obscure and esoteric properties of a situation (Cheng, 2002). It is not simply that these diagramming processes can be used to explicate abstractions. The process of creating and perceptually interpreting the diagrams can substitute for the abstraction.

In light of RUPS, we can return to the traditional position that abstract reasoning is often times opposed to, and must overcome, potentially misleading perceptual resemblances. The results reviewed here suggest the alternative position that formal sensibilities can educate and enhance perceptual resemblances. Increasing reliance on rules and formal principles in lieu of phased-out perceptual information is not the only developmental trajectory. We can also modify our perceptual encodings so as to accord better with our formal understandings. Returning to the example of marsupial and placental wolves introduced in the first paragraph, the possibility raised by RUPS is for the biologist to develop informed, expert perceptual encodings that distinguish between these genetically distant species. Just as parents of identical twins show little difficultly distinguishing their children, learning perceptual relations between the ureter, bladder, and birth canal allows biologists to immediately and intuitively distinguish these animals. Scientific understanding does not merely trump the perception of resemblances. Scientific understanding shapes the perception of resemblances.

The phenomena of RUPS offer exciting possibilities for novel methods of improving educational outcomes. Together with a growing number of other researchers (Glenberg et al., 2004; Kellman, Massey, & Son, 2009; Kellman et al., 2008; Martin & Schwartz, 2005; Schwartz & Black, 1996), we argue that a highly effective way of facilitating sophisticated responding is by systematically training perception and action systems. Pedagogical methods that focus only on deep conceptual understanding, without supporting the perceptual-motor grounding of these understandings, risk creating inefficient and possibly cognitively inert knowledge. Pedagogical practices motivated by RUPS would include well-designed activities to alter the perceived similarity of situations, so that once dissimilar but importantly related situations become phenomenologically similar to one another with learning (Kellman et al., 2008). With techniques designed to promote perceptual learning, an educated student would truly be experiencing and creating new worlds.

Acknowledgments

The authors wish to express thanks to Robbie Jacobs, Phillip Kellman, and an anonymous reviewer for helpful suggestions on this work. This work was funded by National Science Foundation REESE grant 0910218. More information about the laboratory can be found at http://cognitrn.psych.indiana.edu

References

Anderson, J. R. (2005). Human symbol manipulation within an integrated cognitive architecture. *Cognitive Science*, *29*, 313–341.

Anderson, J. R. (2007). *How can the human mind exist in the physical world?* Oxford, England: Oxford University Press.

Barsalou, L. W. (2005). Abstraction as dynamic interpretation in perceptual symbol systems. In L. Gershkoff-Stowe & D. Rakison (Eds.), *Building object categories* (pp. 389–431), Carnegie Symposium Series. Majwah, NJ: Erlbaum.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, *59*, 617–645.

Beech, A., Powell, T., McWilliam, J., & Claridge, G. (1989). Evidence of reduced ''cognitive inhibition'' in schizophrenia. *British Journal of Clinical Psychology*, *28*, 109–116.

Bukach, C. M., Gauthier, I., & Tarr, M. J. (2006). Beyond faces and modularity: The power of an expertise framework. *Trends in Cognitive Sciences*, *10*, 159–166.

Cajori, F. (1928). *A history of mathematical notations*. La Salle, IL: Open Court Publishing Company.

Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.

Carraher, T. N., Carraher, D. W., & Schliemann, A. D. (1985). Mathematics in the streets and in schools. *British Journal of Developmental Psychology*, *3*, 21–29.

Carraher, D., & Schliemann, A. D. (2002). The transfer dilemma. *Journal of the Learning Sciences*, *11*, 1–24.

Cheng, P. C.-H. (2002). Electrifying diagrams for learning: Principles for effective representational systems. *Cognitive Science*, *26*(6), 685–736.

Chi, M. T. H., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121–152.

Clark, A. (2009). *Supersizing the mind*. Oxford, England: Blackwell.

DeLoache, J. S. (1995). Early understanding and use of symbols: The model model. *Current Directions in Psychological Science*, *4*, 109–113.

Detterman, D. R. (1993). The case for prosecution: Transfer as an epiphenomenon. In D. K. Detterman & R. J. Sternberg (Eds.), *Transfer on trial: Intelligence, cognition, and instruction*, (pp. 1–24) Westport, CT: Ablex.

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a non-search task. *Perception & Psychophysics*, *16*, 143–149.

Fahle, M., & Morgan, M. (1996). No transfer of perceptual learning between similar stimuli in the same retinal position. *Current Biology*, *6*, 292–297.

Fahle, M., & Poggio, T. (2002). *Perceptual learning*. Cambridge, MA: MIT Press.

Fine, I., & Jacobs, R. A. (2002). Comparing perceptual learning across tasks: A review. *Journal of Vision*, *2*, 190–203.

Fodor, J. A. (1992). *A theory of content and other essays*. Cambridge, MA: MIT Press.

Gauthier, I., Tarr, M. J., & Bubb, D. (2009). *Perceptual expertise: Bridging brain and behavior*. Oxford, England: Oxford University Press.

Gibson, E. J. (1991). *An odyssey in learning and perception*. Cambridge, MA: MIT Press.

Gibson, J. J., & Gibson, E.. J. (1955). Perceptual learning: Differentiation or enrichment? *Psychological Review*, *62*, 32–41.

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*, *15*, 1–39.

Glenberg, A. M., Gutierrez, T., Levin, J. R., Japuntich, S., & Kaschak, M. P. (2004). Activity and imagined activity can enhance young children's reading comprehension. *Journal of Educational Psychology*, *96*, 424–436.

Goldstone, R. L., & Barsalou, L. (1998). Reuniting perception and conception. *Cognition*, *65*, 231–262.

Goldstone, R. L., & Sakamoto, Y. (2003). The transfer of abstract principles governing complex adaptive systems. *Cognitive Psychology*, *46*, 414–466.

Goldstone, R. L., & Son, J. Y. (2005). The transfer of scientific principles using concrete and idealized simulations. *The Journal of the Learning Sciences*, *14*, 69–110.

Goldstone, R. L., & Wilensky, U. (2008). Promoting transfer through complex systems principles. *Journal of the Learning Sciences*, *17*, 465–516.

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (Ed.), *Problems and projects*, (pp. 437–447). New York: The Bobbs-Merrill Co.

Kaminski, J. A., Sloutsky, V. M., & Heckler, A. F. (2008). The advantage of abstract examples in learning math. *Science*, *320*, 454–455.

Kellman, P. J., Massey, C. M., Roth, Z., Burke, T., Zucker, J., Saw, A., Aguero, K., & Wise, J. (2008). Perceptual learning and the technology of expertise: Studies in fraction learning and algebra. *Pragmatics & Cognition*, *16*(2), 356–405.

Kellman, P. J., Massey, C. M., & Son, J. Y. (2009). Perceptual learning modules in mathematics: Enhancing students' pattern recognition, structure extraction, and fluency. *Topics in Cognitive Science*. DOI: 10.1111/j.1756-8765.2009.01053.x.

Landy, D., & Goldstone, R. L. (2007a). How abstract is symbolic thought? *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *33*, 720–733.

Landy, D., & Goldstone, R. L. (2007b). Formal notations are diagrams: Evidence from a production task. *Memory & Cognition*, *35*, 2033–2040.

Landy, D. H., Jones, M. N., & Goldstone, R. L. (2008). How the appearance of an operator affects its formal precedence. In *Proceedings of the Thirtieth Annual Conference of the Cognitive Science Society* (pp. 2109–2114). Washington, DC: Cognitive Science Society.

Leeper, R. (1935). A study of a neglected portion of the field of learning: The development of sensory organization. *Journal of Genetic Psychology*, *46*, 41–75.

Lovaas, O. I., Koegel, R. L., & Schreibman, L. (1979). Stimulus overselectivity in autism: A review of research. *Psychological Bulletin*, *86*, 1236–1254.

Martin, T., & Schwartz, D. (2005). Physically distributed learning: Adapting and reinterpreting physical environments in the development of the fraction concept. *Cognitive Science*, *29*, 587–625.

McNeil, N. M., & Jarvin, L. (2007). When theories don't add up: Disentangling the manipulatives debate. *Theory into Practice*, *46*, 309–316.

McNeil, N. M., Uttal, D. H., Jarvin, L., & Sternberg, R. J. (2009). Should you show me the money? Concrete objects both hurt and help performance on mathematics problems. *Learning and Instruction*, *19*, 171–184.

Newell, A., & Simon, H. A. (1963). GPS: A program that simulates human thought. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought* (pp. 312–344). New York: McGraw-Hill.

Newell, A., & Simon, H. A. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the ACM*, *19*, 113–126.

Notman, L. A., Sowden, P. T., & Özgen, E. (2005). The nature of learned categorical perception effects: A psychophysical approach. *Cognition*, *95*, B1–B14.

Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 259–303). Hillsdale, NJ: Erlbaum.

Postman, L. (1955). Association theory and perceptual learning. *Psychological Review*, *62*, 438–446.

Pylyshyn, Z. (2000). Situating vision in the world. *Trends in Cognitive Science*, *4*, 197–207.

Quine, W. V. (1977). Natural kinds. In S. P. Schwartz (Ed.), *Naming, necessity, and natural kinds* (pp. 155–175). Ithaca, NY: Cornell University Press.

Rumelhart, D. E., & Zipser, D. (1985). Feature discovery by competitive learning. *Cognitive Science*, *9*, 75–112.

Schwartz, D. L., & Black, J. B. (1996). Analog imagery in mental model reasoning: Depictive models. *Cognitive Psychology*, *30*, 154–219.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, *84*, 127–190.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*, 3–22.

Son, J. Y., & Goldstone, R. L. (2009a). Fostering general transfer with specific simulations. *Pragmatics and Cognition*, *17*, 1–42.

Son, J. Y., & Goldstone, R. L. (2009b). Contextualization in perspective. *Cognition and Instruction*, *27*, 51–89.

Sowden, P. T., Davies, I. R. L., & Roling, P. (2000). Perceptual learning of the detection of features in X-ray images: A functional role for improvements in adults' visual sensitivity*? Journal of Experimental Psychology: Human Perception and Performance*, *26*, 379–390.

Sterelny, K. (2003). *Thought in a hostile world: The evolution of human cognition*. Oxford, England: Blackwell.

Ullman, S. (1984). Visual routines. *Cognition*, *18*, 97–159.

Uttal, D. H., Liu, L. L., & DeLoache, J. S. (1999). Taking a hard look at concreteness: Do concrete objects help young children learn symbolic relations? In C. S. Tamis-LeMonda (Ed.), *Child psychology: A handbook of contemporary issues* (pp. 177–192). Philadelphia, PA: Psychology Press.

Verschaffel, L., Greer, B., & De Corte, E. (2000). *Making sense of word problems*. Lisse, The Netherlands: Swets & Zeitlinger.