

In K. J. Holyoak & R. G. Morrison (Eds.) 2005. *Cambridge Handbook of Thinking and Reasoning*. New York: Cambridge University Press.

Complex Declarative Learning

Micheline T.H. Chi

Department of Psychology
Learning Research & Development Center
University of Pittsburgh

Stellan Ohlsson

Department of Psychology
University of Illinois at Chicago

Preparation of this paper was supported in part by Grant 200100305 from the Spencer Foundation, and in part by Grant 9720359 from the National Science Foundation (Center for Interdisciplinary Research on Constructive Learning Environments, CIRCLE, <http://www.pitt.edu/~circle>) to the first author, and Grant BCS-9907839 from the National Science Foundation to the second author.

INTRODUCTION

How do people acquire a complex body of knowledge, such as the history of the Panama Canal, the structure of the solar system, or the explanation for how the human circulatory system works? Complex learning takes longer than a few minutes and requires processes that are more complicated than the associative processes needed to memorize pairs of words. The materials that support complex learning – such as texts, illustrations, practice problems, instructor feedback -- presented in classrooms and elsewhere, are often difficult to understand, and might require extensive processing. For example, learning about the human circulatory system requires many component processes, such as integrating information from several sources, generating inferences, connecting new information with existing knowledge, retrieving appropriate analogies, producing explanations, coordinating different representations and perspectives, abandoning or rejecting prior concepts that are no longer useful, and so forth. Many of these component processes are still poorly understood, so that we have even less understanding of the complex process of *learning* a large body of knowledge.

Complex knowledge can be partitioned into two types: declarative knowledge and procedural knowledge. Declarative knowledge has traditionally been defined as knowledge of facts or knowing *that*; whereas procedural knowledge is knowing *how* (Anderson, 1976; Winograd, 1975). Declarative knowledge is descriptive and use-independent. It embodies concepts, principles, ideas, schemas, and theories (Ohlsson, 1994; 1996). Examples of declarative knowledge are the laws of the number system, Darwin's theory of evolution, and the history of the Panama Canal. The sum total of a person's declarative knowledge is his or her understanding of the way the world, or some part or aspect of the world, works, independently of the particular tasks the person undertakes. Procedural knowledge, such as how to operate and troubleshoot a machine, how to solve a physics problem, or how to use a computer text editor, is prescriptive and use-specific. It consists of associations between goals, situations, and actions. Research in cognitive neuroscience supports the reality of this distinction between declarative and procedural knowledge (Squire, 1987).

The acquisition of complex procedural knowledge has been extensively investigated in laboratory studies of skill acquisition, problem solving, and expertise (Ericsson, 1996; Feltovich, Ford & Hoffman, 1997), and in field studies of practitioners (Hutchins, 1995; Keller & Keller, 1996). Issues that have been explored include the role of perceptual organization in expert decision making, the breakdown of goals into sub-goals, the effect of ill-defined goals, the nature of search strategies, choices between competing strategies, the conditions of transfer of problem solving strategies from one problem context to another, the effect of alternative problem representations, the role of collaboration in complex tasks, and so on. As will become obvious in this chapter, the issues relevant to the study of complex procedural learning are different from those relevant to the study of complex declarative learning. Because the acquisition of procedural knowledge has been researched so extensively in the past few decades, there are several recent reviews (Holyoak, 1995; Lovett, 2002; Lovett & Anderson, this volume; VanLehn, 1989). Therefore, this chapter will focus primarily on the acquisition of a body of declarative knowledge.

The study of complex declarative learning is still in its infancy and has not yet produced a unified theory or paradigmatic framework. The organization of this chapter is meant to suggest one form that such a framework might take. In the first section, we describe basic characteristics of complex declarative knowledge. In the second section, we classify the different types of changes that occur in declarative knowledge as one learns. This classification is the main contribution of the chapter. The third section is a brief treatment of the so-called learning paradox (Bereiter, 1985). We end with a few concluding remarks.

BASIC CHARACTERISTICS OF DECLARATIVE KNOWLEDGE

Size of Knowledge Base

The most basic observation one can make about declarative knowledge is that human beings have a lot of it. There are no precise estimates of the amount of knowledge a person possesses but two attempts at an estimate seem well grounded. The first is an estimate of the size

of the mental lexicon. The average college educated adult knows between 40,000 – 60,000 words (Miller, 1996, pp. 136-138). The total number of words in the English language is larger than 100,000. Because concepts only constitute a subset of declarative knowledge, this represents a lower bound on the size of a person's declarative knowledge base. Second, Landauer (1986) had estimated how much information, measured in bits, people can remember from a lifetime of learning. His estimate is 2×10^9 bits by age 70. It is not straightforward to convert bits to concepts or pieces of knowledge, but even very fast computers use only 32 or 64 bits to encode one basic instruction. If we make the conservative assumption that it requires 1,000 bits to encode one piece of knowledge, Landauer's estimate implies that a person's declarative knowledge base eventually approximates a million pieces of knowledge.

These estimates apply to the size of the knowledge base as a whole. At the level of individual domains, estimates of the size of domain-specific knowledge bases tend to result in numbers that are comparable to estimates of the mental lexicon. For example, Simon and Gilmarin (1973) estimated the number of chess piece configurations – chunks or patterns – known by master players to be between 10,000 and 100,000. We do not know whether this is a coincidence or a symptom of some deeper regularity.

In short, even without a precise definition of what is to count as a unit of knowledge, the average person's declarative knowledge base must be measured in tens of thousands, more likely hundreds of thousands of units. How all this knowledge – the raw material for reasoning and thinking -- is acquired is clearly a non-trivial, but under-researched, question.

Organization

Knowledge does not grow as a set of isolated units but in some organized fashion. To capture the organization of the learners' declarative knowledge, cognitive scientists operate with

three distinct representational constructs: semantic networks, theories, and schemas (Markman, 1999).

The key claim behind *semantic networks* is that a person's declarative knowledge base can be thought of as a gigantic set of nodes (concepts) connected by links (relations). All knowledge is interrelated and cognitive processes, such as retrieval and inferencing, operate by traversing the links. Early computer simulations of long-term memory for declarative knowledge explored variants of this network concept (Abelson, 1973; Anderson & Bower, 1973; Norman & Rumelhart, 1975; Quillian, 1968; Rips & Medin, this volume; Schank, 1972).

Because the distance between two nodes in a semantic network is determined by the number of relations one must traverse to reach from one to the other, semantic networks implicitly claim that declarative knowledge is *grouped by domain*. We use the term "domain" to refer to both informal areas of knowledge such as home decorating, eating at a restaurant, and watching sports, and formal disciplines like botany, linguistics, and physics. Pieces of knowledge that belong to the same domain are similar in meaning and therefore cluster together functionally. Consistent with this notion, membership in the same domain tends to produce higher similarity ratings, stronger priming effects, and other quantitative behavioral consequences; descriptions of these well-known effects can be found in textbooks in cognitive psychology (e.g., Ashcraft, 2002; Reisberg, 2001).

The structure of any domain representation depends on the dominant relations of that domain. If the dominant relation is set inclusion, the representation is organized as a *hierarchy*. The standard taxonomies for animals and plants are prototypical examples. In contrast, relations like *cause-effect* and *before-after* produce chain-like structures. In general, the representations of domains are *locally structured* by their dominant relations.

The semantic network idea claims that all knowledge is interrelated, but it does not propose any single, overarching structure for the network as a whole. Concepts and assertions are

components of domains, but domains are not components of a yet higher level of organization. Domains relate to each other in a contingent rather than systematic way. Informal observations support this notion. We have one concept hierarchy for *tools* and another for *furniture*, but the node *lamp* appears in both. Home decorating is not a subset of cooking, nor vice versa, but the two share the *kitchen*. The concept of *tangled hierarchies* (Hofstadter, 1999) describes one aspect of local, unsystematic contact points between internally structured domains. These comments are somewhat speculative, because there is little cognitive research aimed at elucidating the structure of the declarative knowledge base as a whole.

Domains can also be represented as *theories*. Theories are “deep” representations (borrowing a term from social psychologists, see Rokeach, 1970) in the sense of having well-articulated *center-periphery* structures. That is, a theory is organized around a small set of core concepts or principles – big ideas – on which the rest of the elements in the domain are dependent. The core knowledge elements are typically fundamental and abstract, while the peripheral ones are based on, derived from, or instances of the core ones. The most pristine examples of center-periphery structures are the formal axiomatic systems of mathematics and logic, where a small set of chosen axioms provide a basis for the proofs of all other theorems in a particular formal theory; and natural science theories such as Newton's theory of mechanical motion, Darwin's theory of biological evolution, and the atomic theory of chemical reactions. These theories are obviously experts' representations and novices' representations of those same domains may or may not exhibit a similar structure, indicating that change in structure is one dimension of complex learning. For example, DiSessa (1988, 1993) has argued that novice knowledge of mechanical motion is not theory-like at all, but is better thought of as an irregular collection of fragments (see Smith, DiSessa & Roschelle, 1995, for a modified version of this view).

Other cognitive scientists, however, prefer to represent the novices' understandings of the natural world as *intuitive theories*, in deliberate analogy with the explicit and codified theories of scientists and mathematicians (Gopnik & Wellman 1994; Gopnik & Meltzoff, 1997; McCloskey, 1983; Wiser & Carey, 1983). By referring to someone's naïve representation as a theory, one implies specifically that the representation shares certain characteristics with explicit theories; most prominently that it has a center-periphery structure.¹

A well-developed center-periphery structure is often the hallmark of an expert's representation of a domain, and a comparison between novices' and experts' representations of the same domain often reveals differences in the "depth" of their representations. However, one can raise the question whether "depth" should also be construed as a characteristic of the domain itself. That is, are some domains intrinsically "deep" while others not, so that a center-periphery structure is not an appropriate representation for some domains? If so, we would expect neither experts nor novices to construct "deep" representations of those domains. For example, in informal everyday domains such as *home decorating* or *eating at a restaurant*, the center-periphery structure is certainly less salient. (However, even if an everyday domain such as *Entertaining* might not have a principled theory, its sub-domain of *formal table setting* does; Bykofsky & Fargis, 1995, pp. 144-146; Tuckerman & Dunnan, 1995, pp. 176-177.) Moreover, even for informal domains such as *cooking* that we as novices might claim to lack deep principles, many professional chefs would disagree. Thus, to what extent is the pervasive striving for a center-periphery structure with increasing expertise a law of mental representation, and to what extent it is an adaptation to the objective structure of domains, remains an open question.

The network concept codifies the intuition that everything is related to everything else, and the theory concept codifies the intuition that some knowledge elements are more important than others. The concept of a *schema*, on the other hand, codifies the intuition that much of our declarative knowledge represents recurring patterns in experience. Although the term "schema"

have never been formally defined, the key strands in this construct are nevertheless clear. To a first approximation, a schema is a set of relations among a set of slots or attributes, where the slots can be thought of as variables that can take values within a specified range (Bobrow & Collins, 1975; Brewer & Nakamura, 1984; Marshall, 1995; Minsky, 1975; Norman & Rumelhart, 1975; Thorndyke, 1984). Take the concept of "cousin" as an example. A cousin can be defined by a schema containing slots such as children, parents, and siblings, along with a collection of relations such as parent-of and sibling-of:

$$(\text{cousin-of } y \text{ } w) = \text{def } [(\text{parent-of } x \text{ } y)(\text{sibling-of } z \text{ } x)(\text{parent-of } z \text{ } w)]$$

To say that a person understands that Steve (slot y) and Bob (slot w) are cousins is to say that he or she knows that Steve (slot y) is the son of Carl (slot x), Carl is the brother of John (slot z), and John is the father of Bob (slot w). The slots are associated with ranges of appropriate values. Being a child, Steve must be younger than Carl, so slot y might have an age range of 1-50 years old, and slot x might have an age range of 21 to 85-years-old. Similarly, slot y can have the values of being either a male (a son) or a female (a daughter).

Schemas are bounded units of knowledge, and it is essential to their hypothesized function that they are retrieved or activated as units. That is, if one part of a schema (relation or slot) is activated, there is a high probability that the rest of the schema will be retrieved as well. Schemas are typically abstract, precisely because they represent recurring patterns in experience. Level of abstraction can vary (Ohlsson, 1993).

There are many variants of the schema idea in the cognitive literature. In the classic chess studies of deGroot (1965) and Chase and Simon (1973), chess experts were found to know by heart thousands of board patterns (each pattern consisting of a few chess pieces arranged in a meaningful configuration), and these familiar patterns altered their perception of the board so as to suggest promising moves. Similar findings regarding the power of perceptual patterns to influence high level cognition can be seen in physician's ability to read x-rays (Lesgold,

Rubinson, Feltovich, Glaser, Klopfer, & Wang, 1988) and fire fighter's ability to seize up a fire (Klein, 1998). Similarly, there is evidence to show that experts' programming knowledge includes frame-like structures called *plans* (Soloway & Erhlich, 1984), which are stereotypical situations that occur frequently in programming: looping, accumulating values, and so forth. These basic plans not only serve as the building blocks when writing programs, but they are also necessary for comprehension of programs. *Scripts* are higher-order knowledge structures that represent people's knowledge of informal or everyday events, such as eating in a restaurant or visiting the dentist's office (Schank & Abelson, 1977). *Explanation patterns* are schemas for how to construct explanations of particular types (Kitcher, 1993; Ohlsson, 2002; Ohlsson & Hemmerich, 1999; Schank, 1986). Yet other schema-like constructs have been proposed (e.g., Collins & Ferguson, 1993; Keegan, 1989; Machamer & Woody, 1992). Chunks, explanation patterns, frames, plans, and scripts are variants of the basic idea that much declarative knowledge consists of representations of recurring patterns. For simplicity, we will use the term *schema* throughout this chapter to refer to all of these constructs.

Although the three constructs of networks, theories, and schemas appear side by side in the cognitive literature, the relations between them are unclear. First, it is not clear how a schema should be understood within the larger notion of a semantic network. For a schema to be a distinct representational entity, there has to be a well-defined boundary between the schema and the rest of the knowledge network. (If not, activation will spread evenly across the nodes and links in the schema and the nodes and links that are not in the schema, which contradicts the central claim of schema theory that the probability of spreading from one node within the schema to another node within the schema is higher than spreading to a node outside the schema.) However, the concept of a network does not provide any obvious way to explain what would constitute such a boundary, other than to assume that links among nodes within a schema are more strongly connected than links among nodes between schemas (Chi & Ceci, 1987;

Rumelhart, Smolensky, McClelland, & Hinton, 1986). The differentiation in the strength of linkages can create clusters that can be conceived of as schemas (Chi & Koeske, 1983).

The relations between a schema and a theory are equally unclear. One can conceptualize a schema as a tool for organizing information, but it is not obvious whether a schema makes assertions or claims about the world. In this conception, schemas are not theories, but people obviously have theories. Finally, any explication of the relation between networks and theories must specify how the center-periphery structure that is intrinsic to theories can be embedded within networks.

In this chapter, we take the stance that networks, theories and schemas, are three partially overlapping but distinct theoretical constructs. Different aspects of the organization of declarative knowledge are best understood with the help of one or the other of these constructs, or with some mixture of the three.

In summary, declarative knowledge bases are very large and they exhibit complex organization. The notion of semantic networks captures the fact that every part of a person's knowledge is related, directly or indirectly, to every other part. Representations of particular domains vary in "depth", that is, the extent to which they are characterized by a central set of fundamental ideas or principles to which other, more peripheral knowledge units are related. Declarative knowledge also represents recurring patterns in experience with schemas, small packets of abstract structural information that are retrieved as units and used to organize information. These three types of organization cannot easily be reduced to each other, and explanations of change in complex knowledge draw upon one or the other of these constructs or on some mixture of the three.

TYPES OF CHANGES

The purpose of this section is to describe different types of changes in the knowledge base as one learns a body of declarative knowledge. There exists no widely accepted taxonomy of changes in a body of declarative knowledge. We chose to characterize changes as potentially occurring in seven dimensions. Presumably, different cognitive mechanisms are responsible for changes along different dimensions, but the field has not specified with any precision learning mechanisms for every dimension. In each section below, we specify a dimension of change, summarize some relevant empirical evidence, and describe the cognitive processes and mechanisms, if any, that have been proposed to explain change along that dimension.

Larger Size

Cumulative growth in size is a basic dimension of change in a body of declarative knowledge. Adults obviously know more about the world in general than do children (Chi, 1976); so that children are often referred to as universal novices (Brown & DeLoache, 1978). Similarly, experts obviously know more about their domains of expertise than novices (Chi, Glaser and Farr, 1988). People routinely accumulate additional facts about the world from sources such as news programs, texts, pictures, and conversations. These sources present people with some factual information that they did not know before, and some of those facts are retained. The declarative knowledge base continues to grow in size throughout the life span, albeit perhaps at a slower rate as a person ages (Rosenzweig, 2001). Rumelhart and Norman (1978) have referred to this type of cumulative addition of pieces of knowledge as *accretion*.

For adults, cumulative acquisition of individual pieces of knowledge – facts -- must be pervasive and account for a large proportion of all learning. There is little mystery as to the processes of acquisition. People acquire them via perception and observation, via comprehension of oral and written discourse, and via inductive (Sloman & Lagnado, this volume) and deductive (Evans, this volume) reasoning (i.e., by inferring new facts from prior

knowledge, or integrating new facts with old knowledge and making further inferences from the combination).

A particularly interesting property of accretion is that it is self-strengthening. Many psychology studies have confirmed that what is encoded, comprehended, and inferred depends on the individual learner's prior knowledge. For example, Spilich, Vesonder, Chiesi and Voss (1979) presented a passage describing a fictitious baseball game. Not only was the amount of recall of the individuals with high prior baseball knowledge greater (suggesting that the information was properly encoded), but the pattern of recall also differed. The high knowledge individuals recalled more information directly related to the goal structure of the game (Spilich et al, 1979) as well as the actions of the game and the related changes in the game states (Voss, Vesonder & Spilich, 1980), whereas the low knowledge individuals recalled the teams, the weather, and other less important events, and confused the order of the actions. Moreover, high knowledge individuals were better than low knowledge individuals at integrating a sequence of sentences (Chiesi, Spilich & Voss, 1979, Exp. V). In short, prior knowledge leads to more effective accretion, which in turn generates more prior knowledge.

Although encoding, comprehending, and inference processes augment the knowledge base, they do not necessarily cause deep changes in prior knowledge. Consider once again a baseball fan reading a newspaper article about a game. He or she will acquire facts that are obviously new – the score in the 8th inning cannot have been known before the game has been played – but the facts about past games are not altered, and he or she is unlikely to acquire a new and different conception of the game itself, although additional facts about baseball games per se may be acquired. The key characteristic that makes this an instance of accretion is that the learner already has a schema for a baseball game, which presumably has slots for the basic actions (throwing the ball), the highest level goal (winning the game), as well as other aspects of the game (Soloway, 1978). Once that schema has been acquired, to become increasingly

knowledgeable is largely to acquire more knowledge that fits into those slots, as well as knowledge of sub-goals and relations between the basic actions and the goal (Means & Voss, 1985). Similarly, readers of narratives might acquire facts about some fictional events, but they are unlikely to change their conceptions of causality, time, or human motivation, arguably three central schemas in comprehending narratives (Graesser, Singer, & Trabasso, 1994; Kintsch, 1998).

These observations imply that we need to distinguish between two levels of learning. Comprehension as normally understood results in the construction of a specific instance of a schema or the accretion of schema-relevant facts. New information is *assimilated* to existing schemas. This is the basic mechanism of accretion. The size of the relevant declarative knowledge base increases without fundamental changes in structure.

Deeper learning, on the other hand, results in some structural modification of the learner's prior schema. The same distinction can easily be expressed within the other two theoretical frameworks that we use in this chapter. In network terms, accretion adds nodes and links without deleting or altering any prior ones, while deeper learning requires a reorganization of the network. In terms of intuitive theories, cumulative growth might develop the relations between the core principles and peripheral knowledge items, while deeper learning either develops the core principles, replaces or alters one or more of the core principles. We discuss deeper learning processes later in this chapter.

Denser Connectedness

In network terms, *connectedness* can be defined as the density of relations between the knowledge elements. We would expect the density of connections in a representation to increase as the learner acquires more knowledge. This implication was supported by a study in which we compared the node-link representation of a single child's knowledge of 20 familiar dinosaurs

with his representation of 20 less familiar dinosaurs (Chi and Koeske, 1983, see Figures 1 and 2). The nodes and relations of the network were captured from the child's generation protocols of dinosaurs and their attributes. The representation of the 20 more familiar dinosaurs was better connected into meaningful clusters in that it had more links relating the dinosaurs that belonged to the same family, as well as relating the dinosaurs with their attributes of diet and habitat. The representation of the 20 less familiar dinosaurs had fewer links within clusters, and thus the cluster were less densely connected, so that they appear less differentiated and more diffused. In short, the better learned materials were more densely connected in an organized way, even though overall, the two networks represented the same number of nodes and links.

-----Insert Figures 1 and 2 here-----

A special case of connectedness is the mapping between layers. Layers can be defined in different ways in different domains. For example, in the context of computer programming we can conceive of the specification (the goals) as the highest layer, and the implementation (the data structures and primitive actions of the program) as the lowest level. Designing and comprehending a program require building a bridge between the specification and the implementation (Brooks, 1983). This bridge maps the implementation to the specification through a series of layers. Expert programmers are skilled at linking high-level goals to specific segments of programming code, whereas less skilled programmers are more likely to link program goals to triggers like variable names (Pennington, 1987). Once again, we see that a person's knowledge base appears to become more densely connected with increased knowledge acquisition.

Another special case of connectedness is between the conditions (declarative knowledge) and the actions (procedural knowledge). For example, experienced and inexperienced pilots knew equivalent number of facts, but the inexperienced pilots failed to apply them in the context

of actions (Stokes, Kemper & Kite, 1997). One can interpret this to mean that the facts that the inexperienced pilots knew were not connected to their actions.

Although the cited studies involved highly domain-specific relations, there are many types of connections that play central roles in declarative knowledge bases. For example, causal relations play a central role in the comprehension of narratives (Buehner & Cheng, this volume; Trabasso & van den Broek, 1985) as well as scientific theories, and hierarchical relations such as set-subset relations form the backbone of taxonomic or classificatory knowledge structures (Rips & Medin, this volume). The general point is that as knowledge acquisition proceeds in a domain, the learner's representation of that domain will increase in connectedness in a meaningful way.

Increased Consistency

The *consistency* of a knowledge representation refers to the degree to which the multiple assertions embedded in an intuitive theory can, in fact, be true at the same time. A person who claims that the Earth is round but who refuses to sail on the ocean for fear of falling over the edge is inconsistent in this sense.

The concept of consistency has been explored for decades in many areas of psychology, philosophy, and education. Social psychologists investigated the consistency of belief systems in the 50's and 60's (Abelson, Aronson, McGuire, Newcomb, Rosenberg & Tannenbaum, 1968; Heider, 1944; Festinger, 1962/1957; Fishbein & Ajzen, 1975; McGuire, 1968), and it remains an area of active research (Eagly & Chaiken, 1993; Harmon-Jones & Mills, 1999). In the wake of Thomas Kuhn's influential book *The Structure of Scientific Revolutions* (Kuhn, 1970), the philosophical debate about theory change in science came to focus on how scientists react to inconsistencies (anomalies) between theory and data, and this perspective carried over into contemporary approaches to science education (Hewson & Hewson, 1984; Posner, Strike, Hewson & Gertzog, 1982; Strike & Posner, 1985). Education researchers were already primed

for this focus by the traditional concern in the Piagetian tradition with contradictions and inconsistencies as driving forces for cognitive development (Piaget, 1985). Unfortunately, the social, philosophical, educational, and developmental literatures on cognitive consistency are not as tightly integrated as they ought to be in the light of the nearly identical ideas that drive research in these fields.

It is reasonably certain that people prefer consistent over inconsistent beliefs, at least locally, and that the discovery of local inconsistency (or conflict, Ames & Murray, 1982) triggers cognitive processes that aim to restore consistency, just as Piaget, Festinger, Kuhn, and others have hypothesized. For example, Thagard (1989, 2000) has explored a computational network model called ECHO in which consistency is defined as the lack of contradictions between assertions and hypotheses. ECHO has successfully predicted human data from a variety of situations, including the evaluation of scientific theories in the light of data (Thagard, 1992a) and the outcome of court cases (Thagard, 1992b).

However, the relation between experienced inconsistency and cognitive change is complex. Several investigators have suggested that conflict triggers efforts to restore consistency only when the conflict is recognized by the learner himself or herself through reflection (Chi, 2000; Ohlsson, 1999; Strike & Posner, 1992). When learners are alerted to inconsistencies and conflicts by an external source, they are more likely to either assimilate or dismiss them (Chinn & Brewer, 1993). Contradiction highlighted by an external source is likely to trigger change processes only if the learner is dissatisfied with his or her current conception (Posner et al, 1982). Furthermore, there are many ways to respond to inconsistency (Chinn & Brewer, 1993; Darden, 1992; Kelman & Baron, 1968) and not all modes of response increase consistency (as opposed to bypassing the problem); we return to this topic in the section on the learning paradox.

Consistency should not be confused with veridicality. It is possible for a knowledge representation to be locally consistent and yet be inaccurate. For example, we have argued that

the naive conception of the circulatory system as a single loop system is flawed but nevertheless constrained by a consistent set of identifiable but inaccurate principles. The learner can use such a flawed conception systematically to generate incorrect explanations. (Chi, 2000). Historically, the Ptolemaic epicycle theory of the solar system was as internally consistent as the Keplerian theory, but obviously not as accurate.

Consistency should also not be confused with level of expertise. A more knowledgeable person does not necessarily have a more consistent domain representation than someone who knows less. Ability to operate with inconsistency has often been proposed as a sign of intellectual sophistication, while insistence on total consistency has long been associated with dogmatism and lack of intellectual flexibility (Ehrlich & Leed, 1969; Rokeach, 1960). A famous historical example is the resolution – or lack of resolution – within quantum mechanics between the wave and particle models of photons. These annoying entities insist on behaving as both waves and particles, and since the time of Niels Bohr physicists have been content to let them be that way.

Consistency is sometimes used synonymously with the term *coherence*, as in Thagard's (1992a) use of the term explanatory coherence to refer to the consistency between a hypothesis and evidence and other hypotheses. However, consistency is distinct from coherence in that, as a measure of a representation, coherence can be used to refer to the more well-defined connectedness in a semantic representation, in which the notion of contradiction or conflict is not an issue (Chi & Koeske, 1983). There is not enough evidence nor agreement about the concept of coherence to warrant discussing it as a separate dimension of change.

To summarize, increased consistency is an important type of change in a declarative knowledge base, but it is distinct from the concepts of higher veridicality, more advanced knowledge, and coherence.

Finer Grain of Representation

Reality is not simple, and almost any aspect of it can be described or represented at different levels of grain. As one learns more about something, one often comes to understand it at a finer grain. For example, learning how the human circulatory system works involves learning the components of the system, such as the heart, the lungs, blood, and blood vessels, and the relation that the contraction of the heart sends blood to different parts of the body.

Given this level of representation, one can then ask, *how does the heart contract?* To answer this question, one would have to learn about the constituents of the heart: the properties of contractive muscle fibers, the role of ventricle pressure, and so on. The learner might push yet towards another level by asking how individual muscle fibers contract. At each level the system is understood in terms of its constituent parts, and further knowledge acquisition expands each component into its constituent parts. This type of process expands the knowledge base, but in a particular way: It moves along *part-of* links (as opposed to *kind of* links). In network terms, what was formerly a single node is expanded downwards into an entire sub-tree.

Miyake (1986) collected protocol data that illustrated this type of change. She showed that dyads, in attempting to understand how a sewing machine works, would move to lower and lower levels when they recognized that they had not understood the mechanism. For example, in figuring out how a stitch is made, one can understand it by explaining that the needle pushes a loop of the upper thread through the material to the underside, so that the upper thread loops entirely around the lower thread. However, in order to understand how this looping mechanism works, one has to explain the mechanism at a yet finer level, namely in terms of how the bottom thread go through the loop of the upper thread.

Knowledge expansion via finer grain of representation is quite common in the sciences. The ultimate example is perhaps the reduction by chemists of material substances to molecules, which in turn are described in terms of atoms, which in turn are re-represented by physicists in

terms of elementary particles. We should keep in mind though that it is the experts' representations of these domains that are refined, and novices' representations do not necessarily follow suit.

In analyzing biological systems like the circulatory system and machines such as the sewing machine, the parts are objects of the *same kind* as the system itself so that they embody the *part-of* relations. In these examples, valves and veins are of the same kind and they are both parts of the cardiovascular system, and thread and a stitch are both of the same kind and they are both parts of the sewing process. The link between the behavior of the parts and the behavior of the whole can often be understood in terms of direct cause and effect, or in terms of mechanical constraints that force movement in one direction rather than another, such as the valves in the veins.

However, there are systems where the relation between the finer and coarser levels of analysis is not of the same kind and the behavior of the system is *emergent* (Chi, submitted; Wilensky & Resnick, 1999). A traffic jam is an example. A traffic jam is a gridlock of cars such that cars can no longer move at normal speed. But the cars are not of the same kind as the traffic jam. In this kind of system, the (often) observable macro level behavior (the traffic jam) can be represented independently of the micro level objects (the moving cars). Each individual car may be following the same simple rule, which is to accelerate if there is no car in front within a certain distance and to slow down when there is another car within that distance. But the jam itself can move backward even though the individual cars move forward. Thus, the behavior of the individual cars in a jam is independent of the jam. Nevertheless, the macro level pattern (the jam) arises from local interactions among the micro level individual cars.

Learning about systems of this kind does not necessarily proceed by unpacking parts into yet smaller parts, but might more often occur by acquiring the two representations of the system separately and then linking them. This type of learning process re-represents the macro in terms

of the relationship between the micro and the macro levels, in order to explain the macro level phenomenon (Chi, submitted; Chi & Hausmann, 2003).

It is not clear how often people are driven to expand their representations downward to finer grain of analyses. In everyday life, people do not always feel the necessity to connect phenomena at one level to phenomena at more fine-grained levels. For example, people appear content to understand the weather at the level of wind, temperature, clouds, humidity, rain, and snow, without re-representing them at the finer levels of molecular phenomena available to the professional meteorologist (Wilson & Keil, 2000). We do not yet understand the factors and processes that drive people to expand but the possibility of such expansion is one important dimension of change in declarative knowledge.

Greater Complexity

A distinct type of change in the knowledge structure is needed when the learner's current concepts are not sufficient to represent the phenomenon or system as a whole. The thing to be understood cannot be assimilated within any schema the learner has available. The learner can respond by creating a more complex schema (Halford, this volume). Although very little is known about how more complex schemas are developed, one plausible hypothesis is that they are created by combining or assembling several existing schemas (Ohlsson & Hemmerich, 1999; Ohlsson & Lehtinen, 1997).

The creation of the theory of evolution by natural selection is a case in point. In the 19th century, many biologists knew that there were variation within species and that many species produce more offspring than survive into adult (reproductive) age, and the fact (as opposed to the explanation) of inheritance was of course commonly accepted. The theory of evolution is the result of assembling or combining these three schemas in a very particular way into a new, more complex schema. The change process here does not move along either *kind-of* or *part-of*

relations and it does not refine the grain of representation. Instead, it moves to greater complexity. The resulting schema is more complex than either of the prerequisite schemas. Such a move does not necessarily require a higher level of abstraction (see next section). The prior principles of intra-species variation, inheritance, and differential survival were already abstract, and there is no significant increase in abstraction in the theory that combines them.

The assembly process can be prompted. In one study, Ohlsson and Regan (2001) studied a laboratory version of the problem of the structure of DNA. Based on published historical accounts of the discovery of DNA, we extracted eight different component concepts that had to be combined to represent the double-helix structure. These turned out to be concepts that most educated adults can be expected to possess, e.g., parallel, pairwise, inverse, complement, etc. We found a linear relationship between the proportion of these eight concepts that were primed by exercises prior to problem solving and the time it took undergraduate students to solve the laboratory version of the DNA problem.

The assembly process can be understood as a combination of schemas. The key step in combining schemas must be to align the slots of one schema to those of another. Natural selection schema does not work unless the species that exhibit variation is also the species that is subject to selective pressure. The assembly process might share features with *conceptual combination*, although the latter process refers to single lexical concepts consisting of unfamiliar noun-noun or adjective-noun pairs, such as *pet fish* (Costello & Keane, 2000; Hampton, 1997; Medin & Shoben, 1988; Rips & Medin, this volume; Smith, Osherson, Rips and Keane, 1988). We know little about the frequency and prevalence of moves towards creating greater complexity at either the single concept or schema levels, and less about the conditions that prompt people to engage in such moves.

The concept of abstraction, in terms of where it comes from or how it is derived, continues to be controversial after two millennia of scholarship. Besides the issue of how abstractions are formed, there is a second, frequently overlooked meaning of moving towards higher abstraction: Given a pre-existing set of abstractions, it is possible to re-represent an object or a domain at a higher level of abstraction. For example, Chi, Feltovich and Glaser (1981) showed that physicists represented routine physics problems in terms of the deep principles that would be needed to construct a solution, whereas physics novices (those who have taken one course in college with an A grade), tend to represent the same problems according to their concrete surface components such as pulleys and inclined planes. The point is that one and the same problem tends to be represented at these different levels of abstraction by two groups *both of whom know the relevant principles*. The novices in the Chi et al (1981) study knew the relevant principles in the sense that they could both state them and use them. However, they did not spontaneously represent problems in terms of those principles instead of concrete properties. Somewhere along the path to expertise, the physicists came to do so.

Re-representing at a higher level of abstraction (using already acquired abstractions) is a very interesting dimension of change, but relevant empirical studies are scarce. As is the case with most other types of changes, we lack knowledge of the conditions that prompt people to move along this dimension and the exact nature of the relevant cognitive mechanism.

Shifted Vantage Point

Changing the level of abstraction is closely related to, but different from, the process that we in normal parlance call *change of perspective*. A classic study by R. Anderson demonstrate that this phrase does not merely refer to a metaphor but to a concrete psychological process. Anderson and Pichert (1978) gave subjects a text to read that described a home. They instructed subjects to take the perspective of either a burglar or a prospective home buyer. The results

showed that the instructions led the subjects to remember different details, even when the perspective taking instructions were given *after* the subjects had read the text.

Shifting one's point of view can facilitate problem solving. For example, Hutchins and Levin (1981) used the occurrence of deictic verbs such as "come", "go", "take", "send", "bring", and place adverbs such as "here", "there", "across" in think-aloud protocols to determine the point of view of subjects solving the Missionaries and Cannibals problem. They found that problem solvers shift perspective as they solve the problem. Initially, they view the river that the Missionaries and Cannibals have to cross from the left bank. Later in the problem solving process, they view from the right bank. One of their most interesting findings was that when solvers were in an impasse in the sense that they have made two non-progressive moves out of their current problem solving state, they could resolve the impasse if they shifted their point of view. In short, the somewhat mysterious process of 'taking' a particular perspective should not be understood as purely metaphorical; this form of re-representation has real consequences for cognitive processing.

In these cases discussed, the perspective shift was transient. There is some evidence to suggest that children become more able to shift perspective as they grow older (Halford, this volume). For example, Shatz and Gelman (1973) showed that young 2-year-olds could not adjust their speech to the age of the listener, whereas 4-year-olds did adjust their speech depending on whether they were speaking to another peer or an adult. This suggests that older (but not younger) children are capable of shifting their perspectives to that of the listeners. Similarly, Piaget and Inhelder (1956) have shown that older but not younger children are capable of understanding what another viewer might see, when the other person views it from another perspective. Although one might assume that as children mature, they acquire more knowledge that enables them to shift perspective, this next study confirms this interpretation since it manipulates knowledge directly. We gave high school students opportunities to play with a

computer simulation that allows them to take different roles in a business context, such as being the vice president of a bank. Students were much more able to take the perspective of the client after playing with the simulation, whereas they were only able to take the perspective of the bank before playing with the simulation (Jeong, Taylor & Chi, 2000).

In another series of studies, we attempted to teach first-grade children about the shape of the Earth (Johnson, Moher, Ohlsson & Gillingham, 1999; Johnson, Moher, Ohlsson & Leigh, 2001; Ohlsson, Moher & Johnson, 2000). Deep understanding of this topic requires that a person can coordinate the normal – we call it *egocentered* -- perspective of a person walking around on the Earth with an *exocentered* perspective from a hypothetical (and physically unattainable) vantage point in space. Such perspective coordinations can be very complex. For example, consider sunsets. What in the egocentered perspective appears as the sun disappearing behind the horizon appears in the exocentered perspective as movement of the border between light and shadow across the surface of the Earth due to the latter's rotation. Clearly, the mapping between these two views of the event is far from natural, simple, or direct, and it requires considerable learning and instruction to develop the exocentered perspective and to link it to everyday perception.

These and related studies demonstrate the occurrence of shifting vantage points and document the advantages they bring. This type of change must be an important dimension of growth of declarative knowledge.

Discussion

We suggest that a complex body of declarative knowledge over time moves along multiple dimensions of change: size, connectedness, consistency, grain, complexity, abstraction, and vantage point. Undoubtedly there are other dimensions along which declarative knowledge

changes during learning as well, such as coherence, but each of these have at least some support in empirical studies.

Although we separate these seven dimensions analytically for purposes of this chapter, we do not suggest that a cognitive change typically moves along a single dimension. Most complex knowledge acquisition processes will involve simultaneous movement along more than one dimension. For example, learning about chemistry involves thinking of material substances as *solids*, *liquids*, and *gases* instead of, e.g. *iron*, *water*, and *air*; this is a move towards higher abstraction. At the same time, the chemistry student acquires a finer grained analysis of material substances in terms of atoms and molecules, and a large number of previously unknown isolated facts about such substances, e.g., their melting points. He or she might have to assemble a new schema such as *dynamic equilibrium*, which involves shifting vantage point between the atomic level (where there are continuous processes) and the emergent macro level (where there is, nevertheless, stability). A year of high-school chemistry is likely to require movement along all seven of these dimensions. We suggest that this is typical in the acquisition of complex declarative knowledge.

Given that a representation can change in all the ways that we have described above, research on the acquisition of complex declarative knowledge encounters a particular difficulty: How to assess the effects of different learning scenarios and training procedures. The study of declarative knowledge contrasts in this respect with the study of procedural knowledge. Learning of procedural knowledge such as problem solving can be assessed relatively straightforwardly by measuring the degree to which a learner's representation of the procedure approximates the correct solution procedure, in terms of the rules and strategies. Learning of declarative knowledge, on the other hand, must be measured in light of the seven dimensions mentioned above. This is perhaps the most important methodological problem in the study of complex declarative knowledge.

Although we understand the character of these seven dimensions relatively well, we know very little about what triggers people to move along one or the other dimension. What are the factors that trigger someone to move to a finer grain or to another level of abstraction; under which conditions will a learner move to an alternative vantage point? Similarly, we do not fully understand the nature of the processes that bring about the changes in each dimension. Empirical research has been focused on documenting the psychological reality of each type of change, and has not sufficiently pursued the question of triggering conditions and the processes of change.

The seven types of changes discussed so far expand the learner's prior knowledge base in a *monotonic* way in that the prior knowledge need not be rejected or overwritten. It is possible to move towards larger size, denser connectedness, finer grain of representation, greater complexity, higher abstraction, and a different vantage point without rejecting or replacing one's prior knowledge representation. The one exception is a move towards increased consistency. To achieve increased consistency, one might have to reject or abandon some prior knowledge or belief. The next section discusses such *non-monotonic* changes.

THE LEARNING PARADOX: MONOTONIC AND NON-MONOTNIC CHANGE

It is tempting to think of a novice as primarily lacking knowledge; the learning process is then naturally seen as a process of accretion: filling a void or adding information. Some of the other types of changes described in the previous sections, such as increased connectedness and moves towards finer grain of representation, also have this cumulative nature since they significantly extend prior knowledge. On the other hand, several of the other types of changes, such as greater complexity, higher level of abstraction, and shifting vantage point, do not have this cumulative nature. Rather, they go further in that they re-represent the domain rather than merely add to it. However, in either the cumulative cases or the re-representation cases, the changes do not require that prior knowledge be rejected or replaced. For example, re-

representing something at a higher level of abstraction does not require rejection of the prior representation, because abstract and concrete representations of the same thing are not mutually incompatible. We can switch back and forth between conceptualizing something as a *hammer* and as a *tool* without any need to make a permanent choice between these two concepts. Thus, in these types of re-representation process, the old and the new representation can co-exist. Likewise for re-representing two component concepts or schemas into a more complex concept or schema via assembly. The representations for the original concepts remain. In short, these types of cumulative and re-representational changes are monotonic.

However, there are learning scenarios in which (a) the learner has a well-developed intuitive theory of the target domain, and (b) the subject matter to be acquired directly contradicts one or more of the core principles or beliefs of that intuitive theory. Successful learning in scenarios with these properties requires that the learner goes beyond mutually compatible representations. The learner has to re-represent the domain in the more fundamental sense of abandoning or rejecting – i. e., stop believing -- what he or she believed before, and replacing it with something else. We refer to this as *non-monotonic* change.

Science education provides numerous examples of prior conceptions that must be abandoned. Research on so-called misconceptions have documented that people have complex and rich conceptions about domains in which they have not received explicit instruction, but for which everyday experience provides raw material for intuitive theory formation (Confrey, 1990). Research on such spontaneous science theories have focused on physics, chemistry, and biology, although social science and non-science domains have also been investigated (Limon, 2002). (The older social psychology work on belief systems focused primarily on intuitive theories of society and religion; see, e.g., Abelson et al, 1968; Rokeach, 1970.)

Mechanics (forces and motion) is by far the most investigated domain. The dominant misconception in this domain is that motion implies force (Clement, 1982; diSessa, 1983, 1988;

Halloun & Hestenes, 1985; McCloskey, 1983; Minstrel, 1982). Students assume that when an object is in motion, the motion is caused by a force being applied to the object, the object's motion is in the direction of the force, and an object will move with constant velocity as long as it is under the influence of a constant force, and the velocity of an object is proportional to the magnitude of the applied force. When there is no force, an object will either slow down, if it is moving, or remain at rest. Motion is thus misconceived as produced by force, as opposed to the more accurate view that motion is a natural (i.e., equilibrium) state that will continue indefinitely unless some force interferes with it. Students' intuitive theory is more like the impetus theory held by Jean Buridan and other 14th century thinkers (Robin & Ohlsson, 1989) than like the inertia principle that is central to the Newtonian theory. Misconceptions about other topics, such as biological evolution, are also well documented (Bishop & Anderson, 1990; Brumby, 1984; Demasters, Settlage & Good, 1995; Ferrari & Chi, 1998; Lawson & Thompson, 1988).

The empirical findings not only show that novices possess well-developed misconceptions about many domains (Reiner, Slotta, Chi & Resnick, 2000), but that these misconceptions persist in the face of instruction and other innovative kinds of intervention. For example, many science misconceptions in Newtonian mechanics are robust and remain after instruction, even at very selective academic institutions (DiSessa, 1982; Caramazza, McCloskey & Green, 1980). With respect to mechanics, innovative instructional interventions include the use of carefully chosen analogies (Clement, Brown & Zietsman, 1989; Driver, 1987), deliberately invoking cognitive conflict (Posner, et al, 1982), engaging in deliberate confrontation (Licht, 1987), or using a succession of increasingly sophisticated models (White & Frederiksen, 1990). Although it is difficult to evaluate the outcomes of such interventions, it appears that students at best acquire the scientific conception, perhaps in an encapsulated form, while maintaining their initial intuitive conception (Johsua & Dupin, 1987), which is not quite

the intended outcome. There are at least three reasons (presented below) why misconceptions are so resistant to instruction so that non-monotonic change often fails.

Distortion Via Assimilation

As was mentioned earlier, in learning, new information is typically assimilated to existing schemas. Thus one reason that misconceptions persist is that when an instructor states the more veridical theory so that it contradicts the learner's prior misconceived knowledge, the new information is typically distorted in the process of being assimilated to the prior misconceived knowledge. To illustrate, consider a young child who believes that the Earth is as flat as it looks to the unaided eye. What happens if he or she is told that the Earth is round? Nussbaum (1979; 1985), Nussbaum and Novak (1976), Vosniadou (1994a, 1994b), and Vosniadou and Brewer (1992) have observed two intuitive schemas that we are tempted to interpret as consequences of distortion by assimilation. Some children draw the Earth as a flat entity with a circular periphery (like a pancake); others claim that the Earth is spherical but hollow and half filled with dirt (thus providing a flat surface for people to walk on). In both cases, the Earth is both flat and round. Instruction to the effect that the Earth is round was thus assimilated to a prior flat Earth conception without any significant changes in the latter.

Evasion of Conflicts

Distortion via assimilation is most plausible when the learner is unaware of the conflict between his or her prior knowledge and new information. The previous example involving the shape of the Earth illustrates this well; the young child is not aware that he or she is interpreting the adjective "round" in a different way than that intended by the adult speaker. This type of distortion can be reliably triggered in the laboratory by deliberately creating texts that violate a normal reader's world view (Graesser, Kasselner, Kreuz & Mclain-Allen, 1998).

However, even if the conflict between prior knowledge and new information is detected, it does not necessarily trigger productive change processes. Social psychologists (Abelson et al, 1968) and cognitive researchers (Chinn & Brewer, 1993; Darden 1992) have converged on very similar lists of potential modes of response to inconsistency. They agree that inconsistency often triggers evasive maneuvers that dismiss the inconsistency in some other way than by revising the relevant knowledge. The most basic mode of response is *abeyance*, that is, to postpone dealing with a contradiction on the grounds that not enough information is available to decide what, if anything, follows. One step removed from doing nothing is *bolstering*: The person who encounters information that contradicts some concept or belief X hastens to seek out supporting or confirming evidence that supports X. Festinger (1962/1957) and others hypothesized that the need to reduce an inconsistency is proportional to the ratio of supporting to contradicting pieces of information. Thus, by drowning the contradicting piece of information in a flood of confirming ones, it is possible to lower the need to resolve the contradiction, and hence to keep going without altering one's knowledge. Another process with a similar outcome is *recalibration*, that is, to lower the importance one attaches to the conflicting thoughts, thus making the conflict itself less important and easier to ignore. (A student might decide that he or she is not interested in science after all, so it does not matter what they teach in science courses.) These processes constitute evasive modes of response to inconsistent information, but they are not *learning* processes because there is no constructive change in the person's knowledge.

Lack of Computational Power

In describing the seven dimensions of changes, we sometimes speculated on the processes of change. What would happen if the inconsistent information triggered one or more of the learning processes that we proposed in previous sections? Take the process of creating greater complexity via assembly as example. In that process, a more complex representation is

created by combining two or more existing representations. It is doubtful whether this process could lead to a new, more veridical theory. Each of the assembled representations will presumably be consistent with the learner's prior intuitive theory, so they will lack veridicality. One cannot combine two non-veridical representations to create a third, veridical representation. For example, learners' naïve conception of heat and temperature, when combined, do not add up to the correct scientific conception of heat (Wiser & Carey, 1983), nor can teleological and Lamarckian ideas combine to form the principle of natural selection.

Although we do not spell out each argument here, a similar case could be made regarding the processes responsible for each of the seven types of changes discussed in the previous section. None of them has the computational power to create a new conception that goes beyond its own conceptual inputs, since, by definition, they are non-monotonic changes.

To summarize, the mere presence of contradictory information is not sufficient to trigger productive cognitive change of the non-monotonic kind. A conflict between prior knowledge and new information might go undetected, in which case the learner might blithely assimilate the new information to prior knowledge, probably distorting it in the process. Even if the learner detects the conflict, he or she might hold the new information in abeyance rather than respond to it. If he or she feels a need to deal with the contradiction, there is a repertoire of evasive maneuvers, including bolstering and recalibration of subjective importance, that will make the contradiction less disturbing without any revisions in prior knowledge. Finally, the productive learning processes discussed previously do not have the computational power to create a new conception that goes beyond the conceptual inputs to those processes. The prevalence of these three kinds of responses to encounters with contradictory information—distortion via assimilation, evading conflicts, and lacking computational power--raises the question of how can an intuitive theory ever be replaced? That is, how can a truly new theory or idea, that is not an

extension of old theories or ideas, ever be acquired? Bereiter (1985) referred to this as the *learning paradox*.

CONCLUSIONS AND FUTURE DIRECTIONS

Despite the prevalence of distortion via assimilation to prior knowledge, evasion of conflicts, and lack of computational power, non-monotonic change does happen.

Children do replace their childhood conceptions with adult ones, some physics students do succeed in learning Newtonian mechanics, and scientists do sometimes replace even their most fundamental theories in the face of anomalous data. Thus there must be cognitive mechanisms and processes that can overcome the learning paradox. A theory of complex learning should explain both why non-monotonic change has such low probability of occurring, and how, by what processes, it happens when it does happen.

The study of such non-cumulative learning processes is as yet in its infancy. In this section, we offer a small number of speculative proposals about how non-monotonic learning processes can occur. These brief proposals are intended to serve as inspiration for further research.

Pathways to Non-Monotonic Change?

We describe below four mechanisms along with some empirical support. We then consider whether each of them can potentially achieve non-monotonic change.

Transformation via Bootstrapping

One hypothetical path to a new theory is to edit or revise one's existing theory piece by piece until the theory says something significantly different from what it said originally. We can conceptualize such a bootstrapping process as a series of *local repairs* of a knowledge structure.

Local repairs require simple mechanisms such as adding links, deleting links, reattaching links, and so forth. The critical condition for local repairs is that the student recognizes that the repairs are needed, by reflecting on the differences between his or her existing knowledge and new knowledge. We have some evidence that the accumulation of local repairs can lead to a significant transformation of a person's mental model of the circulatory system, from a flawed single-loop model to the correct double loop model (Chi, 2000).

As a second example of bootstrapping, Thagard (1992a) analyzed the changes in the French chemist Lavoiser's conception of matter during the critical years of the development of the oxygen theory of combustion. Thagard shows how Lavoiser's conception of combustion can be modeled by a semantic network, and how that network is gradually transformed over several years as the scientist is reflecting on the outcomes of empirical experiments. By adding and deleting nodes and re-drawing links, Thagard depicts Lavoisier's knowledge network as undergoing a gradual transformation such that its initial state represents the phlogiston theory of combustion but its final state represents the oxygen theory.

How much can transformation via local repairs explain? There are multiple explanations for why local repairs succeed in the case of the circulatory system. One reason is that the transformation from a single-loop type of model to a double-loop crosses no ontological categories (Chi & Roscoe, 2002). Another reason might be the relative lack of "depth" of this domain, in the sense that it cannot be represented by a center-periphery structure. The single-loop principle does not deductively imply the other relevant facts about the circulatory system in the manner in which Newton's three laws of motion imply more peripheral statements within the domain of motion. The looser connection between center and periphery might make the single-loop principle easy to tinker with. Finally, there is a question of commitment (Ohlsson, 1999). Although students believe that there is a single circulatory loop, this is not one of their most cherished beliefs and they probably do not experience it as important to their world view.

Tinkering even with the core principle of this domain might therefore come easier than in domains with a stronger center-periphery structure and deeper commitment to the core principles. Rokeach (1970) has presented evidence from other than scientific domains that knowledge elements are more resistant to change the more central they are. It is plausible that transformation via bootstrapping a sequence of local repairs is less applicable the “deeper” the domain, at least as long as the change has to encompass the core principles to be complete. So perhaps this bootstrapping process cannot be considered a true non-monotonic change mechanism.

Replacement

If stepwise revisions can only go so far to explain non-monotonic change, what alternative is there? Knowledge structures can be *replaced*. That is, an alternative representation of a domain is constructed in parallel with a prior one, through processes that do not use the prior one as input. The old and the new representations then compete for the control of discourse and behavior in the course of question answering, explanation, reasoning, and problem solving. The new, presumably more veridical representation frequently wins, and the old one eventually fades from disuse.

Bottom-Up Replacement

Replacement can proceed either bottom-up or top-down. First, consider a new representation built bottom-up. This might occur when the new knowledge is encountered in a context that does not necessarily evoke the conflicting prior knowledge. For example, students might experience science instruction as so distant from everyday experience that they build representations of what is taught in class that are independent from, and unconnected to the former. The outcome of such encapsulated knowledge is an ability to solve textbook problems, without enriched understanding of relevant phenomena encountered in other contexts (everyday experience, news reports, etc.). Due to the compartmentalization of contexts, the conflict

between the prior intuitive theory and the new theory is not salient to the learner, and the construction of the new theory can proceed without interference from prior knowledge.

If matters remain in this state, it is doubtful whether this can be considered successful non-monotonic learning. The crucial question is whether the new theory, once constructed, can migrate into and usurp the territory of the prior intuitive conception. Successful non-monotonic learning requires that a phenomenon previously understood within the intuitive theory begins to be understood within the new theory instead.

Top-Down Replacement

Consider the possibility of top-down generation of a new knowledge structure. An abstract schema might be acquired in an alternative domain and transferred wholesale to a new domain. An example of this hypothetical process is provided by recent attempts to understand the operation of the immune system in Darwinian terms. Philosophers and theoretical biologists have attempted to formalize Darwin's theory of evolution (Thompson, 1989), and the resulting abstract schema has been applied to the question of how the immune system could produce antibodies for a wide variety of antigens. The Darwinian answer is that the immune system continually generates more or less random antibodies; high fit between antibodies and antigens triggers increased production of the former; thus, the antigens themselves function as an environment that selects for the antibodies that fight them (Gazzinga, 1992). The accuracy of this theory of the immune system is not the issue here. It is an example of a process in which a complex abstract schema was transferred as a whole to provide a cognitive template for a novel theory of a physiological process far removed from the evolutionary processes of speciation and adaptation for which the schema was originally constructed.

This top-down process is limited in that it relies on the prior existence of an appropriate abstract schema, which raises the question of where abstractions come from. This issue has

remained controversial for over two millennia. The standard suggestions include induction over exemplars (Rips & Medin, this volume) and social interaction (Greenfield, this volume). Since the topic of abstraction is discussed elsewhere in this volume, we do not intend to answer this question here.

Side-stepping the issue of where an abstract schema comes from in the first place, we first need to know whether top-down replacement is possible, given that an abstract schema exists. To test the feasibility of this top-down replacement process, we are instructing students about a domain-general abstract schema that might serve as a template for understanding multiple concepts in many domain. One example is the schema of emergence (Chi, submitted), which has applications in biology, chemistry, and physics. It is plausible that direct instruction of this sort results in the de novo construction of an alternative conception, as opposed to gradual transformation of a prior conception.

Transfer via Analogy

Existence of an abstract schema may not be a necessary requisite for the top-down process to work. A concrete schema from another domain might serve as template, if the two domains are easy enough to align so that the transfer process can operate via analogy (Holyoak, this volume). In this hypothetical process, the learner acquires a schema in some source domain S; later, he or she is learning about some target domain T for which he or she already has an intuitive theory. The new information about T contradicts his or her prior intuitive theory about T, but is analogous to what is known about S. If the learner creates a new representation for T based on what is known about S instead of building directly on his or her current intuitive theory of T, then he or she might avoid distortion by assimilation.

We tested the reality of this *transfer of concrete schema* process in a virtual reality based scenario for teaching children that the Earth is round (Johnson, et al 1999; Johnson, et al 2001; Ohlsson, et al 2000). We created a virtual planet that was small enough so that the consequences

of sphericity were immediately perceivable. For example, even minor movement through the virtual world made objects visibly 'appear' or 'disappear' over the horizon. Having acquired a notion of living on a spherical planet in the context of this fictional asteroid (about which the children were not expected to have any distorting prior views), we then supported, via a one-on-one dialogue, the analogical transfer of that schema to the context of the Earth. Pre- to posttest comparisons between the treatment and a dialogue-only control group showed that the effect of prior learning in the virtual environment was positive (albeit small in magnitude). We infer that the schema for the virtual asteroid to some extent served as template for the new conception of the Earth that we tried to teach them. Hence, the learning paradox was overcome by stimulating the children to build a representation of what life on a sphere is like independent of their prior knowledge of the Earth, and then encouraging the use of that representation as a template for building a new representation of the Earth.

Ontological Shift

Ontological categories refer to a set of categories to which people partition the world in terms of its most fundamental features (as opposed to characteristic and defining features; Chi, 1997). For example, two high level categories that people are likely to partition the different types of entities in the world into are substances and processes. Each type of entity is conceptualized as having certain fundamental properties. For example, substances such as sand can be contained in a box, but processes such as a baseball game, cannot; on the other hand, processes can last for two hours but substances cannot. Misconceptions are mis-categorization of an entity into a wrong ontological category. For example, students typically misconceive of heat or electricity as a kind of stuff or substances that can move from one location to another (Chi, Slotta, & De Leeuw, 1994). Continued study of some entity that is initially thought of as belonging to category X might reveal properties that are not consistent with its ontological status.

In those cases, successful learning requires that the learner re-represents the entity as belonging to another ontological category, such as from a kind of substance to a kind of process (Slotta, Chi, & Joram, 1995).

This kind of ontological shift replaces a prior conception with a new conception in terms of an entity's ontological status. Thus, this process of ontological shift may qualify as a kind of a non-monotonic mechanism.

Toward A Theory of Learning

In 1965, Robert M. Gagne published a book, *The Conditions of Learning*, that summarized what was known about learning at the time. His approach was the unusual one of assuming that there are multiple distinct types of learning processes, distinguishable with respect to their prerequisites, processes, and results. He presented these in order of increasing complexity, beginning with "signal learning" (simple conditioning) and ending with "problem solving" (Gagne, 1965). The most noteworthy feature of his approach is signaled by the book's title: For each type of learning, Gagne asked under which conditions that type of learning might occur.

In our efforts to summarize what is known about the acquisition of complex declarative knowledge, we, too, have been led to present a list of different types of learning. In the realm of monotonic learning, we distinguish between seven different dimensions of change: size, connectedness, consistency, grain, complexity, abstraction, and vantage point. In the realm of non-monotonic change, we have specified numerous non-learning modes of response to contradictory information such as assimilation and evasive processes of abeyance, bolstering, recalibration, and why many of the learning mechanisms cannot in principle produce true non-monotonic learning. Finally, even our proposals with respect to non-monotonic learning breaks down into multiple processes like transformation via local repairs, bottom-up compartmentalized

replacement, and top-down replacement with the help of abstract schemas, transfer of concrete schema via analogies, and ontological shift. It seems likely that as the study of complex learning progresses, cognitive scientists will further our understanding of these replacement processes.

However, as Gagne clearly saw 40 years ago, a list of learning processes is by itself an incomplete theory of learning. One would expect such a theory to support explanation of learning outcomes, to allow us to say why one subject matter is more difficult to acquire than another, to predict the success rate of particular instructional scenarios, and so on. However, to accomplish these and other theoretical tasks, we need to know when, under which circumstances, one or the other learning process is likely to occur. A predictive science of complex learning requires that we can specify the *when* and *wherefore* of the many process hypotheses that spring from the imagination of the cognitive scientist. Nowhere is this more obvious than in the case of non-monotonic learning. This, we suggest, is the research front in the study of complex declarative learning.

REFERENCES

- Abelson, R. P. (1973). The structure of belief systems. In: R.C. Schank and K.M. Colby (Eds.), *Computer models of thought and language* (pp. 287-339). San Francisco, CA: W. H. Freeman.
- Abelson, R.P., Aronson, E., McGuire, W.J., Newcomb, T.M., Rosenberg, M.J., & Tannenbaum, P.H. (Eds.), (1968). *Theories of cognitive consistency: A sourcebook*. Chicago, IL: Rand McNally.
- Ames, G.J., & Murray, F.B. (1982). When two wrongs make a right: Promoting cognitive change by social conflict. *Developmental Psychology*, 18, 894-897.
- Anderson, J.R. (1976). *Language, memory, and thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J.R., & Bower, G.H. (1973). *Human associative memory*. New York: Wiley.
- Anderson, R.C., & Pichert, J. (1978). Recall of previously unrecallable information following a shift in perspective. *Journal of Verbal Learning and Verbal Behavior*, 17, 1-12.
- Ashcraft, M.H. (2002). *Cognition* (3rd ed.). Upper Saddle River, NJ: Prentice Hall.
- Bereiter, C. (1985). Toward a solution of the learning paradox. *Review of Educational Research*, 55(2), 201-226.
- Bishop, B., & Anderson, C. (1990). Student conceptions of natural selection and its role in evolution. *Journal of Research in Science Teaching*, 27, 415-427.
- Bobrow, D.G., & Collins, A. (1975). *Representation and understanding: Studies in cognitive science*. New York: Academic Press.
- Brewer, W.F., & Nakamura, G.V. (1984). The nature and functions of schemas. In: R. Wyer and T. Srull (Eds.), *Handbook of social cognition* (pp. 119-160). Hillsdale, NJ: Erlbaum.
- Brooks, R. (1983). Towards a theory of the comprehension of computer programs,

International Journal of Man-Machine Studies, 18, 543-554.

Brown, A.L. & DeLoache, J.S. (1978). Skills, plans and self-regulation. In: R.S. Siegler (Ed.), *Children's Thinking: What Develops?*, Hillsdale, NJ: Erlbaum.

Brumby, M. (1984). Misconceptions about the concept of natural selection by medical biology students. *Science Education*, 68, 493-503.

Buehner, M. & Cheng, P. (in press). Causal reasoning. In: Holyoak, K.J. and Morrison, B.G. (Eds.), *Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press.

Bykofsky, S., & Fargis, P. (1995) *The big book of life's instructions*. New York: Galahad.

Caramazza, A., McCloskey, M. and Green, B. (1980). Naive beliefs in "sophisticated" subjects: Misconceptions about trajectories of objects. *Cognition* 9:117-123.

Chase, W.G., & Simon, H.A. (1973). Perception in chess. *Cognitive Psychology*, 4, 55-81.

Chi, M.T.H. (1976). Short-term memory limitations in children: Capacity or processing deficits? *Memory and Cognition*, 4, 559-572.

Chi, M.T.H. (1992). Conceptual change within and across ontological categories: Examples from learning and discovery in science. In: R.N. Giere (Ed.), *Cognitive models of science*. Minneapolis, Minnesota: University of Minnesota Press.

Chi, M.T.H. (1997). Creativity: Shifting across ontological categories flexibly. In: T.B. Ward, S.M. Smith, & J. Vaid (Eds.), *Conceptual Structures and processes: Emergence, Discovery and Change*. (Pp. 209-234). Washington, D.C: American Psychological Association.

Chi, M.T.H. (2000). Self-explaining: The dual processes of generating inferences and repairing mental models. In: R. Glaser (Ed.), *Advances in Instructional Psychology*. Pg. 161-237. Mahwah, NJ: Lawrence Erlbaum Associates.

Chi, M.T.H. (Submitted). Emergent and commonsense direct processes: Potential

schemas for overcoming misunderstandings in science. *Journal of the Learning Sciences*.

Chi, M.T.H. & Ceci, S.J. (1987). Content knowledge: Its role, representation, and restructuring in memory development. In: H.W. Reese (Ed.), *Advances in Child Development and Behavior* (Vol. 20, Pp. 91-142). New York: Academic Press.

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

Chi, M.T.H., & Koeske, R.D. (1983). Network representation of a child's dinosaur knowledge. *Developmental Psychology*, 19: 29-39.

Chi, M.T.H., Glaser, R. & Farr, M. (1988). *The Nature of Expertise*, Hillsdale, NJ, Erlbaum.

Chi, M.T.H., Slotta, J.D. and de Leeuw, N. (1994). From things to processes: A theory of conceptual change for learning science concepts. *Learning and Instruction*, 4, 27-43.

Chi, M.T.H. & Hausmann, R.G.M. (2003). Do radical discoveries require ontological shifts? In: L.V. Shavinina (Ed.), *International Handbook on Innovation*, Oxford: Elsevier.

Chi, M.T.H., & Roscoe, R.D. (2002). The processes and challenges of conceptual change. In: M. Limon and L. Mason (Eds). *Reconsidering Conceptual Change: Issues in Theory and Practice*. Kluwer Academic Publishers, The Netherlands, pp 3-27.

Chiesi, H.L., Spilich, G.J., & Voss, J.F. (1979). Acquisition of domain-related Information in relation to high and low domain knowledge. *Journal of Verbal Learning and Verbal Behavior*, 18, 257-274.

Chinn, C.A., & Brewer, W.F. (1993). The role of anomalous data in knowledge acquisition: A theoretical framework and implications for science instruction. *Review of Educational Research*, 63, 1-49.

Clement, J. (1982). Students' preconceptions introductory mechanics. *American Journal*

of Physics, 50, 66-71.

Clement, J., Brown, D., & Zietsman, A. (1989). Not all preconceptions are misconceptions: Finding anchoring conceptions for grounding instruction on students' intuitions. *International Journal of Science Education*, 11, 554-565.

Collins, A., & Ferguson, W. (1993). Epistemic forms and epistemic games: Structures and strategies to guide inquire. *Educational Psychologist*, 28, 25-42.

Confrey, J. (1990). A review of the research on student conceptions in mathematics, science, and programming. In: C.B. Cazden (Ed.), *Review of Research in Education* (Vol. 1, pp. 3-56). Washington, D.C.: American Educational Research Association.

Costello, F.J., & Keane, M.T. (2000). Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*, 24, 299-349.

Darden, L. (1992). Strategies for anomaly resolution. In: R. Giere (Ed.), *Cognitive models of science*. Minneapolis, MN: University of Minnesota Press.

DeGroot, A.D. (1965). *Thought and Choice in Chess*. The Hague, the Netherlands: Mouton.

Demasters, S., Settlage, J., & Good, R. (1995). Students' conceptions of natural selection and its role in evolution: Cases of replication and comparison. *Journal of Research in Science Teaching*, 32, 535-550.

diSessa, A. (1988). Knowledge in pieces. In: G. Forman & P. Pufall (Eds.), *Constructivism in the computer age*. Hillsdale, NJ: Erlbaum.

diSessa, A. (1993). Toward an epistemology of physics. *Cognition and Instruction*, 10, 101-104.

diSessa, A. (1982). Unlearning Aristotelian physics: A study of knowledge-based learning. *Cognitive Science* 6:37-75.

diSessa, A. (1983). Phenomenology and the evolution of intuition. In: D. Gentner and A.

L. Stevens, (Eds), *Mental Models*. Hillsdale, NJ: Erlbaum, pp. 15-33.

Driver, R. (1987). Promoting conceptual change in classroom settings: The experience of the children's learning in science project. In: J.D. Novak (Ed.), *Proceeding of the Second International Seminar on Misconceptions and Educational Strategies in Science and Mathematics*, (Vol. II, pp. 97-107). Ithaca, NY: Department of Education, Cornell University.

Eagly, A.H., & Chaiken, S. (1993). *The psychology of attitudes*. Ft. Worth, TX: Harcourt, Brace & Jovanovich.

Ehrlich, H.J., & Leed, D. (1969). Dogmatism, learning, and resistance to change: A review and a new paradigm. *Psychological Bulletin*, 71, 249-260.

Ericsson, K.A. (1996). *The road to excellence*. Mahwah, NJ: Erlbaum.

Evans, J. (in press). Deductive Reasoning. In: Holyoak, K.J. and Morrison, B.G. (Eds.), *Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press.

Feltovich, P.J., Ford, K.M., & Hoffman, R.R. (1997). *Expertise in context*. Menlo Park, CA: AAI Press.

Ferrari, M., & Chi, M.T.H. (1998). The nature of naïve explanations of natural selection. *International Journal of Science Education*, 20, 1231-1256.

Festinger, L. (1962/1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.

Fishbein, M., & Ajzen, I. (1975). *Belief, attitude, intention and behavior*. Reading, MA: Addison-Wesley.

Gagne, R.M. (1965). *The conditions of learning*. New York: Holt, Rinehart and Winston.

Gazzinga, M.S. (1992). *Nature's mind*. New York: Basic Books.

Gick, M.L. & Holyoak, K.J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.

Gopnik, A., & Wellman, H. M. (1994). The theory theory. In: L.A. Hirschfeld and S.A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 255-293). Cambridge, UK: Cambridge University Press..

Gopnik, A., & Meltzoff, A.N. (1997). *Words, thoughts, and theories*. Cambridge MA: MIT Press.

Graesser, A., Kassar, M.A., & Kreuz, R.J., & Mclain-Allen, B. (1998). Verification of statements about story worlds that deviate from normal conceptions of time: What is true about Einstein's dreams. *Cognitive Psychology*, 35, 246-301.

Graesser, A., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, 101, 186-209.

Greenfield, P. (in press). Paradigms of cultural thought. In: Holyoak, K.J. and Morrison, B.G. (Eds.), *Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press.

Halford, G. (in press). Development of thinking. In: Holyoak, K.J. and Morrison, B.G. (Eds.), *Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press.

Halloun, I.A., & Hestenes, D. (1985). Common sense concepts about motion. *American Journal of Physics*, 53, 1056-1065.

Hampton, J.A. (1997). Emergent attributes in combined concepts. In: T. Ward, S.M. Smith, and J. Vaid (Eds.), *Creative Thought: an Investigation of Conceptual Structures and Processes*, American Psychological Association, Washington, D.C., pp. 83-110.

Harmon-Jones, E., & Mills, J. (1999). *Cognitive dissonance: Progress on a pivotal theory in social psychology*. Washington, DC: American Psychological Association.

Heider, F. (1944). Social perception and phenomenal causality. *Psychological Review*,

51, 358-374.

Hewson, P.W. & Hewson, M.G.A. (1984). The role of conceptual conflict in conceptual change and the design of science instruction. *Instructional Science*, 13, 1-13.

Hofstadter, D. (1999). *Godel, Escher, Bach - An Eternal Golden Braid*. Basic Books.

Holyoak, K.J. (1995). Problem solving. In: E.E. Smith & D.N. Osherson (Eds.), *Thinking: An Invitation to Cognitive Science, Volume 3*, pp. 267-296, Cambridge: MIT Press.

Holyoak, K.J. (in press). Analogy. In: Holyoak, K.J. and Morrison, B.G. (Eds.), *Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

Hutchins, E.L., & Levin, J.A. (1981). Point of view in problem solving. CHIP Tech. Rep. No. 105, University of California at San Diego.

Jeong, H., Taylor, R., & Chi, M.T.H (2000). Learning from a computer workplace simulation. *Proceedings of the 22nd annual meeting of the Cognitive Science Society*, pp. 705-710, Mahwah, NJ: Lawrence Erlbaum Associates.

Johnson, A., Moher, T., Ohlsson, S., & Gillingham, M. (1999). The Round Earth Project -- collaborative VR for conceptual learning. *IEEE Computer Graphics and Applications*, 19(6), 60-69.

Johnson, A., Moher, T., Ohlsson, S., Leigh, J. (2001). Exploring multiple representations in elementary school science education. In *Proceedings of IEEE VR 2001*, Mar 13-17, 2001, Yokohama, Japan, pp. 201-208.

Johsua, S. and Dupin J.J. (1987). Taking into account student conceptions in instructional strategy: An example in physics. *Cognition and Instruction*, 4(2), 117-135.

Keegan, R.T. (1989). How Charles Darwin became a psychologist. In: D.B. Wallace & H.E. Gruber (Eds.), *Creative people at work: Twelve cognitive case studies* (pp. 107-125).

Keller, C.M., & Keller, J.D. (1996). *Cognition and tool use: The blacksmith at work*.

Cambridge, UK: Cambridge University Press.

Kelman, H.C., & Baron, R.M. (1968). Determinants of modes of resolving inconsistency dilemmas: A functional analysis. In: R. P. Abelson, E. Aronson, W. J. McGuire, T. M. Newcomb, M. J. Rosenberg & P. H. Tannenbaum, (Eds.), *Theories of cognitive consistency: A sourcebook* (pp. 670-683). Chicago, IL: Rand McNally.

Kintsch, W. (1998). *Comprehension*. Cambridge, UK: Cambridge University Press.

Kitcher, P. (1993). *The advancement of science*. New York: Oxford University Press.

Klein, G. (1998). *Sources of Power: How People Make Decisions*. Cambridge: The MIT Press.

Kuhn, T. (1970). *The structure of scientific revolutions*. Chicago: University of Chicago Press.

Landauer, T. K. (1986). How much do people remember? Some estimates of the quantity of learned information in long-term memory, *Cognitive Science*, 10, 477-493.

Lawson, A., & Thompson, L. (1988). Formal reasoning ability and misconceptions concerning genetics and natural selection. *Journal of Research in Science Teaching*, 25, 733-746.

Lesgold, A., Rubinson, H., Feltovich, P., Glaser, R., Klopfer, D., Wang, Y. (1988). Expertise in a complex skill: Diagnosing x-ray pictures. In: M. Chi, R. Glaser, & M. Farr (Eds.), *The Nature of Expertise* (pp. 311-342). Hillsdale NJ: Erlbaum.

Licht, P. (1987). A strategy to deal with conceptual and reasoning problems in introductory electricity education. In: Novak, J. (Ed), Proceedings of the 2nd International Seminar "Misconceptions and Educational Strategies in Science and Mathematics", Vol. II. Ithaca: Cornell University, 275-284.

Limon, M. (2002). Conceptual change in history. In: M. Limon & L. Mason (Eds.), *Reconsidering Conceptual Change: Issues In Theory and Practice*, pp. 259-289, Kluwer, Academic Publishers.

Litman, L. & Reber, A. (in press). Implicit and explicit thought. In: Holyoak, K.J. and Morrison, B.G. (Eds.), *Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press.

Lovett, M.C. (2002). Problem solving. In: H. Pashler & D. Medin (Eds.), *Stevens' Handbook of Experimental Psychology, Volume 2: Memory and Cognitive Processes*, New York: Wiley, pp. 317-362.

Lovett, M.C. & Anderson, J. (in press). Thinking as a production system. In: Holyoak, K.J. and Morrison, B.G. (Eds.), *Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press.

McCloskey, M. (1983). Naïve theories of motion. In: D. Gentner & A.L. Stevens (Eds.), *Mental models* (pp. 299-323). Hillsdale, NJ: Erlbaum.

McGuire, W.J. (1968). Theory of the structure of human thought. In: R.P. Abelson, E. Aronson, W.J. McGuire, T.M. Newcomb, M.J. Rosenberg, & P.H. Tannenbaum (Eds.), *Theories of Cognitive Consistency: A Sourcebook* (pp. 140-162), Chicago: Rand McNally and Company.

Machamer, P., & Woody, A. (1992). Models of intelligibility in science: Using the balance as a model for understanding the motion of bodies. In: S. Hills (Ed.), *The history and philosophy of science in science education* (Vol 2, pp. 95-111). Kingston, Ontario: Queen's University.

Markman, A.B. (1999). *Knowledge representation*. Mahwah, NJ: Erlbaum.

Marshall, S.P. (1995). *Schemas in problem solving*. Cambridge, UK: Cambridge University Press.

Means, M.L. & Voss, J.F. (1985). Star wars: A developmental study of expert and novice knowledge structures. *Journal of Memory and Language*, 24, 746-757.

Medin, D.L., & Shoben, E.J. (1988). Context and structure in conceptual combination.

Cognitive Psychology, 20, 158-190.

Miller, G.A. (1996) *The science of words*. New York: Scientific American Library.

Minsky, M.A. (1975). A framework for the representation of knowledge. In: P. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill.

Minstrel, J. (1982a, January). Explaining the "at rest" condition of an object. *The Physics Teacher*, pp. 10-14.

Miyake, N. (1986). Constructive interaction and the iterative process of understanding, *Cognitive Science*, 10, 151-177.

Norman, D.A., & Rumelhart, D.E. (1975). *Explorations in cognition*. San Francisco, CA: W. H. Freeman.

Nussbaum, J. (1979). Children's conception of the Earth as a cosmic body: a cross-age study. *Science Education*, 63, 83-93.

Nussbaum, J. (1985). The Earth as a cosmic body. In: R. Driver, E. Guesne, & A. Tiberghien (Eds.), *Children's ideas in science*. Milton Keynes, UK: Open University Press.

Nussbaum, J., & Novak, J. D. (1976). An assessment of children's concepts of the Earth utilizing structured interviews. *Science Education*, 60, 535-550.

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

Ohlsson, S. (1993) The interaction between knowledge and practice in the acquisition of cognitive skills. In: A. Meyrowitz and S. Chipman, (Eds.), *Foundations of knowledge acquisition: Cognitive models of complex learning* (pp. 147-208). Norwell, MA: Kluwer Academic Publishers.

Ohlsson, S. (1994). Declarative and procedural knowledge. In: T. Husen & T. Neville-Postlethwaite, (Eds.), *The International Encyclopedia of Education* (Vol. 3, 2nd ed., pp.1432-1434). London, UK: Pergamon Press.

Ohlsson, S. (1996). Learning from performance errors. *Psychological Review*, 103, pp. 241-262.

Ohlsson, S. (1999). Theoretical commitment and implicit knowledge: Why anomalies do not trigger learning. *Science & Education*, 8, 559-574.

Ohlsson, S. (2002). Generating and understanding qualitative explanations. In: J. Otero, J.A. Leon, & A.C. Graesser, (Eds.), *The psychology of science text comprehension* (pp. 91-128). Mahwah, NJ: Erlbaum.

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

Ohlsson, S., & Hemmerich, J. (1999). Articulating an explanation schema: A preliminary model and supporting data. In: M. Hahn & S. Stones (Eds.), *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society* (pp. 490-495). Mahwah, NJ: Erlbaum.

Ohlsson, S., Moher, T. G., & Johnson, A. (2000). Deep learning in virtual reality: How to teach children that the Earth is round. In L. R. Gleitman and A. K. Joshi, (Eds.) *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society* (pp. 364-368). Mahwah, NJ: Erlbaum.

Ohlsson, S., & Regan, S. (2001). A function for abstract ideas in conceptual learning and discovery. *Cognitive Science Quarterly*, 1, 23-277.

Pennington, N. (1987). Comprehension strategies in programming. In G.M. Olson, S. Sheppard, & E. Soloway (Eds.), *Empirical Studies of Programmers: Second Workshop*. Norwood, NJ: Ablex.

Piaget, J. (1985). *The Equilibration of Cognitive Structures*. Chicago, IL: University of Chicago Press.

Piaget J. & Inhelder B. (1956). *The Child's Conception of Space*. London: Routledge and Kegan Paul.

Posner, G.J., Strike, K.A., Hewson, P.W., & Gertzog, W.A. (1982). Accommodation of a scientific conception: Toward a theory of conceptual change. *Science Education*, 66, 211-27.

Quillian, M. R. (1968). Semantic memory. In M. Minsky, (Ed.), *Semantic information processing* (pp. 227-270). Cambridge, MA: MIT Press.

Reisberg, D. (2001). *Cognition: Exploring the science of the mind* (2nd ed.). New York: Norton.

Rips, L. & Medin, D. (in press). Concepts and Categories: Memory, Meaning and Metaphysics. In: Holyoak, K.J. and Morrison, B.G. (Eds.), *Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press.

Rokeach, M. (1960). *The open and closed mind*. New York: Basic Books.

Rokeach, M. (1970). *Beliefs, attitudes, and values; A theory of organization and change*. San Francisco, CA: Jossey-Bass.

Robin, N., & Ohlsson, S. (1989). Impetus then and now: A detailed comparison between Jean Buridan and a single contemporary subject. In D. E. Herget (Ed.), *The history and philosophy of science in science teaching* (pp. 292-305). Tallahassee, FL: Florida State University.

Rosenzweig, M. R. (2001). Learning and neural plasticity over the life span. In P. E. Gold and W. T. Greenough (Eds.), *Memory consolidation* (pp. 275-294). Washington, DC: American Psychological Association.

Rumelhart, D. E., & Norman, D. A. (1978). Accretion, tuning and restructuring: Three modes of learning. In J. W. Cotton & R. Klatzky (Eds.), *Semantic factors in cognition*. Hillsdale, NJ: Erlbaum.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986). Schemata and sequential thought processes in PDG models. In: J. McClelland and D.E. Rumbelhart (Eds.), *Parallel distributed processing: Exploration sin the microstructure of cognition* (Vol 2, pp. 7-

57). Cambridge, MA: MIT Press.

Schank, R. (1972). Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, 3, 552-631.

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

Schank, R.C., & Abelson, R.P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Schultz, T. R., & Lepper, M. R. (1996). Cognitive dissonance reduction as constraint satisfaction. *Psychological Review*, 103, 219-240.

Shatz, M. and Gelman, R. (1973). The development of communication skills: Modifications in the speech of young children as a function of listener. Monographs of the society for research in child development. Serial number 152 38(5).

Simon, H.A. & Gilmarin, K. (1973). A simulation of memory for chess positions, *Cognitive Psychology*, 5, 29-46.

Slovan, S. & Lagnado, D.A. (in press). The Problem of Induction. In: Holyoak, K.J. and Morrison, B.G. (Eds.), *Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press.

Slotta, J.D., Chi, M.T.H. & Joram, E. (1995). Assessing students' misclassifications of physics concepts: An ontological basis for conceptual change. *Cognition and Instruction*, 13(3), 373-400.

Smith, III, J. P., DiSessa, A. A., & Roschelle, J. (1995). Misconceptions reconceived: A constructivist analysis of knowledge in transition. *The Journal of the Learning Sciences*, 3, 115-163.

Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, 12, 485-527.

Soloway, E.M. (1978). Learning = interpretation + generalization: A case study in

knowledge-directed learning. Unpublished doctoral dissertation, University of Massachusetts.

Soloway, E. & Erhlich, K. (1984). Empirical studies of programming knowledge. *IEEE Transactions on Software Engineering*, 10, 595-609.

Spilich, G.J., Vesonder, G.T., Chiesi, H.L., & Voss, J.F. (1979). Test processing of domain-related information for individuals with high and low domain knowledge. *Journal of Verbal Learning and Verbal Behavior*, 18, 275-290.

Squire, L. R. (1987). *Memory and brain*. New York: Oxford University Press.

Stokes, A.F., Kemper, K. & Kite, K. (1997). Aeronautical decision making, cue recognition, and expertise under time pressure. In C.E. Zsombok & G. Klein (Eds.), *Naturalistic Decision Making. Expertise: Research and Applications*. Mahwah, NJ: Lawrence Erlbaum.

Strike, K. A., & Posner, G. J. (1985). A conceptual change view of learning and understanding. In L. West and L. Pines (Eds.), *Cognitive structure and conceptual change* (p. 211-231). New York: Academic Press.

Strike, K. A., & Posner, G. J. (1992). A revisionist theory of conceptual change. In R. A. Duschl & R. J. Hamilton (Eds.), *Philosophy of science, cognitive psychology, and educational theory and practice*. New York: State University of New York Press.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, 12, 435-467.

Thagard, P. (1992a). *Conceptual revolutions*. Princeton: Princeton University Press.

Thagard, P. (1992b). Adversarial problem solving: Modeling an opponent using explanatory coherence. *Cognitive Science*, 16, 123-149.

Thagard, P. (2000). *Coherence in thought and action*. Cambridge, MA: MIT Press.

Thompson, P. (1989). *The structure of biological theories*. New York: State University of New York Press.

Thorndyke, P. W. (1984). Applications of schema theory in cognitive research. In J. R. Anderson and S. M. Kosslyn (Eds.), *Tutorials in learning and memory* (pp. 167-). San Francisco, CA: W. H. Freeman.

Trabasso, T., & van den Broek, P. (1985). Causal thinking and the representation of narrative events, *Journal of Memory and Language*, 24, 612-630.

Tuckerman, N. & Dunnan, N. (1995) *The Amy Vanderbilt complete book of etiquette*. New York: Doubleday.

VanLehn, K. (1989). Problem solving and cognitive skill acquisition. In: M.I. Posner (Ed.), *Foundations of Cognitive Science*, pg. 527-579, Cambridge: MIT Press.

Vosniadou, S. (1994). Capturing and modeling the process of conceptual change. *Learning and Instruction*, 4, 45-69.

Vosniadou, S. (1994). Universal and culture-specific properties of children's mental models of the earth. In L. Hirschfeld and S. Gelman (Eds.), *Mapping the mind*. Cambridge, MA: Cambridge University Press.

Vosniadou, S., & Brewer, W.F. (1992). Mental models of the earth: A study of conceptual change in childhood. *Cognitive Psychology*, 24, 535-585.

Voss, J.F., Vesonder, G.T., & Spilich, G.J. (1980). Text generation and recall by high-knowledge and low-knowledge Individuals. *Journal of Verbal Learning and Verbal Behavior*, 19, 651-667.

White, B. & Frederiksen, J. (1990). Causal model progressions as a foundation for intelligent learning environments. *Artificial Intelligence*, 42, 99-157.

Wilensky, U., & Resnick, M. (1999). Thinking in levels: A Dynamic Systems Perspective to making sense of the world. *Journal of Science Education and Technology*, 8, 3-19.

Wilson, R.A. & Keil, F.C. (2000). The shadows and shallows of explanation. In F.C. Keil & R.A. Wilson (Eds.), *Explanation and Cognition*. A Bradford Book, The M.I.T. Press, Cambridge, Mass.

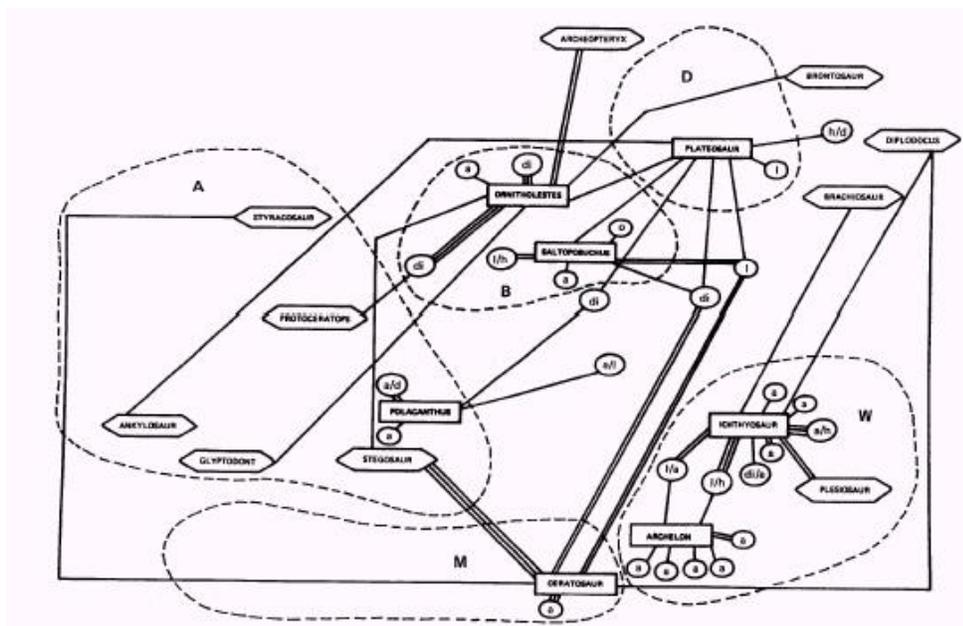
Winograd, T. (1975). Frame representations and the declarative/procedural controversy. In D. Bobrow and A. Collins (Eds.) *Representation and Understanding: Studies in Cognitive Science* (pp. 185-210). New York, NY: Academic press.

Wiser, M. & Carey, S. (1983). When heat and temperature were one. In: D. Gentner & A. Stevens (Eds.), *Mental Models*, pp. 267-297, Hillsdale, NJ: Erlbaum.

Figures

Figure 1: A child's representation of 20 familiar dinosaurs (taken from Chi & Koeske, 1983).

Figure 2: A child's representation of 20 less familiar dinosaurs (taken from Chi & Koeske, 1983).



Endnotes

¹ In social cognition research, intuitive theories are called *belief systems* (Fishbein & Ajzen, 1975; Rokeach, 1960, 1970). Although the two constructs of intuitive theory and belief system are essentially identical, this connection between social and cognitive psychology has been overlooked on both sides (but see Schultz & Lepper, 1996).