

Forthcoming in Vosniadou, S. (in press) Handbook of Conceptual Change. Erlbaum.

DYNAMIC TRANSFER AND INNOVATION

Daniel L. Schwartz and Sashank Varma

Stanford University

Lee Martin

University of California, Davis

Daniel L. Schwartz  
485 Lasuen Mall  
Stanford, CA 94305-3096  
(650) 736-1514  
[Daniel.Schwartz@stanford.edu](mailto:Daniel.Schwartz@stanford.edu)

Sashank Varma  
450 Serra Mall, Building 160  
Stanford, CA 94305-2055  
sashank@stanford.edu

Lee Martin  
School of Education, UC Davis  
One Shields Ave  
Davis, CA 95616  
(650) 804-4539  
leemartin@ucdavis.edu

## ABSTRACT

Transfer occurs when people use learning from one situation in another, for example, when they use their knowledge of water flow to understand electrical current. Transfer is relevant to conceptual change, because new concepts build on a foundation of prior learning. The contribution of the transfer literature to conceptual change comes from its explanations of why people retrieve and change their ideas in some contexts, but not others. Because transfer can precipitate changes to the environment as well as concepts, we use the broader term innovation of which conceptual change is a subset. Transfer addresses two critical issues for innovation: the knowledge problem and the inertia problem. The knowledge problem asks how prior knowledge of one sort can contribute to creating new knowledge of another sort. The inertia problem asks why people often fail to innovate, even though they have the relevant prior knowledge. People often transfer for repetition rather than innovation. The chapter proposes two types of transfer to address these problems. Similarity transfer occurs when people are able to recognize that they have well-formed prior ideas that can be profitably used to describe another situation in a new way. Dynamic transfer occurs when component competencies are coordinated through interaction with the environment to yield novel concepts or material structures. The two types of transfer can work together so that people transfer the idea of being innovative when it is appropriate to do so.

**KEYWORDS:** Similarity Transfer, Dynamic Transfer, Mechanisms of Transfer, Measures of Transfer, The Inertia Problem, The Knowledge Problem, Preparation for Future Learning, Sequestered Problem Solving, Conceptual Change, Innovation, Cognition, Context of Innovation

Conceptual change occurs when people learn a new way to think about a class of situations. There are gradations of conceptual change from the simple accumulation of enriching examples to total conceptual replacement (e.g., Tyson, Venville, Harrison, & Treagust, 1997). The history of science provides instances of grand conceptual change, for example, the discovery that disease is due to germs (Thagard, 1996). Developmental psychology has also noted a number of recurrent conceptual changes, for example, children's developing notion of "alive" (Carey, 1985). Adults too, can experience conceptual change, for example, when becoming parents.

Transfer occurs when people use learning from one situation in another. For example, individuals may transfer learning from school to kitchen and from algebra to chemistry. Transfer is relevant to conceptual change, because new ideas build on a foundation of prior learning. Yet, at the same time, people need to go beyond their original learning to accomplish a conceptual change. The product cycle of Intel, for example, requires a new product release every three months (COSEPUP, 2006). If the chip designers at Intel could not transfer in their prior learning to facilitate the design of new chips, Intel would be too inefficient to keep up with its competitors. At the same time, if the chip designers did not go beyond their prior knowledge, there would be no innovative product designs. Understanding transfer's relation to conceptual change may help create conditions that foster innovation.

Transfer research addresses two challenges for conceptual change: the *knowledge problem* and the *inertia problem*. The knowledge problem asks how prior knowledge can make innovative new knowledge. Through what processes can old concepts of one sort possibly create new concepts of another sort? The transfer literature proposes two

different solutions called *similarity transfer* and *dynamic transfer*. In similarity transfer, people apply well-formed concepts from one situation to explain another situation in a novel way. For example, people often use causal explanations for the “sophomore slump” in sports, where rookies who excel in their first season do worse the next year (e.g., early success led to complacency). By transferring in their knowledge of statistics, people may re-conceptualize the drop in performance as an instance of regression to the mean. In dynamic transfer, the context helps people coordinate component abilities to create a novel concept in the first place. Thelen and Smith (1994), for example, describe how infants learn to coordinate their physical actions of looking and reaching, which leads to the development of object permanence (the appreciation that an object is still there even though one may not see it). We describe the two forms of transfer more fully below.

While the knowledge problem asks how people change concepts, the inertia problem addresses why people do not always do so. Transfer is a double-edged sword for change. People can transfer in prior learning to support change, but they can also transfer in routines that prevent change. In Luchins and Luchins’ (1959) classic studies of *einstellung*, or rigidity of behavior, people learned a complex method for measuring water using several jars in sequence. Once they had mastered the multi-jar method, they received simpler problems that could be solved more efficiently with only one or two jars. People transferred the complex method and never considered searching for an alternative.

The inertia problem highlights the tension between using, yet going beyond prior learning. An excellent example comes from children learning fractions. Children could hardly learn to compare two-out-of-three wins versus four-out-of-seven wins, if they had

never learned about natural numbers. New ideas do not arise in a vacuum. At the same time, children have difficulty interpreting “ $2/3$ ” as a ratio. Instead, they interpret the “2” and “3” as natural numbers (Kerslake, 1986). They over-rely on their well-practiced and well-understood natural number schema.

The inertia problem asks what conditions lead people to transfer for change rather than repetition. Before people have made a conceptual innovation, it is hard to know whether it is worth pursuing – people cannot know the concept before they have it. Moreover, people’s prior learning typically yields routines that are good enough for “getting by”. Consequently, people often do not recognize, seek, or risk alternatives. Similarity and dynamic transfer provide different routes by which people can overcome the problem of inertia. In the final section of the paper, we consider the inertia problem in the context of creating educational experiences that can help people transfer the idea of being innovative.

To develop the relation between transfer and innovation, the chapter relies on two sets of empirical examples. One set involves mathematics. We choose mathematics because it has been an important focal point for testing theories of transfer. Mathematics, like logic, is general and widely applicable across many domains, but people do not always transfer mathematics when relevant. Moreover, the development of mathematical reasoning depends on many conceptual changes, such as the shift from counting whole quantities to working with ratios. Thus, mathematics is an ideal content area for examining the mechanisms of transfer for innovation. Our second set of examples involves tools, including physical and representational tools. Whereas mathematics leads us to focus on conceptual innovations, tools lead us to focus on material innovations.

The first section sketches the relations between innovation, conceptual change, and transfer. The second section offers a brief review of similarity transfer, which has been the primary focus of transfer research. From there, the chapter moves into new terrain. The third section consolidates a number of findings to introduce dynamic transfer and argue for its relevance to both transfer and conceptual change. The final section discusses the inertia problem by introducing a representation of context that helps delineate different trajectories that do and do not lead to conceptual change.

### INNOVATION AND TRANSFER

Transfer is unique among the many approaches to conceptual change, because transfer is intimately concerned with the context of cognition. Transfer research tries to explain why people can retrieve and change their ideas in some contexts, but not others. In his illuminating discussion on the evolving concept of disease, Thagard (1996) concludes, “A full theory of conceptual change must integrate its representational, referential, and social aspects” (p. 477). By referential, Thagard means how concepts interact with context, as opposed to how concepts are represented. Here, we address the referential aspects of conceptual change. In fact, we broaden the discussion of conceptual change to include structural changes to the environment, which can feed back and change people’s mental representations. To avoid diluting the term conceptual change, we adopt the broader term innovation, which can refer to new ideas and new material structures.

The strength of the transfer literature for innovation is in its explication of how people and contexts interact to achieve an innovation. Unlike the conceptual change literature, a precise definition of a conceptual change or innovation is not central to the

transfer literature. Nevertheless, it is useful to have a working definition, so we propose *prior inconceivability* and *generativity* as criteria for innovations.

Prior inconceivability means that the changes that result from an innovation could not be prefigured prior to the change. For example, for children who have already learned two-digit addition, learning three-digit addition is a conceivable extension. For children who have only learned addition, comparing ratios is incomprehensible. Learning to do so would constitute an innovation. Generativity means that innovations can extend beyond the specific situations and problems that led to their initial development. For example, a fully developed concept of ratio can extend to quantities that did not appear in the conditions of original learning – from fractions to proportions and percentages.

The term innovation can refer to an outcome or to a process. Innovative outcomes come at two scales. Grand innovations are those which are novel on a global scale, and which were previously unknown to most everyone, such as the telephone and the theory of evolution. Petite innovations are local and new to a given person. Over history, many different people can achieve the same petite innovation. For example, most children learn that the amount of water in wide glass does not change when poured into a narrow glass, even though the visual height changes. The prevalence of this “conservation” concept does not diminish the innovative value of the concept for any one child.

The process of innovation refers to the mechanisms involved in restructuring thought or the environment. The process of innovation applies to both grand and petite innovations. This is important, because grand innovations are rare events, and their low frequency and unpredictability makes it hard to develop an empirically grounded account

of innovation (Johnson-Laird, 1989). However, because petite innovations regularly recur across people, it is possible to study the processes of these innovations. Ideally, findings from research on petite innovations can help create conditions for grand innovation. One important goal of education is to make the next generation of innovators – innovators who can make petite innovations in response to a rapidly changing world, and those who can make grand innovations that rapidly change the world.

Transfer mechanisms are relevant to the process of innovation, because they enable the application of prior learning for the purposes of innovation. The main body of transfer research has not examined how mechanisms of transfer support innovation. Detterman (1993), for example, defined transfer as the “the degree to which a new behavior will be repeated in a new situation” (p. 4). Replicating a behavior is quite different from innovating a new one. Transfer research has typically focused on the transfer of prior knowledge or behavior to improve speed and accuracy on a novel task (Schwartz, Bransford, & Sears, 2005). Bassok and Holyoak (1989), for instance, taught students an algorithm to compute arithmetic progressions. Half of the students learned the algorithm as a math lesson and half as a physics lesson. Students who received the algorithm in the math lesson were more likely to transfer it to a new problem domain (e.g., finances). Students who did not transfer the algorithm could still solve the new problem. They were just slower, because they used an iterative technique. Their use of an inefficient method indicated they had failed to transfer. Asking how people become more efficient by repetition is different from asking how people innovate by building on prior knowledge. Nevertheless, there is research on transfer that is relevant to innovation, which we describe next.

## SIMILARITY TRANSFER AND THE KNOWLEDGE PROBLEM

### Processes of Innovation and Similarity Transfer

In similarity transfer, people already possess well-formed prior knowledge developed for one situation, and they use this knowledge to understand a different situation in a new way. For example, scientists might learn to view traffic congestion in terms of their knowledge of fluid dynamics. The conceptual change does not occur to their existing knowledge of fluid dynamics; rather, they change their concept of freeway congestion. We call this similarity transfer, because the key move involves recognizing that two situations or ideas are similar, even though they may not appear so at first.

Similarity transfer occurs when people realize that what they learned for one situation can be used for another situation. Similarity transfer depends on well-formed prior knowledge. Many purported demonstrations of failed transfer actually were not failed transfer – people never learned the concepts or skills to start with, so there was no way they could fail to transfer. A genuine failure of similarity transfer occurs when people have “inert knowledge” (Whitehead, 1929). Inert knowledge is a description of the situation where people have relevant knowledge but do not spontaneously apply it. The field has made several important discoveries about the types of experiences and knowledge structures that help people avoid inert knowledge. These include the opportunity to discern commonalities and differences using multiple examples (e.g., Gick & Holyoak, 1983; Schwartz & Bransford, 1998), explicate the general principal behind the examples (e.g., Brown & Kane, 1988), and experience the problem for which the transferable knowledge would be the solution (e.g., Bransford, Franks, Vye, & Sherwood, 1989).

In similarity transfer, people have knowledge gained in one context that is sufficient for application to another. The question is whether they will use that knowledge in the new context. To make the transfer, people need to detect the similarity between one situation and another. One type of similarity occurs at the level of perceptual *surface features* or “identical elements” (Thorndike & Woodworth, 1901). Chi, Feltovich, and Glaser (1981), for example, reported that undergraduates taking introductory physics classified problems by surface features (e.g., pulley problems, spring-mass problems). A focus on surface features can cause individuals to transfer the wrong ideas across problems. For example, the undergraduates would likely *fail* to transfer solutions across spring and pulley problems, because they see them as unrelated. They are also likely to exhibit the *negative* transfer of a solution from one pulley problem to another, even if the problems do not involve the same principles (cf., Ross, 1989)

A second type of similarity involves identical relations or *deep features*. For example, in contrast to undergraduates, the physics graduate students in Chi, Feltovich, and Glaser’s (1981) study classified spring and pulley problems according to the underlying principles needed to solve the problems. The graduate students could appreciate that the relations between the objects within each problem involved the same principles, even though they looked quite different. When people can recognize relational or structural similarities, they are able to make the transfer across situations that have surface level differences. For example, they can transfer principles from pulley to spring problems (when appropriate).

Similarity transfer, particularly when the nature of the similarity is relational, can support innovation. The most extensively studied use of similarity transfer for innovation

comes from studies of *analogical transfer* (for representative studies on analogical transfer see Gick & Holyoak, 1983; Reed, Ernst, & Banerji, 1974). Kepler managed to develop his theory of planetary motion by drawing an analogy from the light of the sun (Gentner et al., 1997). He posited that the sun projected a force to the planets in a similar manner that light emanated from the sun. From there, he further drew an analogy between the amount of light that reached a planet through its movement and the amount of force that kept it in place. This is an example of a grand innovation. Kepler's introduction of the concept of force explained elliptical orbits. Kepler's innovation was also generative. His theory could predict orbits of planets wherever they might be, so it went beyond the specific instances he worked with.

Researchers often distinguish *near* and *far* transfer, which is another way of stating the experienced similarity of two situations (for factors that affect experienced similarity, see Barnett & Ceci, 2002). Generally speaking, transfer based on surface features is considered near, and transfer based on deep features is considered far. Innovations, almost by definition, depend on far transfer. If the transfer were too near, then it would not illuminate a different situation in a novel way.

Another useful distinction involves *spontaneous* and *prompted* transfer. Instructors can often prompt far transfer by introducing analogies that help students innovate new (to them) concepts. For example, an instructor might use the analogy of water in a pipe to explain the invisible processes of electricity. In contrast, spontaneous transfer occurs when people make the analogy themselves without any external support. Using an analogy to help explain a concept to someone else is quite different from learners spontaneously generating the analogy on their own. Spontaneous, self-directed

transfer across highly different domains or contexts is infrequent (Detterman, 1993). It took Kepler a lifetime to work out the analogy between light and gravity.

One possible reason that spontaneous and innovative far transfers are infrequent is that people often do not have enough knowledge about a second domain. Scholars, for example, often have very precise knowledge about the topic for which an innovation would be useful. However, they may not have sufficient expertise in any other domains to generate candidate analogies. The analogies they can generate from other domains will not have sufficient precision and structure to map into the complexity of the situations they are trying to explain.

If we expand the unit of analysis for similarity transfer beyond the individual, we find more instances of innovation. Biologists, for example, have extremely detailed knowledge about the problems they need to solve and the constraints on what would make a good solution. Dunbar (1997) found that biology labs have computer programs that help scientists generate candidate analogies. The biologists enter specific characteristics of the biological compounds and interactions they are trying to understand, and the computer program generates homologs (i.e., other biological structures that have similar properties), and the biologist can draw the analogy between the homologs to see if they work. A second example comes from companies that use on-line programs to solicit analogies from different disciplines (Feller, Fitzgerald, Hissam, & Lakhani, 2005). For example, a pharmaceutical company might post a drug compound that has side effects they cannot explain. People from other disciplines can look at the posting. If they can generate a solution, the company might pay them for it. For example, a mathematician might recognize the drug problem as an instance of mathematical knot theory.

## Similarity Transfer as a Measure of Innovative Outcomes

The term “transfer” has two uses. Thus far, we have described transfer as a mechanism that leads to an innovative process. Transfer can also mean a transfer task specifically designed to measure outcomes.<sup>1</sup> Transfer tasks are useful because they can measure the generality of an innovative outcome or conceptual change. For example, if a child shows conservation of liquid, measures of transfer allow us to ask whether the conservation concept extends to the conservation of solids (e.g., a ball of clay that thins as it stretches).

Similarity transfer uses a specific task structure to measure learning outcomes. Bransford and Schwartz (1999) called this type of measure Sequestered Problem Solving (SPS). In SPS assessments of transfer, people receive a new problem without any resources or interactive opportunities to learn. The assessment detects the robustness and generality of people’s learning. SPS measures provide a useful index for evaluating a purported conceptual change. We consider two kinds of conceptual change – *global* and *domain-specific*.

Global conceptual changes are the result of broad cognitive reorganizations that necessarily transfer across all relevant situations. A simple analogy can be drawn to the literature on perceptual development. As infants develop binocular vision over the first weeks of life, their brain reorganizes to see the world in depth, and children subsequently see depth in all contexts. Similar arguments have been made about conceptual changes that involve reasoning. Piaget (1952), in particular, argued that children go through a series of stages that involve increasingly complex reasoning structures. Once children achieve a particular stage in their reasoning, they will see the world in those terms. For

example, one important development is the ability to conserve quantities. An example of conservation involves knowing that an amount of clay does not change, even as one stretches it from a ball into a long rope. To test the claim that the ability to conserve is a global change, one would use SPS measures to see if children transfer conservation across a number of diverse tasks, such as when water is poured from a wide glass into a narrow one. If children fail to transfer conservation across problems, then the claim that there is a global conceptual change fails.

Transfer tasks can also illuminate domain-specific changes. Researchers can learn about the generality and precision of a concept by seeing how far it transfers. For example, very young children develop a rudimentary concept of aliveness (Keil, 1989), an instance being the development of the distinction between inanimate and dead (Carey, 1985). What is the nature of children's concepts? Some authors propose that children's early concepts have a theoretical quality, in the sense there is a causal coherence, underlying abstractions, and sensitivity to negative evidence (Gopnick & Meltzoff, 1997). Okita and Schwartz (2006) examined this proposition by using SPS transfer tasks. They asked 3- to 5-year-olds questions that probed whether robotic dogs were alive. It was an SPS assessment because the children had no opportunity to learn during the assessment, for example, by interacting with the robots to get further information. The researchers showed children a dancing robotic dog and asked questions including whether it felt hunger, whether it would run away if there were a fire, whether it could "wake-up" without a remote control, and so forth.

Three-year-old children extended the concept of alive to the robotic dogs by giving the same answers they give for regular dogs. These younger children also

extended the concept of alive to stuffed animals. In contrast, 5-year-olds did not extend all aspects of alive to the robotic dogs or stuffed animals. They believed the robotic dogs needed a remote control and would not grow, but they still thought the robotic dogs felt hunger and would wake up in a fire. These older children were also less likely to indicate attributes of aliveness for stuffed animals. By measuring what children spontaneously transferred to make sense of the robotic dogs, the researchers found that children begin with an undifferentiated concept of alive that did not exclude anything with surface features resembling dogs. Thus, the concept was not a theory in the sense of being applied differentially to empirical situations. Moreover, with development, the children's concept of alive changed in piecemeal fashion, rather than exhibiting the coherence of a true theory. The researchers concluded that children's concept of alive was more script-like than theory-like. That is, children recalled interaction scripts with living things, and this led to the organization of their attributions of alive. For example, older children knew the script for using remote controls, so they merged this script into their other scripts for interacting with pets (e.g., feeding them).

Similarity measures of transfer that use SPS assessments can reveal the nature of people's concepts and the outcomes of intellectual innovation. However, SPS assessments are not ideal for measuring the process of conceptual change, because they do not include opportunities to learn and change during the assessment. Below, we describe a different measure of transfer that is more suited to evaluating the process of innovation.

## DYNAMIC TRANSFER AND THE KNOWLEDGE PROBLEM

### The Processes of Innovation and Dynamic Transfer

Dynamic transfer tries to explain how prior knowledge can create concepts that did not previously exist. While similarity transfer explains how people map well-formed concepts to structure new situations, it does not explain how initial concepts develop in the first place. In dynamic transfer, people bring component competencies and situations into coordination to learn a new concept.

Dynamic transfer is one of many different mechanisms by which people learn new ideas. Dynamic transfer has not been the focus of traditional transfer research, which has instead focused on similarity transfer. Although dynamic transfer is different from traditional investigations of transfer and seeks to explain different phenomena, we still maintain the label transfer for three reasons. The first is that it depends on reusing prior knowledge in new contexts, which is the hallmark of transfer. The second reason is that the term is already in use by physics educators who have been working on the challenge of changing students' intuitive physical concepts into normative physics concepts (e.g., Dufresne, Mestre, Thaden-Koch, Gerace, & Leonard, 2005; diSessa & Wagner, 2005; Hammer, Elby, Scherr, & Redish, 2005; Mestre, 2004; Rebellow et al., 2005). The third reason is that there are a number of critics of similarity transfer who have urged a broader consideration of how the environment, which includes social and physical aspects, helps people coordinate productive actions (e.g., Carraher & Schliemann, 2002; Greeno, Smith, & Moore, 1993; Hutchins, 1995; Pea, 1993; Suchman, 1987). We agree with these investigators that similarity transfer does not address these considerations, but we do not think this is sufficient reason to abandon the many important findings and theories from similarity transfer. By juxtaposing dynamic and similarity transfer, we want to emphasize that they are not mutually exclusive or opposing theoretical camps. They are

simply different mechanisms of innovation, and both can be at play in the same innovation, as we describe in the final section on learning to be innovative.

There are a number of distinctions between similarity and dynamic transfer. The first involves the prior knowledge that precipitates the transfer. Similarity transfer depends on well-formed prior knowledge, so people can re-conceptualize the new situation in terms of the prior knowledge. In contrast, dynamic transfer depends on component knowledge, skills, or abilities that become well-formed during the process of transfer.

The second distinction involves the crux move in achieving the transfer. In similarity transfer this involves seeing that “this is like that,” while dynamic transfer involves seeing that “this goes with that.” Similarity transfer hinges on detecting the similarity between situations so people can apply their ideas. Dynamic transfer hinges on coordinating systems that may be quite dissimilar.

The third distinction involves the role of the context in causing transfer. In similarity transfer, the environment cues the retrieval of intact prior knowledge. In dynamic transfer, the environment coordinates different components of prior knowledge through interaction. It is important to note that similarity transfer includes a period “mapping” – once people have a candidate analogy, they figure out which parts of the analogy work for the target situation (Gentner et al., 1997). Mapping differs from coordination, because dynamic coordination depends on an environment that provides feedback and/or structuring resources to join ideas together. In contrast, mapping similarities can occur in thought alone.

Before discussing dynamic transfer in more detail, we offer two examples. The interactive process of dynamic transfer often has the quality of trial and error, and it can look highly inefficient. However, when coupled with the right resources, the interactions may result in a coordination that can be an innovation. In the first example, interactions with the material world lead to an innovation in symbolic thought. In the second example, interactions with symbols lead to an innovation in thinking about the material world.

### Two Examples of Dynamic Transfer

Our first example of dynamic transfer involves 9- and 10-year-old children learning to add simple fractions like  $1/4 + 1/2$  (Martin & Schwartz, 2005). Half of the children learned using pie pieces and half learned with tile pieces, as shown in Figure 1. After instruction, both groups could solve the problems using their respective materials, but they could not solve the problems in their heads. For the transfer task, the children received similar fraction addition problems, but they had to use novel materials (e.g., beans, fraction bars). The question was how the children from the two conditions would transfer their prior learning to work with the new materials. The Tile students exhibited a successful dynamic transfer, and the Pie students exhibited a failed similarity transfer.

---

Insert Figure 1 about here.

---

Figure 2 provides a schematic of the two courses of transfer for the Tile and Pie children. The vertical dimension indicates how frequently children gave accurate verbal answers; for example, given the problem  $1/2 + 1/4$ , they state “three-fourths”. The horizontal dimension indicates how frequently the children arranged the physical materials correctly; for example, for  $1/2 + 1/4$ , the children arrange the one-half as two

groups with two pieces each. The arrows indicate the average trajectory of learning over time for the two groups.

---

Insert Figure 2 about here.

---

The Tile children showed a stable trajectory that ended in the top right corner of near-perfect verbal and physical performance. The Tile children first improved in how they structured the physical materials, before they began to give the correct symbolic answers. They exhibited a dynamic transfer, because their interactions with the environment helped them slowly coordinate their (verbal) understanding. In contrast, the Pie children did not reach high levels of accuracy on either dimension, nor were they stable: they often regressed from one trial to the next. In addition, unlike the Tile children, the Pie children were better at giving verbal answers than arranging the materials. What explains the poor performance of the Pie children? During their initial learning, they had implicitly relied on the part-whole structure of the pies, and the new materials did not exhibit the circular whole. These children could not make a dynamic transfer through interaction, because they had never explicitly learned to impose a whole to turn discrete elements into fractions – they always got the whole “for free” when working with the pies. So instead of completing a dynamic transfer, they attempted a similarity transfer by trying to remember how they solved similar prior problems, and they could not use the structure of the novel materials to help them learn.

Our second example involves a dynamic transfer where children interact with the symbols of mathematics. Figure 3 provides an instance of a balance scale problem used in the research (Schwartz, Martin, & Pfaffman, 2005). The children needed to answer how the balance scale behaves when the hand lets go. Children under 12 typically solve

the problem by focusing on the weights or distances, but not both. For example, they might say that it will tilt to the left, because there are more weights on the left side, or it will tilt to the right, because the weights are farther out on the right side. The children do not coordinate the two dimensions of weight and distance simultaneously, and therefore, they cannot infer that the scale will balance. Learning to think about two dimensions simultaneously heralds an important innovation in how children can reason about many complex situations.

---

Insert Figure 3 about here.

---

Children from 9- to 11-years-old solved a series of increasingly complicated problems. For each problem, they made a prediction about what would happen, and they had to type in a justification for their prediction. They then saw the correct answer. The critical manipulation was that half of the children were told to justify their answers using words. The other half of the children were told to justify their answers using math. The motivation behind this contrast was that the Math children might transfer in their knowledge of math to help coordinate and structure their thinking about the scale. The idea that the developing child, like a scientist, might be able to use mathematics, diagrams, and other explicit representations to discover and organize complex empirical relations, comes from Vygotsky's (1987) foundational insight that cultural forms mediate development. The current experimental manipulation examined the significance of mathematics, a cultural form, in helping children develop complex physical knowledge about balance.

Children in the Word condition did not make much headway. They typically justified an answer based on one dimension, for example, "More weight." If they saw

they had the wrong answer, they would switch to the other dimension, for example, “More distance.” The Math children, in contrast, progressively learned to consider both dimensions simultaneously. Here is a prototypical sample of responses for a 10- to 11-year-old. At first, the child justified her answer with “ $3 > 2$ ,” focusing on only one dimension. Later, on a new problem, the child wrote, “ $3 + 3 = 4 + 2$ .” She was considering both dimensions of information and all four of the relevant parameters simultaneously (two weights and two distances). When she saw that she made the wrong prediction, her justification shifted to “ $3 - 3 < 4 - 2$ .” Mathematics provided candidates for structuring the problems, so she could switch from addition to subtraction. After a few more struggles, this child discovered that multiplication works, and she was able to solve all the subsequent problems (e.g.,  $4 \times 3 = 3 \times 4$ ).

In a posttest, the Word children performed just like other children their age. Only 36% of the 10- to 11-year-old Word children tried to solve the problems by considering both dimensions. In contrast, 95% of the 10- to 11-year-old Math children considered both dimensions in their answers. Results from the 9-year-old Math children are particularly relevant for distinguishing similarity and dynamic transfer. Very few of the 9-year-olds discovered the correct solution, so they never developed a mathematical formula for solving the problem. This means they never succeeded in seeing the similarity between the balance scale and multiplication structures. Nevertheless, 86% of the Math 9-year-olds reasoned about two dimensions simultaneously on the posttest. A common line of reasoning for these children was, “This side has more weight, but this side is closer.” Even though the younger children did not find the similarity between

multiplication equations and the balance scale, in the process of dynamically transferring math, they learned a more complex, qualitative way to approach this physical system.

### Examples of Coordination and Innovation

A nagging question for conceptual change asks, “Where do new concepts come from?” Theories of similarity transfer have difficulty answering this question. Through the mechanisms of similarity transfer, new concepts arise by building bridges between existing concepts, allowing one concept to inform another. However, if we only entertain similarity transfer, then all new concepts are links between existing concepts, and it is impossible to break out of this closed system to create something genuinely new.

Theories of dynamic transfer are better suited to answer the question of where new concepts come from, as the coordination of different systems or bodies of knowledge can yield innovations that are more than the sum of their parts. Dynamic transfer can create new “coordination classes” (diSessa & Wagner, 2005). By examining how basic components can combine to give rise to more complex concepts, we can begin to understand how new concepts arise. We provide examples from three domains: machine learning, animal learning, and human learning.

A suggestive example comes from early work on connectionist networks. Connectionist networks showed impressive learning, but they soon ran into a problem that required the network to undergo a “conceptual change.” Early models used two layers of nodes. When a stimulus is presented to the input nodes, the network gives the response on the output nodes. The network learns by adjusting the strength of the associations among the input and output nodes. These two-layer networks could learn to reason about conjunctions and disjunctions of stimuli, but they could not simultaneously

learn to handle exclusive disjunctions (A or B, but not both; Papert, 1988). No amount of learning enabled the connectionist networks to make the conceptual change necessary to reason about exclusive disjunctions. Later, researchers solved the problem by interspersing a third (so-called “hidden”) layer to coordinate the input and output layers (Rumelhart, Hinton, & Williams, 1986). With the added layer and appropriate feedback during learning, the system learned to compute conjunctions, disjunctions, and exclusive disjunctions. Thus, by coordinating the two layers with a third layer that had the same structure as the first two layers, the system was able to exhibit a “conceptual change.”

A second example comes from work with Macaque monkeys. Macaques do not exhibit tool use in the wild. They have no “innate” concept of tool. Ishibashi, Hihara, Takahashi, Heike, Yokota, and Iriki (2002) successfully trained a monkey to use a small rake to retrieve food. Afterwards, the scientists gave the monkey a transfer task. The monkey received a rake that was too short to reach the food, but it was long enough to reach a longer rake that could not be grabbed directly. The monkey used the short rake to grab the long rake, and then used the long rake to grab the food. The monkey had developed a sufficiently general “concept” of “reaching-tool” that it could think to see whether the long rake could achieve the goal. When the scientists examined the brain of the monkey, they found a set of rich neuronal projections that connected visual and motor regions. Macaques do not normally have such extensive projections (though humans do). However, they do have a latent gene that can express the neurotransmitters for growing such projections, enabling the coordination of these brain regions. The external pressure of trying to use the tool to retrieve the food caused the gene to express itself, and enabled

the monkey to coordinate its visual and motor systems. The monkey coordinated two pre-existing neural systems to develop a “concept” of tool.

A third example comes from Case and colleagues who examined how children developed a robust concept of natural numbers (e.g., Case & Okamoto, 1996). He theorized that young children need to coordinate a number of distinct systems. For example, infants (and animals) have the ability to make magnitude comparisons of perceptual input such as loud and soft, bright and dim. Infants (and some animals) also demonstrate another perceptual system that individuates small sets of objects (subitizing, Trick & Pylyshyn, 1994). Case proposed that systems like these served as the building blocks of the concept of number, such that a child can understand that 5 is a greater magnitude than 3, and that 5 refers to the count of objects in a set. Case hypothesized that some children have not coordinated these distinct systems into a central conceptual structure of number. Although each system was relatively robust in isolation, the lack of integration prevented children from having a solid concept of number, which in turn led to mathematical difficulties. For example, Griffin, Case, and Siegler (1994) found children who could add  $2 + 4$  but did not know that the answer 6 was larger than the addends. These children were at risk for school failure, because they had not developed an integrated concept of number.

Based on his hypothesis, Case and colleagues created a curriculum, called RightStart where students played a number of board games designed to help coordinate different systems that feed into the adult concept of number. For example, the children would count a number of steps on the game board and compare the total number of steps with another value, which helped them relate the ordinal (fifth step) and cardinal (five

total) meanings of number. The at-risk children who played these games for several months ended up equivalent to their peers who were not at risk. Interestingly, matched children who completed more traditional remediation did not show strong gains in their core concept of number. By hypothesis, this is because they did not coordinate different aspects of number. For example, children in traditional math programs simply counted but did not compare magnitudes of the counts in the same activity. More recent brain-based research is further exploring the hypothesis that the development of the concept of number depends on coordinating different perceptual systems (e.g., Butterworth, 1999).

In each of the three preceding examples, innovations in thought resulted from the coordination of lower-level systems. We view these as examples of dynamic transfer, because they unified competencies originally developed for another purpose or historical context. Some might object to viewing these as examples of transfer, because transfer often refers to the transfer of knowledge (as opposed to systems). However, this difference is simply one of disciplinary perspective. In Case's formulation, he did not describe the component abilities that feed into the concept of number as perceptual systems. Rather, he described them as ordinal and cardinal knowledge about quantity. Dynamic transfer involves the coordination of different sources of functional behavior, whether they are mental, social, and physical systems or different types of knowledge states and abilities.

#### The Significance of Interactive Experiences for Innovation and Dynamic Transfer

Dynamic transfer depends on the environment to support coordination. We highlight four characteristics of environments and how they can support innovation. We emphasize characteristics of the environment to rebalance the transfer literature, which

has tended to emphasize internal processes of transfer. Nevertheless, it is important to remember throughout this discussion that we cannot fully ascribe the processes of dynamic transfer to the structure of the environment. People need to engage the environment in ways that lead to innovation. If people have performance goals that stress immediate efficiency, they can block useful exploration into the larger structure of the situation and minimize learning (Vollmeyer, Burns, & Holyoak, 1996). In contrast, a playful attitude allows people to explore new ways of interaction that can yield innovations. Hatano and Inagaki (1986) argued that if the risk attached to the performance of a procedure is minimal, people are more inclined to experiment and adapt new ways of doing things, noting that “when a procedural skill is performed primarily to obtain rewards, people are reluctant to risk varying the skills, since they believe safety lies in relying on the ‘conventional’ version” (p. 269). Adaptive experts may recuse themselves from high-pressure situations, so they can work out new ways of doing things at lower risk. Examples include superstar athletes who continue to change their repertoire in the off-season or professors who take sabbaticals to learn new skills. More generally, the structure of the environment is a necessary component for dynamic transfer, but people still need to transfer in attitudes, values, and other remnants of prior experience that determine whether they will engage (or change) the environment in productive ways. We say more about this when we turn to the inertia problem.

### Distributed memory

The first reason the environment of interaction is valuable is that people have limited working memories, and it is hard to coordinate a jumble of concepts without distributing some of the work to the environment. This is not to say that people cannot

do important intellectual work without interacting in the environment. Self-explanation, for example, is a powerful way for people to check for inconsistencies in their understanding (Chi, de Leeuw, Chiu, & LaVancher, 1994). Thought experiments, such as Maxwell's demon, Schrödinger's cat, and Searle's Chinese room, are an example of this. People also have powerful imaginations that enable them to reconfigure spatial relations into novel structures (Finke, 1990). Kekule famously proposed the structure of the benzene ring after imagining a snake seizing its own tale. Even so, it is difficult to use purely internal processes to coordinate new ideas during dynamic transfer. An important property of the external world is that it can store intermediate structures. Architects, for example, sketch designs to help the process of conceptualizing a new design (Goel, 1995). We suspect that one of the important conceptual changes in meta-cognition involves learning when and how to use the environment to amplify thinking and memory in ways that lead to innovation.

An example of how alleviating memory demands supports discovery comes from research by Zhang and Norman (1994). These authors had people solve variants of the Tower of Hanoi puzzle, which involves moving a stack of disks, one-by-one, from one peg to another. Zhang and Norman created variants of the Tower of Hanoi task that eased the burden of remembering the rules of the task. For example, one rule is that a larger disk cannot be placed on a smaller one. By replacing the disks of different sizes with cups of different sizes, people did not have to remember the abstract verbal rule, because the environmental constraints (cups fitting into each other) enforced the rule. Participants who received the cup version of the task were more effective at innovating an efficient algorithm for solving the Tower of Hanoi problem.

Distributed memory is a particularly powerful support for innovation in young children. Piaget (1976) conducted the first developmental research with the Tower of Hanoi task. He found that 5- and 6-year-old children could not solve simple problems involving as few as two disks. Klahr and Robinson (1981) reasoned that this was due to the challenge of remembering the task rules. They constructed a version of the Tower of Hanoi quite similar to Zhang and Norman's (1994), for example, replacing disks with cups so that the children did not have to remember the ordering rules. As a result, 6-year-old children were able to innovate solutions for problems with as many as three cups requiring as many as five moves.

External changes that leave a trail also support backtracking. Similarity transfer, as the act of an individual mind, often has an all-or-nothing quality. If people make a far transfer that does not succeed, they can learn why the analogy did not work. But, they need to start over looking for a new analogy. In contrast, dynamic transfer is an iterative process: a sequence of near transfers that accumulate into innovations. When the process of dynamic transfer hits a dead end, people can backtrack, unrolling near transfers by looking at a record in the environment, until a choice point is reached and another possible trajectory revealed and pursued. For example, the Math children who were learning the balance scale often made faulty predictions, even late in the learning session. When a prediction failed, they changed their justifications. They did not start over by considering only weight or distance. They maintained their use of two dimensions of information, and changed the mathematical operators.

#### Alternative interpretations and feedback

The second reason that interactive experiences are valuable is that they can help people let go of pre-existing ideas that can block innovation. One obstacle to innovation is that prior interpretations can transfer in and interfere with developing new interpretations. Bruner and Potter (1964), for example, asked people to view pictures of common objects (e.g., a dog, silverware). The pictures were initially blurry and the experimenters slowly brought them into focus. The researchers measured how much focus the participants required before they could name the object of the picture. The researchers found that a high level of initial blurriness and a slow rate of focusing interfered with people's abilities to identify the pictures. Under these conditions, participants needed more focus before they could name the picture. The authors' explanation was that people formed interpretations of the initial pictures and had a hard time letting go of them. The blurrier the initial photo, the worse their initial interpretations. Moreover, the slower the focusing, the more entrenched participants became in their original interpretation. This study is often cited by philosophers of science as evidence that existing hypotheses can blind scientists to novel interpretations of the data (e.g., Greenwald, Pratkanis, Leippe, & Baumgardner, 1986). Here, it illustrates how initial interpretations can interfere with the development of alternative ways of seeing.

---

Insert Figure 4 about here.

---

Interactions with the environment generate feedback and variability that can help people shake free of their initial interpretations. Chambers and Reisberg (1985) studied how prior interpretations can interfere with developing new ones. The researchers asked people to look at a picture that can be seen as a duck or a rabbit (e.g., Fig. 4). Once

people had an interpretation, they closed their eyes. Chambers and Reisberg asked the people if they could come up with a second interpretation. Could they overcome their original interpretation (e.g., duck) and see a second interpretation (e.g., rabbit)? Not a single participant over several studies could do the re-interpretation. They were stuck with their first interpretation. However, if people are allowed to open their eyes and jiggle the picture a bit, they can find the alternative interpretation.

A second example of the value of interaction comes from a study by Martin and Schwartz (2005). They asked 9-year-olds to solve equivalent fraction problems, for example, “What is one-fourth of eight?” Children received eight small plastic tiles and had to indicate their answers. When children could only look at the pieces without touching, they tended to indicate one and/or four pieces as the answer. They mapped the pieces in one-to-one fashion to the numerals “1” and “4” of the fraction  $1/4$ . They were interpreting the problem with their prior concept of natural number. However, when children could push the pieces around, they were nearly four times more successful. By moving the pieces, the children began to notice grouping structures. They discovered new mathematical interpretations of the tiles, including that it is possible to make “4” groups where each (“1”) group has two pieces. The children were on their way to innovating a new understanding of rational number.

#### Candidate structure

The third reason interactivity is highly important for innovation is that the environment constrains and structures possible actions. This point was made forcibly by Simon (1969), who introduced the metaphor of an ant walking on a beach to a food source. To the external observer, the ant appears to be engaged in a sophisticated

representational process of navigation while keeping in mind the goal of obtaining food. Simon, however, proposed that the ant had no such representation. Instead, the ant had a few simple rules of behavior that led to complex behaviors due to the complexity of the beach.

Unlike ants, people can learn from interacting with complex, well-structured environments. This is one of the core ideas behind scaffolding. A scaffold permits people to engage in the structure of a mature performance. Through this engagement, people appropriate or internalize the structure of the environment. Ideally, they learn the structure well enough that the scaffolds can be removed and people do not become dependent on them. If the structures of the environment are new to the learner, they can yield petite innovations. The benefit of structuring environments is not limited to physical environments. As noted above, children used symbolic mathematics to generate candidate structures for predicting the behavior of a balance scale. Without the candidate structures of mathematics, the children could not coordinate the two dimensions of weight and distance.

New structures can also be imported into people's environments, and people can learn to innovate new ways of interacting with those artifacts. Lin (2001), for example, describes the case of teacher in Hong Kong who decided to use a form of instruction imported from the United States. The teacher used an extended mathematics problem-solving task based on a 20-minute video narrative called, *The Adventures of Jasper Woodbury* (CTGV, 1997). To solve the problem in the video, students often have to work for several days, and they work in groups. This manner of instruction was foreign to the Hong Kong teacher, who had not used group learning and who had always taught

modular units that ended with each class. As the students worked on the problem, the teacher confronted a number of problems: students turned out to be competitive in groups; the extended nature of the problem disrupted her classroom routines; and the teacher was concerned about losing authority as the students worked on the problem without her direct guidance. The structure of the novel instructional materials were disruptive, but over time, the teacher found innovative solutions that built on the materials' structure and that helped her to restructure the classroom and how she thought about the importance of teaching social skills.

#### A focal point for coordination

Finally, interactivity provides a natural focal point for coordinating different sources of knowledge. One powerful example of coordination taking place in the environment comes from the study of so-called “split-brain” patients who have had the neural connections between the right and left hemispheres of their brains severed as a treatment for severe epilepsy. Although this procedure creates two highly isolated cognitive systems within a single individual, some patients who have had this procedure can function quite normally in everyday life. One reason for their success may be their ability to use the external world to coordinate thinking in their isolated hemispheres. Kingstone and Gazzaniga (1995) report that one split-brain patient almost fooled them into thinking there was covert communication between his two hemispheres. When they presented the words “bow” and “arrow” to opposite hemispheres (via opposite visual fields), the patient was able to produce a coordinated drawing of a bow and arrow. Although this seemed like neural communication between the hemispheres, it was actually occurring through the drawing surface itself, as each hemisphere had access to

the emerging drawing on the page. The patient substituted the external world for neural pathways as the focal point of coordination.

Although less dramatic, neurologically normal people can find themselves in a similar situation as the split-brain patients. Innovations depend on integrating knowledge, but people often lack the structures necessary to relate the knowledge. With insufficient resources to make progress “inside the head,” coordination of mental systems often must occur outside the brain. Consequently, external interactions are central to this coordination. Sometimes, connections made in the external world can be internalized so that that environmental supports are no longer needed. We can return to the example of children’s early number learning reported by Case and colleagues (Griffin, Case, & Siegler, 1994). The simple board games were designed so that children had to coordinate ordinal and cardinal conceptions of number. For example, imagine a game where children roll a die that shows a “6”. The children have to count out six spaces in order to reach a particular position. As they land on each space, they pick up a chip. When they reached the final space, they have to decide if they have more chips than another student who rolled a five and ended up on the fifth space. In this example, the environment is designed to help the child coordinate the cardinal value of the digit “6” (total amount) with its ordinal value (sixth step). Interactive focal points can bring different pockets of knowledge into alignment, and for the children learning math, it leads to an understanding of number where ordinal and cardinal properties are related to one another.

#### Dynamic Transfer as a Measure of Innovative Experiences

Innovation can present a problem for educational settings. Innovative processes are often slow and filled with trial and error. In the balance scale study, the children were

trying almost any type of mathematical operation to find an answer. Moreover, innovative processes often fail to yield an innovative outcome. Many of the younger children never learned to use multiplication to solve the balance scale problems. These perceived inefficiencies raise challenges for educational practice. For example, discovery curricula can look inefficient, which leads to objections like “Wouldn’t it just be more efficient to tell students the answer?” At the same time, schools should be an ideal setting for designed experiences in innovation, which learners can then transfer to life beyond school.

The relation of education to innovation hinges on the types of assessments used to measure the benefits of particular experiences. Sequestered Problem Solving (SPS) assessments of transfer are excellent for detecting the scope of a successful innovation. However, they are not ideal for evaluating innovative experiences that do not come to fruition with an innovative outcome. An alternative approach is to use Preparation for Future Learning (PFL) assessments (Bransford & Schwartz, 1999). PFL assessments measure dynamic transfer; students have an opportunity to learn during the assessment with the support of useful external resources. In a PFL assessment, using the environment to learn during a test does not contaminate the test results or mean the student has been cheating. Instead, the assessment asks whether people can learn with the help of the environment.

One example of a PFL assessment that detected the benefits of innovative experiences comes from classroom research on learning statistics (Schwartz & Martin, 2004). Over multiple class periods, 9<sup>th</sup>-grade students received numerous opportunities to innovate ways to graph and measure variability. Consider a task where students had to

innovate a way to measure variability. Students received the grids shown in Figure 5. Each black circle indicates where a pitching machine through a ball when aimed at the  $X$  in the center. The students' task was to find a method to compute a reliability index for each machine, so people might decide which one they would like to buy. The goal was not for the students to invent the canonical solutions for measuring variability (though that would have been a fine outcome). Instead, the goal was for the students to have the opportunity to innovate in a well-structured environment. Students were explicitly told they were inventing a solution so they would engage the environment with an innovative mindset.

---

Insert Figure 5 about here.

---

Notice that the task includes the four elements of environments useful for supporting innovation. Students worked in pairs and they could draw on the sheets to test out ideas, thus the task is designed to help them distribute their cognition. The pitching grids create “contrasting cases” to help students generate alternative interpretations. For example, the lower right-hand grid shows a number of tightly clustered pitches that are far from the target. This helps students distinguish variability and error. The task requires students to use the candidate structures of symbolic mathematics to help them innovate. So, rather than simply ask the students to rank the reliability of the four pitching machines, they have to use the structure of mathematics to help them coordinate the different aspects of variability (e.g., sample size, density, distance from the mean, etc.). The visual structure of the grids naturally creates a point of coordination for organizing discussions and tentative solutions. For example, one student kept his finger on the ball of one pitching grid, while talking about another.

For tasks like these, very few students innovate the conventional solution. This fact returns us to the question of whether there is any value in engaging innovative processes even though they may not yield an innovative outcome. This is where PFL measures can be helpful: they can measure the extent to which innovative experiences prepare students to appreciate an innovative (to them) solution when it appears. The second phase of this experiment tackled this question. The target innovation was to learn to make comparisons across unlike distributions by normalizing data, as one done does when comparing the performance of athletes from different eras, even though they may have used different equipment.<sup>2</sup>

On the last day of two weeks of instruction, students were divided into two instructional conditions. Students in the Invent-a-Measure condition received raw data and tried to innovate a way to compare specific individuals from different distributions (e.g., given two groups of students who took different exams, which student got the highest relative score). No students innovated a satisfactory solution. Students in the Tell-and-Copy condition received the same data, but they learned and practiced a graphical procedure for solving the problems. A few days later, all students completed a long paper-and-pencil test that covered two weeks of instruction. The last problem on the test involved comparing athletes from different eras, which requires finding and using normalized scores. The students had not compared athletes during the instruction. Moreover, the target problem only provided summary statistics, whereas the students in both conditions had worked with raw data. Thus, it did not share surface features with the instruction and it made a difficult transfer problem.

Two complete the experimental design, the two instructional conditions were further subdivided into two test conditions. Half of the students in each condition had a worked example embedded in the middle of their test (the PFL condition), and the other half did not (the SPS condition). The worked example held the keys for how to solve normalizing problems. Students in the PFL conditions had to copy the worked example as part of the test, which they did quite accurately. The question was whether they would learn from the worked example, and then apply this learning to solve the target problem at the end of the test. Figure 6 summarizes the design of the study and shows the results.

---

Insert Figure 6 about here.

---

As measured by the SPS assessment, the innovative experiences of the Invent-a-Measure condition did not yield any special benefits over the Tell-and-Copy condition. It was the PFL assessment that revealed the benefits of the innovative experiences: the students learned from the worked example, as indicated by their performance on the target transfer problem. In contrast, the Tell-and-Copy students who received the worked example showed little benefit. The results from this study indicate that innovative experiences can prepare students to learn, even if students never innovate the correct outcome themselves, and that PFL assessments can detect the fruits of the innovation process. This is useful, because it helps to overcome the worry that innovative experiences, while fun for students, may not have any particular value. It also shows that dynamic transfer measures, which include opportunities to learn at transfer, can be more sensitive to important types of prior learning than standard similarity-based transfer measures that do not permit learning during the assessment. In this case, the PFL

assessment showed that the innovation experiences prepared students for the conceptual change that it is possible to compare unlike entities.

Unlike much of life, education can design innovation experiences that maximize productive processes and minimize losses. Sears (2006), for example, compared college students learning the Chi-Square statistic. The students received two kinds of instruction: one group read how to solve a series of contrasting cases involving the Chi-Square statistic; the other tried to innovate a solution to handle the contrasting cases and then they were shown how to solve the problems. Only the instructional goal differed; the students received the exact same materials and the same amount of time on task. The only difference was whether the students learned how to compute Chi-Square before or after they received the series of contrasting cases. There was a second, orthogonal factor: half of the students in each group worked in pairs, and half worked individually. On a subsequent SPS measure of how well the students had learned Chi-Square, all the conditions performed quite well with no differences. Thus, the designed innovation experience did not hurt the students, so long as they ultimately received the correct procedure. A different pattern of results showed up on a PFL assessment that involved learning Cohen's Kappa, which is a variant of Chi-Square. The students who innovated around Chi-Square made the dynamic transfer to Cohen's Kappa more frequently. Notably, the group work enhanced the effects of innovation experiences on the PFL assessment relative to working alone, but group work made no difference for the students who had simply solved (but not innovated solutions to) the Chi-Square problems. This makes sense. Groups working on standard, efficiency-driven tasks tend to partition and "hand off" results to one another for checking. In contrast, well-functioning groups

engaged in innovation are more likely to negotiate one another's ideas and learn from one another. Said another way, their discourse environment offers distributed memory, alternate interpretations, multiple candidate structures, and a focal point for coordinating their efforts – characteristics that we have argued support dynamic transfer for innovation.

### THE INERTIA PROBLEM

A natural threat to innovation is the tendency to assimilate situations into one's pre-existing routines. People may treat new situations like old ones, instead of changing their ideas or actions. The great strength of transfer is that people can reuse what they know to help them handle a new situation. The great risk of transfer is that people may have sufficient prior knowledge that they can operate efficiently enough in the situation. There is minimal impetus to overcome inertia and innovate a new way of doing or thinking about things. In the prior example of people learning to pour between jugs, the complex method was good enough to get the job done, and people did not seek out a potentially more effective solution. Instead of transferring for innovation, people often transfer for repetition. In many situations, repetition is a good thing, and constant innovation would be insufferable.

There is a second source of the inertia problem. Even if people do recognize a situation might be worth a bit of innovation, there is no guarantee that the effort to innovate will yield an innovative outcome. Moreover, the process of innovation typically requires a "productivity dip" as people become temporarily less efficient. Fullan (2001) discusses the "implementation dip" that occurs when people deploy an innovation, for example, in a business setting. The productivity dip refers to the loss of effectiveness while making the innovation itself.

An example of a productivity dip comes from the U-shaped curve in the acquisition of past tense (Ervin & Miller, 1963). Early on, children correctly use regular and irregular past tense forms, for example, “talked” and “bought” and “gave”. At this stage, children have memorized the past tense for each word on a case-by-case basis. In the next stage, children make a conceptual innovation by switching from their ad hoc system to a rule that might be summarized as, *add ‘-ed’ to a verb to make it past tense*. At this stage, children are able to convert novel words into the past tense, but they also add “-ed” to irregular verbs that they previously said correctly. For example, they might say “buyed” or “gived”. They over-generalize the newly acquired rule. It is only after some time that they learn to apply the past tense rule to some words and to use irregular forms for other words.

How do people overcome the inertia of their prior learning so they engage in innovation? Since at least Piaget (1952), the supposition has been that efforts towards innovation are fault-driven. People detect an internal contradiction or an external impasse, and this causes them to search for a new way to resolve the problem. This account is consistent with the scientific practice of rectifying theories given falsifying data – a parallel that Piaget captured in his notion of “genetic epistemology”. A fault-driven account claims that at a minimum transfer supports innovation because people need to transfer in prior learning to recognize there is a fault (cf. Popper, 1968).

Fault-driven transfer for innovation, however, cannot be the complete story. In some cases, people may not experience any fault in performance, but they will still try to coordinate new information with their prior learning, and this will yield new conceptions. Vosniadou and Brewer (1992) showed that children integrate their experience of a flat

world with overheard verbal facts, for example, that the world is round. As a result, they innovate hybrid models of the earth (e.g., a flat earth inside a round bowl). This is true even though their existing concept of a flat world does not lead to performance problems in their everyday lives. In the preceding case of language learning, children change their concept of past tense, even though the change leads to more errors in the short term. They are undeterred by the productivity dip.

### Trajectories through a Space of Innovation

Similarity and dynamic transfer provide different processes by which people interact with their context to overcome the inertia problem, or not. We describe how the two processes can work together, and how educational experiences can help people transfer the idea of being innovative. First, however, we treat each mechanism separately. Figure 7 anchors the discussion of how transfer and context interact with respect to innovation.

---

Insert Figure 7 about here.

---

Figure 7 uses a spatial metaphor to create a space of innovation. The vectors represent a sample of the trajectories people can follow over time; they range from genuine innovation to mere improvements in efficiency to outright failure. Movement up the gradient represents the process of adaptation over time. The height of the gradient at any point captures the adaptiveness of the outcomes at that point in time. (Adaptiveness can be gauged according to different definitions – time, cost, happiness, cultural mandates. The basic abstraction stays the same.) Over time, some trajectories reach higher levels of adaptiveness that constitute innovations. Other trajectories achieve more efficient routines, but they do not yield the long-term adaptive benefits gained by

innovations. The “fault” and “productivity dip” zones indicate that people become locally less adaptive in those areas. The gradient representation is inspired by work in artificial intelligence. Higher peaks represent more innovative outcomes. However, a steeper rise is more difficult to climb, because it implies a faster rate of adaptation. The trajectories of Figure 7 do not cover the complete range of ways people traverse the space of innovation (different endeavors will have their own terrains), but they help to clarify different transfer solutions to the inertia problem.

In similarity transfer, people apply prior learning whole cloth to a new situation. The challenge for similarity transfer is whether the existing knowledge and the new situation are close enough in the innovation space that the similarity will be noticed and leveraged. Trajectory 1 represents a spontaneous, far transfer that yields an innovation. The figure illustrates why spontaneous far transfer is rare: it depends on an unprovoked, prospective leap far beyond one’s position in the space. There is no impetus to do so, unless people can prefigure the innovative outcome of the transfer. And if it fails, people have learned little about the terrain over which they jumped.

Trajectory 2 represents fault-driven similarity transfer. A sequence of near transfers advances up the slope of innovation, until there is a fault that blocks further incremental progress. For example, engineers might work on creating fruit pickers. They transfer techniques of making gentle claws in a series of near transfers that handle increasingly delicate fruit, for example, from melons to oranges to nectarines. However, a situation presents a fault that cannot be traversed by a near transfer. For example, it turns out that their various claws keep puncturing and crushing tomatoes. The engineers might go down the chasm and perseverate on the near transfer of making gentler claws, but the

continued loss of adaptiveness may be too great. In the terms of artificial intelligence, hill-climbing algorithms do not descend into valleys. Instead, the engineers might intentionally search for an innovative analogy. They might draw the analogy to work they had done on claws for lifting boxes. They remember they made the boxes thicker. Given this prior learning, they might jump to the innovation of changing the thickness of the tomato through genetic engineering instead of working on softer and softer claws.

Trajectory 3 shows people stopping at the sign of a major fault. Prior knowledge fails to adapt to the demands of a new situation. Rather than moving forward, people disengage from the problematic situation and the process of adaptation stops. This is not inherently bad; dead ends can signal ill-conceived plans. In some cases, however, people disengage when there might be other paths available given perseverance. The transfer literature would be enriched if it examined the types of dispositions and social supports that can help people avoid the tendency to transfer in the belief that efforts to persevere are fruitless (e.g., Blackwell, Trzesniewski, & Dweck, 2007). Barron (2004) studied interest development in technology and found that some children who exhibit perseverance in innovation have parental models that reveal the other side of fault situations. Their parents provide models of the valued outcomes that expertise in computer science can yield.

Trajectory 4 shows dynamic transfer. Dynamic transfer is the product of a sequence of interactions with a well-structured environment that may include tools, representations, other people, and so forth. People make a number of small dynamic transfers that eventually yield the stable coordination that constitutes an innovation, as in the case of the children learning the balance scale. The children slowly worked through

different applications of math until they ended up with a stable way to coordinate two dimensions of information.

Dynamic transfer cannot jump across the productivity dip (or a fault) with a far transfer. Dynamic transfer depends on coordination with the environment for learning; it cannot fly across the environment with an abstraction. Nevertheless, the ease of the near transfers can lead people into the productivity dip. In fault-driven learning, it is necessary to recognize the fault. This is not the case with dynamic transfer, which depends on coordination rather than fault detection. Therefore, the temporary loss of adaptiveness does not have to be registered for continued learning. People can continue forward without noticing (or caring) that there has been a loss of productivity. We saw an example of this earlier in the productivity dip that children traverse when acquiring the past tense. Interestingly, connectionist models of past-tense acquisition are consistent with Trajectory 4. They demonstrate that small, local, and quantitative changes of the kind licensed by dynamic transfer can lead to innovations – qualitative changes in how knowledge of inflectional forms is organized (e.g., Plunkett & Marchman, 1996).

Dynamic transfer provides a less arduous route to innovation than similarity transfer, because people do not have to make a leap of innovation and because they may not register the productivity dip. Of course, there are situations where the environment is not conducive to entering the productivity dip, and people will avoid it. Children learn the past-tense rule because they are in an environment that encourages on-going interaction and tolerates dips in efficiency. In contrast, Trajectory 5 represents a situation of high risk or pressure to perform in the short term. In this trajectory, people balk at the productivity dip, because any loss of effectiveness can be problematic. One option is to

stop adapting. Here, we show the second alternative, which is to take a shallow gradient that does not involve a short-term loss in effectiveness. This trajectory ultimately merges with the final trajectory in our inventory.

Trajectory 6 represents a series of near similarity transfers, where people generalize and fine-tune their abilities across situations. This leads to an improvement in efficiency, but not a genuine innovation. This shallow region of the space is traversed until a plateau is reached. People achieve a level of “routine expertise” that enables them to accomplish familiar tasks in familiar and efficient ways (Hatano & Inagaki, 1986).

### Transferring the Idea of Dynamic Transfer

Given the space of innovation, we can hypothesize how people learn to be innovative. Our proposal is simple and involves a combination of similarity and dynamic transfer. The similarity transfer involves recognizing that similar situations have been solved by engaging in dynamic transfer. People learn that it is worth taking the temporary productivity dip associated with dynamic transfer, and they learn techniques that can help them make the dynamic transfer. We first describe a study that shows that specific educational experiences correlate with dynamic transfer. We conclude with a final study on transferring the very idea of being innovative through dynamic transfer.

In a study designed to demonstrate that educational experiences can correlate with dynamic transfer, college students completed a diagnosis task (Martin & Schwartz, 2007). The students received twelve reference cases, each on a separate sheet of paper. Each sheet indicated the symptoms and diagnosis for one patient. The students had to use the reference cases to help diagnose a series of ten new patients. For each new patient, they ordered medical tests to reveal symptoms, and when ready, they made their diagnosis.

The primary question was whether the students would make representations of the reference cases to help with diagnosis, for example, creating tables or decision trees. Or, would they simply shuffle through the reference cases to diagnose each new patient? Stepping back to make a representational tool before diagnosing the patients would involve a short-term productivity dip compared to just diving into the diagnosis task. The study compared undergraduate and graduate students. The graduate students were drawn from disciplines that involved complex information management (e.g., computer science, engineering, biology), but none of them had completed diagnoses like these.

All of the graduate students made visual representations of the reference cases, and spent roughly 15 minutes before they tried to diagnose their first patient. All told, they achieved 94% accuracy. The graduate students made a dynamic transfer. They used their general knowledge of data management to fashion representations tailored to the specifics of diagnosis, and half of them changed their representations as they became more familiar with the task. In contrast, less than 20% of the undergraduates made any sort of representation to help solve the problems. They began diagnosing the new patients within two minutes of receiving the reference cases. Even so, they achieved 91% accuracy in their diagnoses by shuffling through the reference cases to help diagnose each new patient. The study also included a third condition. Undergraduates completed the same set of diagnoses, but they had to set aside the reference cases each time they tried to diagnose a new patient. This imposed a heavy memory burden. Under these conditions, 88% of the undergraduates took the time to make representations. Thus, the (otherwise similar) undergraduates who had continuous access to the cases knew how to make representations. They simply did not bother to do so. In contrast, the graduate

students found it worthwhile to make representations. Their experiences in school were associated with their tendency to engage in dynamic transfer and its associated productivity dip.

---

Insert Figure 8 about here.

---

Figure 8 provides a summary of the trajectories of the three conditions. The undergraduates who did not have a memory burden avoided the productivity dip of making a representation. They engaged the task directly and did well enough, though their approach would become relatively inefficient given more diagnoses. (Once they completed their representations, those students who made representations were faster for each individual diagnosis and ordered more informative tests.) In contrast, the undergraduates who had the memory burden encountered a fault condition. They could not remember all the information from the reference cases. They made a similarity-based transfer along the lines of, “in situations of high memory load, it is worth making an external representation.” This drove them into the productivity dip, where they dynamically transferred in their knowledge of representations and adapted it to the particular task. The graduate students also entered the productivity dip to make the dynamic transfer. They had sufficient prior experiences with data management that they did not need to experience a fault to engage in the dynamic transfer of creating and revising representations.

#### Transferring the Idea of Innovating

The graduate students suggest that it is possible for people to learn to engage in dynamic transfer, even when it is not necessary for immediate efficiency. For these graduate students, making a visual representation to help solve these problems hardly

seems like a process of innovation (e.g., see Novick & Morse, 2000). But, for people who have less experience, it would constitute a process of innovation. The next study demonstrates that educational experiences can lead to people to transfer the very idea of innovating through a dynamic transfer.

In this study, seventh-grade students received a pretest hidden in a regular class activity that involved problems where a visual representation of causal chains would be helpful (Schwartz, 1993). The problems were something like, “X can communicate the disease to Y. Q can get the disease from R. F gets infected by Y. If X has the disease, what else can get the disease?” On the pretest, none of the children tried to construct a visual representation to help solve the problems. Two weeks later, the students completed several cycles where they explicitly invented visual representations to solve complex problems, none of which involved causal chains. After each invention effort, they were shown the representations that experts might use to solve the problem (e.g., a matrix, a Cartesian graph). This way they had a chance to experience the innovative process of visualization and they had a chance to experience the outcome of a good visualization. Two weeks after the instructional intervention, the students received a posttest similar to the pretest; it involved causal chains and it was embedded in a regular class activity. The key finding was that half of the students tried to invent new visual representations to help solve the problems on the posttest, compared to none at the pretest. Even though they had not used visualizations for problems of this structure, they recognized that the problems were quite complex, and they made the similarity transfer that visualizations could help. They then engaged in the dynamic transfer of trying to innovate a visual representation that could help solve the problem. They demonstrated a

conceptual change in the sense that before the instruction they would engage a task as given, and after, they learned the concept that “trying to invent a visual representation can help manage information complexity.”

Together, the preceding two studies suggest that appropriate experiences can lead people to transfer the idea of innovating. It is notable that the studies involved innovations for a specific class of problems (complex information management) using a specific external representational tool (visualizations). The children did not demonstrate a general disposition to innovate, for example, by brainstorming. To our knowledge, they also did not suddenly become more creative or innovative at home.

It is an important open question whether innovation can be a general, meta-cognitive mindset that transfers to many situations, or whether it is more confined to a specific set of techniques that are more or less appropriate for specific situations. We are not sanguine about “creativity programs” that try to train business people or children to be creative in general. As the literature on transfer has repeatedly shown, people do not transfer universally applicable skills like logical reasoning (e.g., Nisbett, Fong, Lehman, & Cheng, 1987). Instead, the ability to transfer and innovate grows from experiences where people gain insight into important environmental structures and their possibilities for interaction.

## ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grants No. BCS-0214549 and SLC-0354453. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## REFERENCES

- Barron, B. (2004). Learning ecologies for technological fluency: Gender and experience differences. Journal of Educational Computing Research, 31, 1-36.
- Barnett, S.M., & Ceci, S.J. (2002). When and where do we apply what we learn? A taxonomy for far transfer. Psychological Bulletin, 128, 612-637.
- Bassok, M. and K. J. Holyoak (1989). Interdomain Transfer Between Isomorphic Topics in Algebra and Physics. Journal of Experimental Psychology, 15, 153-166.
- Blackwell, L. S., Trzesniewski, K. H., & Dweck, C. S. (2007). Implicit theories of intelligence predict achievement across an adolescent transition: A longitudinal study and an intervention. Child Development, 78, 246-263.
- Bransford, J. D. & Schwartz, D. L. (1999). Rethinking transfer: A simple proposal with multiple implications. In A. Iran-Nejad & P. D. Pearson (Eds.), Review of Research in Education, 24, 61-100. Washington, D.C.: American Educational Research Association.
- Bransford, J. D., Franks, J. J., Vye, N. J. & Sherwood, R. D. (1989). New approaches to instruction: Because wisdom can't be told. In S. Vosniadou & A. Ortony (Eds.), Similarity and analogical reasoning (pp. 470-497). NY: Cambridge University Press.
- Brown, A. L., & Kane, M. J. (1988). Preschool children can learn to transfer: Learning to learn and learning from example. Cognitive Psychology, 20, 493-523.
- Bruner, J. S., & Potter, M. C. (1964). Inference in visual recognition. Science, 144, 424-425.

- Butterworth, B. (1999). The mathematical brain. UK: MacMillan.
- Carey, S. (1985). Conceptual change in childhood. Cambridge, MA: MIT Press.
- Carraher, D.W. & Schliemann, A.D. (2002). The transfer dilemma. Journal of the Learning Sciences, 11, 1-24.
- Case, R., & Okamoto, Y. (1996). The role of central conceptual structures in the development of children's thought. Monographs of the Society for Research in Child Development, 61, 1-265.
- Chambers, D., & Reisberg, D. (1985). Can Mental Images Be Ambiguous? Journal of Experimental Psychology: Human Perception and Performance, 11, 317-328.
- Chi, M. T. H., de Leeuw, N., Chiu, M-H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. Cognitive Science, 18, 439-477.
- Chi, M. T. H., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.
- COSEPUP [Committee on Science, Engineering, and Public Policy]. (2006). Rising above the gathering storm: Energizing America for a brighter economic future. Washington DC: The National Academies Press.
- CTGV [Cognition and Technology Group at Vanderbilt]. (1997). The Jasper Project: Lessons in curriculum, instruction, assessment, and professional development. Mahwah, NJ: Lawrence Erlbaum Associates.
- Detterman, D. L., (1993). The case for the prosecution: Transfer as epiphenomenon. In D. K. Detterman & R. J. Sternberg (Eds.), Transfer on trial: Intelligence, cognition, and instruction (pp. 1-24). Norwood, NJ: Ablex.

- diSessa, A. A., & Wagner, J. F. (2005). What coordination has to say about transfer. In J. P. Mestre (Ed.) Transfer of learning from a modern multidisciplinary perspective (pp. 121-154). Greenwich, CT: Information Age Publishing.
- Dufresne, R., Mestre, J., Thaden-Koch, T., Gerace, W., Leonard, W. (2005). Knowledge representation and coordination in the transfer process. In J. P. Mestre (Ed.) Transfer of learning from a modern multidisciplinary perspective (pp. 155 - 216). Greenwich, CT: Information Age Publishing.
- Dunbar, K. (1997) How scientists think: On-line creativity and conceptual change in science. In T. B. Ward, S. M. Smith and J. Vaid (Eds.), Conceptual structures and processes: Emergence, discovery, and change (pp. 461-493). Washington, D.C., American Psychological Association Press.
- Ervin, S. M., & Miller, W. R. (1963). Language development. In H.W. Stevenson (Ed.), Child psychology: The sixty-second yearbook of the National Society for the Study of Education, Part 1 (pp. 108–143). Chicago: University of Chicago Press.
- Feller, J., Fitzgerald, B., Hissam, S., & Lakhani K. R., Eds. (2005). Perspectives on Free and Open Source Software. Cambridge: MIT Press.
- Finke, R. (1990). Creative Imagery: Discoveries and Inventions in Visualization. Hillsdale, NJ: Erlbaum.
- Fullan, M. (2001). Leading in a culture of change. San Francisco, CA: Jossey-Bass
- Gentner, D., Brem, S., Ferguson, R. W., Markman, A. B., Levidow, B. B., Wolff, P., & Forbus, K. D. (1997). Analogical reasoning and conceptual change: A case study of Johannes Kepler. Journal of the Learning Sciences, 6, 3-40.

- Gick, M. L., & Holyoak, K. J. (1983). Schema Induction and Analogical Transfer. Cognitive Psychology, *15*, 1-38.
- Goel, V. (1995). Sketches of thought. Cambridge, MA: MIT Press.
- Gopnik, A., & Meltzoff, A. N. (1997). Words, thoughts, and theories. Bradford MIT Press, Cambridge MA.
- Greeno, J. G., Smith, D. R., Moore, J. L. (1993). Transfer of situated learning. In D. K. Detterman & R. J. Sternberg (Eds.) Transfer on trial: Intelligence, cognition, and instruction (pp. 99-167). Norwood, NJ: Ablex.
- Greenwald, A. G., Pratkanis, A. R., Leippe, M. R., & Baumgardner, M. H. (1986). Under what conditions does theory obstruct research progress? Psychological Review, *93*, 216-229.
- Griffin, S. A., Case, R., & Siegler, R. S. (1994). Rightstart: Providing the central conceptual prerequisites for first formal learning of arithmetic to students at risk for school failure. In K. McGilly (Ed.), Classroom lessons: Integrating cognitive theory and classroom practice (pp. 25–49). Cambridge, MA: MIT Press.
- Hammer, D., Elby, A., Scherr, R. E., & Redish, E. F. (2005). Resources, framing, and transfer. In J. P. Mestre (Ed.) Transfer of learning from a modern multidisciplinary perspective (pp. 89-120). Greenwich, CT: Information Age Publishing.
- Hatano, G. & Inagaki, K. (1986). Two courses of expertise. In H. Stevenson, H. Azuma, & K. Hakuta (Eds.) Child development and education in Japan (pp. 262-272). NY: Freeman.

- Hutchins, E. (1995). Cognition in the wild. Cambridge, MA: MIT Press.
- Ishibashi, H., Hihara, S., Takahashi, M., Heike, T., Yokota, T., & Iriki, A. (2002) Tool-use learning selectively induces expression of brain-derived neurotrophic factor, its receptor trkB, and neurotrophin 3 in the intraparietal multisensory cortex of monkey. Cognitive Brain Research, 14, 3-9
- Johnson-Laird, P. N. (1989). Analogy and the exercise of creativity. In Vosniadou, S., and Ortony, A. (Eds.) Similarity and Analogical Reasoning (pp. 313-332). New York: Cambridge University Press.
- Keil, F. C. (1989) Concepts, Kinds, and Cognitive Development. MIT Press, Cambridge, MA.
- Kerslake, D. (1986). Fractions: Children's strategies and errors. Berkshire, England: NFER-Nelson.
- Kingstone, A., & Gazzaniga, M. S. (1995). Subcortical transfer of higher order information: More illusory than real? Neuropsychology, 9, 321-328.
- Klahr, D., & Robinson, M. (1981). Formal assessment of problem-solving and planning processes in preschool children. Cognitive Psychology, 13, 113-148.
- Lin, X.D. (2001). Reflective adaptation of a technology artifact: A case study of classroom change. Cognition & Instruction, 19, 395-440.
- Lobato, J. (2003). How design experiments can inform a rethinking of transfer and vice versa. Educational Researcher, 32, 17-20.
- Luchins, A. S., & Luchins, E. H. (1959). Rigidity of behavior: A variational approach to the effect of einstellung. Eugene, OR: University of Oregon Books.

- Martin, L. & Schwartz, D. L. (2007). Prospective adaptation in the use of external representations. Manuscript submitted for publication. Stanford University, CA.
- Martin, T., & Schwartz, D. L. (2005). Physically distributed learning: Adapting and reinterpreting physical environments in the development of the fraction concept. Cognitive Science, *29*, 587-625.
- Mestre, J.P. (1994, February). Cognitive aspects of learning and teaching science. In S.J. Fitzsimmons & L.C. Kerpelman (Eds.), Teacher Enhancement for Elementary and Secondary Science and Mathematics: Status, Issues and Problems (pp. 3-1 - 3-53). Washington, D.C.: National Science Foundation (NSF 94-80).
- Nisbett, R. E., Fong, G. T., Lehman, D. R., & Cheng, P. W. (1987). Teaching Reasoning. Science, *238*, 625-631.
- Novick, L. R., & Morse, D. L. (2000). Folding a fish, making a mushroom: The role of diagrams in executing assembly procedures. Memory & Cognition, *28*, 1242-1256.
- Okita, Y. S., & Schwartz, D. L. (2006) Young Children's Understanding of Animacy and Entertainment Robots, International Journal of Humanoid Robotics, *3*, 393-412.
- Papert, S. (1988). One AI or many. Daedalus, *17*, 1-13.
- Pea, R.D. (1993). Practices of distributed intelligence and designs for education. In G. Salomon (Ed.), Distributed cognitions (pp. 47-87). New York: Cambridge University Press.
- Piaget, J. (1952). The origins of intelligence in children (M. Cook., Trans.). New York: International Press.
- Piaget, J. (1976). The grasp of consciousness. Cambridge, MA: Harvard University Press.

- Plunkett, K., & Marchman, V. A. (1996). Learning from a connectionist model of the acquisition of the English past tense. Cognition, 61, 299-308.
- Popper, K. R. (1968). The logic of scientific discovery. New York: Harper and Row.
- Rebellow, N. S., Zollman, D. A., Allbaugh, A. R., Englehardt, P. V., Gray, K. E., Hrepic, Z., Itza-Ortiz, S. F. (2005). Dynamic transfer: A perspective from physics education research. In J. P. Mestre (Ed.) Transfer of learning from a modern multidisciplinary perspective (pp. 217- 250). Greenwich, CT: Information Age Publishing.
- Reed, S. K., Ernst, G. W., & Banerji, R. B. (1974). The role of analogy in transfer between similar problem states. Cognitive Psychology, 6, 436-450.
- Ross, B. (1989). Reminders in learning and instruction. In Vosniadou, S., and Ortony, A. (Eds.) Similarity and Analogical Reasoning (pp. 438-469). New York: Cambridge University Press.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. Nature, 323, 533-536.
- Schwartz, D. L. (1993). The construction and analogical transfer of symbolic visualizations. Journal of Research in Science Teaching, 30, 1309-1325.
- Schwartz, D. L. & Bransford, J. D. (1998). A time for telling. Cognition & Instruction, 16, 475-522,
- Schwartz, D. L., Bransford, J. D., & Sears, D. (2005). Efficiency and innovation in transfer. In J. P. Mestre (Ed.) Transfer of learning from a modern

- multidisciplinary perspective (pp. 1-52). Greenwich, CT: Information Age Publishing.
- Schwartz, D. L., & Martin, T. (2004). Inventing to prepare for learning: The hidden efficiency of original student production in statistics instruction. Cognition & Instruction, 22, 129-184.
- Schwartz, D. L., Martin, T., & Pfaffman, J. (2005). How mathematics propels the development of physical knowledge. Journal of Cognition and Development, 6, 65-88.
- Sears, D. A. (2006). Effects of innovation versus efficiency tasks on collaboration and learning. (Doctoral dissertation, Stanford University, 2006). Dissertation Abstracts International, 67, 1652A.
- Simon, H. A. (1969). The sciences of the artificial. Cambridge, MA: MIT Press.
- Suchman, L. (1987). Plans and situated actions: The problem of human machine communication. NY: Cambridge University Press.
- Thagard, P. (1996). The concept of disease: Structure and change. Communication and Cognition, 29, 445-478.
- Thelen, E. & Smith, L.B. (1994). A Dynamic Systems Approach to the Development of Cognition and Action. Cambridge, MA: The MIT Press.
- Thorndike, E. L., & Woodworth, R. S. (1901). The influence of improvement in one mental function upon the efficacy of other functions. Psychological Review, 8, 247-261.

- Trick, L. M., & Pylyshyn, Z. W (1994). "Why are small and large numbers enumerated differently? A limited-capacity preattentive stage in vision". Psychological Review 101, 80-102.
- Tyson, L. M., Venville, G. J., Harrision, A. L., & Treagust, D. F. (1997). A multidimensional framework for interpreting conceptual change events in the classroom. Science Education, 81, 387-404.
- Vollmeyer, R., Burns, B.D., & Holyoak, K.J. (1996). The impact of goal specificity on strategy use and the acquisition of problem structure. Science, 20, 75-100.
- Vosniadou, S. & Brewer, W.F. (1992). Mental models of the earth: A study of conceptual change in childhood. Cognitive Psychology, 24, 535-85.
- Vygotsky, L. S. (1987). The collected works of L. S. Vygotsky. (Eds). R. Rieber and A. Carton. NY: Plenum.
- Whitehead, A. N. (1929). The aims of education. New York: MacMillan.
- Zhang, J., & Norman, D. A. (1994). Representations in distributed cognitive tasks. Cognitive Science, 18, 87-122.

## FOOTNOTES

1. There has been some confusion about transfer, because researchers slip between its meaning as an outcome measure and as a mechanism. This is natural, because people draw inferences about the mechanisms of transfer by looking at measures gathered from transfer tasks. However, as many authors have proposed (e.g., Lobato, 2003), if people fail on a researcher-designed measure of transfer, it does not mean there were not mechanisms of transfer at play that led the people to give their answer. It just means the hypothesis that the experimental manipulation would promote transfer for a specific outcome failed.

2. One computes how far each athlete's performance exceeds the mean of their respective era, and then divides those values by the variability of the era. These measures indicate how far each athlete outstripped others at the time, and these measures can be compared.

## FIGURE CAPTIONS

Figure 1. Pie wedges and tiles often used to teach children fractions. The pies exhibit a part-whole structure. In contrast, to turn the three tiles into three-fourths, it is necessary to impose structure, for example, by adding a fourth tile to the side.

Figure 2. Two trajectories of transfer (adapted from Martin & Schwartz, 2005).

Figure 3. A balance scale problem. When the hand is removed, will the balance scale tilt left or right, or will it balance?

Figure 4. Duck/rabbit image used in many experiments (e.g. Chambers & Reisberg, 1985).

Figure 5. Pitching machine task (Schwartz & Martin, 2004).

Figure 6. The effects of innovative experiences on SPS and PFL measures of learning (adapted from Schwartz & Martin, 2004).

Figure 7. Trajectories of innovation (see text for explanation).

Figure 8. Three trajectories in a medical diagnosis task.

FIGURES

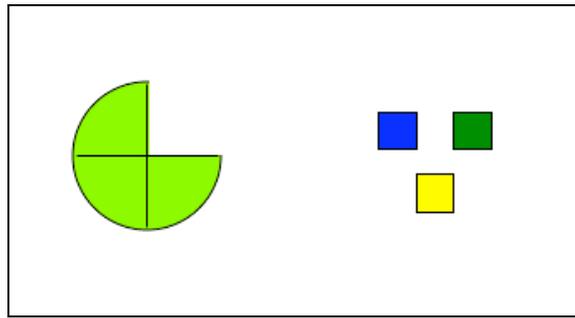


Figure 1

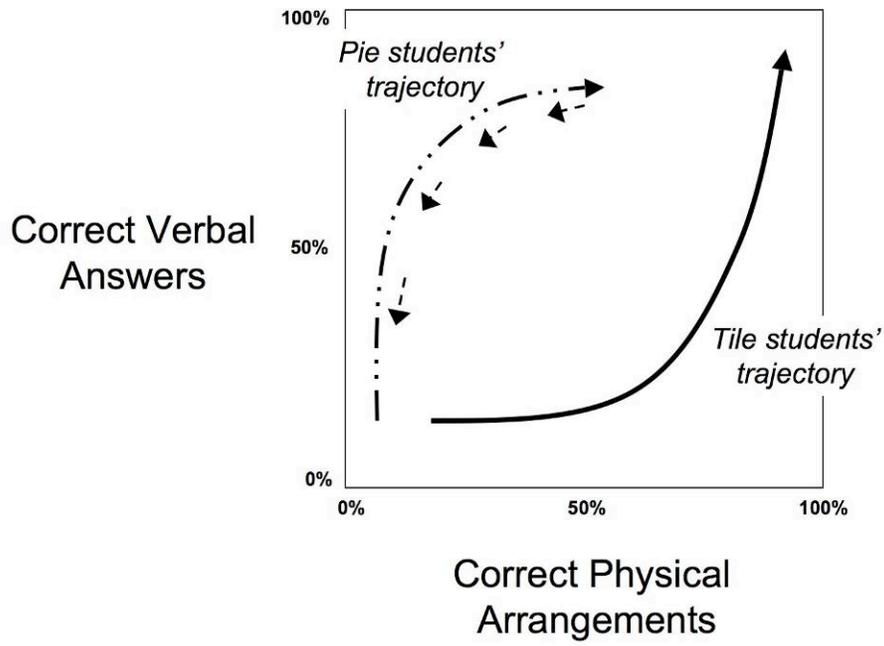


Figure 2

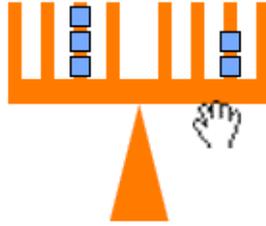


Figure 3

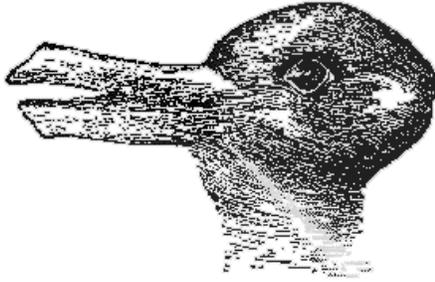


Figure 4

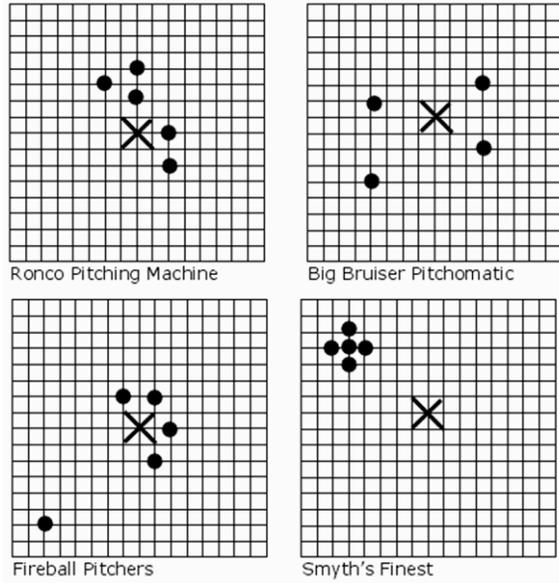


Figure 5

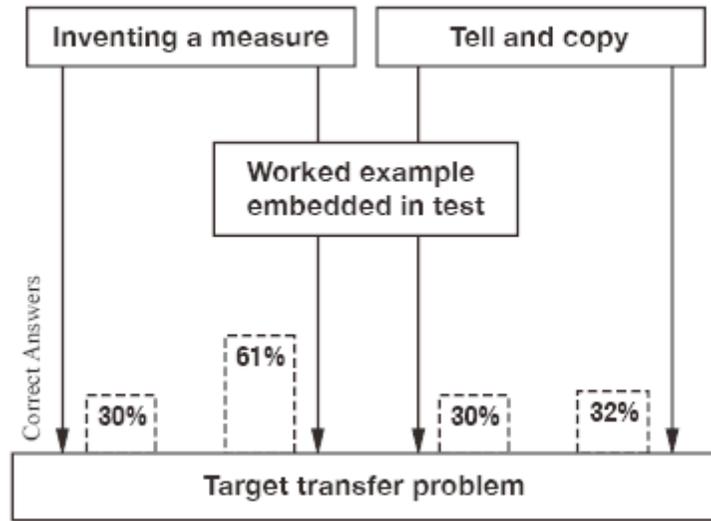


Figure 6

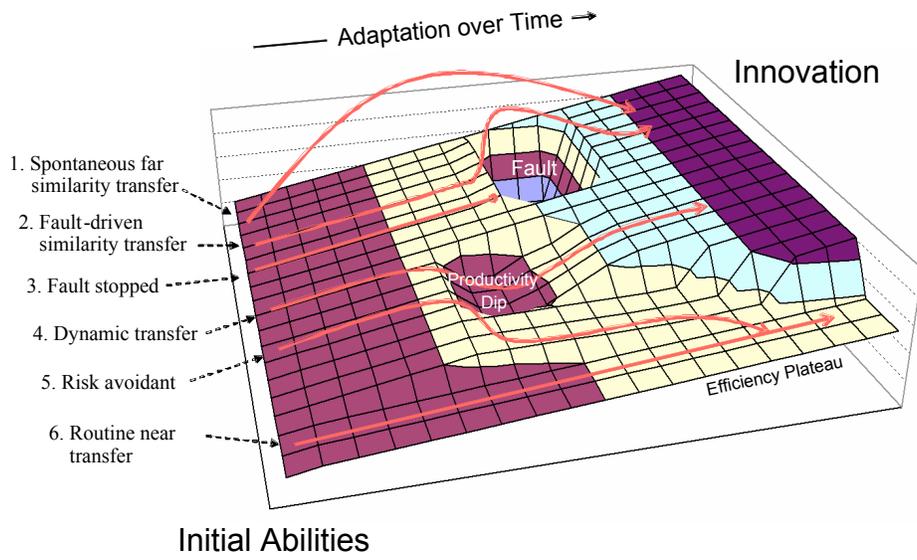


Figure 7

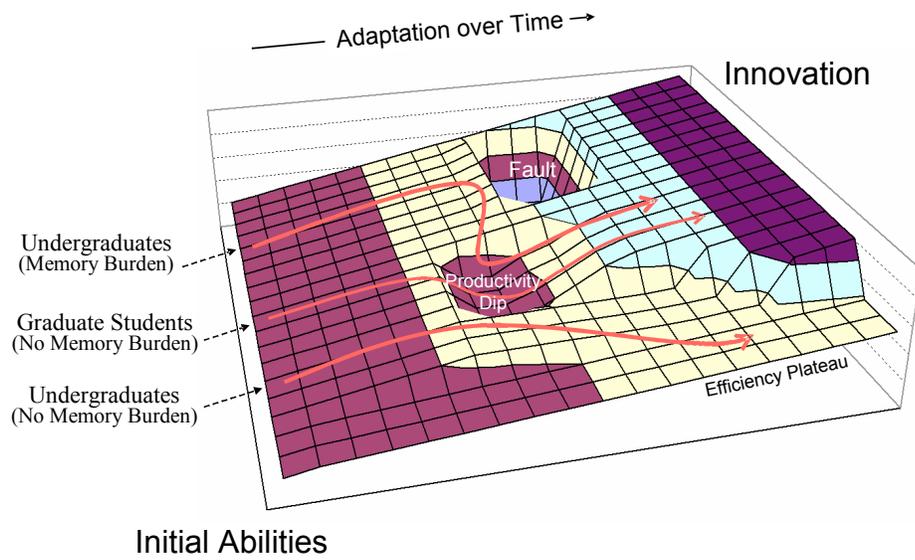


Figure 8