UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Cognitive Flexibility Theory: Advanced Knowledge Acquisition in Ill-Structured Domains

Permalink

Journal

Authors

Spiro, Randolph J.
Coulson, Richard L.
Feltovich, Paul J.
et al.

Publication Date

1988

Peer reviewed

COGNITIVE FLEXIBILITY THEORY:
ADVANCED KNOWLEDGE ACQUISITION IN ILL-STRUCTURED DOMAINS

Rand J. Spiro
University of Illinois at Urbana-Champaign

Richard L. Coulson, Paul J. Feltovich, and Daniel K. Anderson
Southern Illinois University School of Medicine

Advanced knowledge acquisition in a subject area is different in many important ways from introductory learning (and from expertise). In this paper we discuss some of the special characteristics of advanced learning of complex conceptual material. We note how these characteristics are often at odds with the goals and tactics of introductory instruction and with psychological biases in learning. We allude to our research in biomedical cognition that has revealed a substantial incidence of misconception attributable to various forms of oversimplification, and we outline the factors that contribute to suboptimal learning at the advanced stage. We then sketch a theoretical orientation for more successful advanced knowledge acquisition in ill-structured domains, Cognitive Flexibility Theory.

The Goals of Advanced Knowledge Acquisition

In our work we have been interested in "advanced knowledge acquisition"--learning beyond the introductory stage for a subject area, but before the achievement of practiced expertise that comes with massive experience. This often neglected intermediate stage is important because the aims and means of advanced knowledge acquisition are different from those of introductory learning. In introductory learning the goal is often mere exposure to content and the establishment of a general orientation to a field; objectives of assessment are likewise confined to the simple effects of exposure (e.g., recognition and recall). At some point in learning about a knowledge domain the goal must change; at some point students must "get it right." This is the stage of advanced knowledge acquisition [7, 10, 12]: the learner must *attain a deeper understanding of content material, reason with it, and apply it flexibly in diverse contexts.* Obstacles to advanced knowledge acquisition include conceptual complexity and the increasing ill-structuredness that comes into play with more advanced approaches to a subject area. By *ill-structuredness* we mean that many concepts (interacting contextually) are pertinent in the typical case of knowledge application, and that their patterns of combination are inconsistent across case applications of the same nominal type. (See [12] for a more detailed treatment of the nature and consequences of ill-structuredness.)

The methods of education in introductory and advanced learning seem, in many ways, to be at odds. For example, compartmentalizing knowledge, presenting clear instances (and not the many pertinent exceptions), and employing reproductive memory criteria are often in conflict with the realities of advanced learning-- knowledge which is intertwined and dependent, has significant context-dependent variations, and requires the ability to respond flexibly to "messy" application situations. These discrepancies in aims and tactics (along with many others that we have observed) raise the possibility that introductory learning, even when it is "successful," lays foundations in knowledge and in an approach to learning that interfere with advanced acquisition. As we have seen repeatedly demonstrated, that possibility is an actuality [6, 7, 10, 12].

Deficiencies in Advanced Knowledge Acquisition

Medical school is an archetype of an advanced knowledge acquisition setting [7]. Medical students have already had introductory exposure to many of the subject areas of biological science that they go on to study in medical school, but they are certainly not yet expert. Furthermore, the goals of medical education are clearly those of advanced knowledge acquisition. Important aspects of conceptual complexity must now be mastered (superficial familiarity with key concepts is no longer sufficient); and the ability to apply knowledge

from formal instruction to real-world cases is certainly something that is expected of those studying to be physicians.

In our laboratory we have been studying medical students' learning, understanding and application of important but difficult biomedical science concepts. This effort has revealed widely held systematic misconceptions among students, despite their having been exposed to appropriate information [6, 7, 10, 12]. Stubborn misconception, notwithstanding usual classroom efforts at instruction, has been found for difficult concepts in other areas as well (e.g., physics: [3, 13]).

The biomedical misconceptions that we have identified are of various kinds [7, 10]. These include *contentive* errors, often involving overgeneralization; for example, areas of subject matter are seen as being more similar than they really are. Errors attributable to dysfunctional biases in *mental representation* are also observed; for example, dynamic processes are often represented more statically. *Prefigurative "world views"* that underlie learners' understanding processes also cause problems; for example, the presupposition that the world works in such a way that "parts add up to wholes" leads students to decompose complex processes into components that are treated (mistakenly) as independent. Furthermore, at all these levels misconceptions interact in *reciprocally supportive* ways, and combine to yield higher order misconceptions [6, 7]. Failures of understanding *compound* themselves, building up durable chains of larger scale misconception.

Reductive biases: The pervasive role of oversimplification in the development of misconceptions. A predominant share of the misconceptions (and networks of misconception) that we have identified reflect one or another kind of *oversimplification* of complex material--associated with learners' earlier experiences with introductory learning, and even influenced by many experiences with advanced learning. Misconceptions of advanced material result both from interference from earlier, simplified treatments of that material and from a prevailing mode of approaching the learning process in general that fosters simplificational strategies and leaves learners without an appropriate cognitive repertoire for the processing of complexity [7, 10, 12].

We have termed the general tendency to reduce important aspects of complexity the *reductive bias*. Several forms of the bias have been identified, selected examples of which follow (see [6, 7, 10] for *examples of biomedical misconceptions* corresponding to the types of reductive bias listed):

1. Oversimplification of complex and irregular structure. Superficial similarities among related phenomena are treated as unifying characteristics. Interacting components are treated as independent. Incomplete conceptual accounts are presented (or accepted by the learner) as being comprehensive. Instances that are referred to as belonging to the same generic category are treated in a uniform manner despite their being highly diverse. The irregular is treated as regular, the nonroutine as routine, the disorderly as orderly, the continuous as discrete, the dynamic as static, the multidimensional as unidimensional. (This first reductive bias is the most general one, encompassing many of the specific ones listed below.)

2. Overreliance on a single basis for mental representation. A single, encompassing representational logic is applied to complex concepts and phenomena that are inadequately covered by that logic. For example: Understanding of a new concept is reduced to the features of a (partially) analogous concept. New, highly divergent examples are understood by exclusive reference to a single prototype. A single schema or theory is proffered and preferred, despite the fact that its coverage is significantly incomplete. Complexly multifaceted content has its understanding narrowed to just those aspects covered by a single organizational scheme. And so on.

3. Overreliance on "top down" processing. Understanding and decision making in knowledge application situations (i.e., cases) rely too exclusively on generic abstractions (i.e., concepts, theories, etc.); detailed knowledge of *case* structure is not used enough (i.e., knowledge of "how cases go," as well as reasoning from specific case precedents).

4. Context-independent conceptual representation. The contexts in which a concept is relevant are treated as having overly uniform characteristics. This promotes the representation of conceptual knowledge in a manner too abstract for effective application (i.e., without sufficient regard for the specifics of application in context). Concepts are insufficiently tailored to their uses; concepts are not recognized as relevant when, in fact, they are; and concepts are mistakenly judged to be relevant in contexts where they are not.

5. Overreliance on precompiled knowledge structures. *Fixed* protocols or rigidly *prepackaged* schemas are presented to learners and used by them as recipes for what to do in new cases.

6. Rigid compartmentalization of knowledge components. Components of knowledge that are in fact interdependent are treated as being separable from each other. Learners develop mistaken beliefs in the independence of the components. Relatedly, where knowledge components do function independently, it may nevertheless be the case that conveying relationships between their conceptual structures would aid understanding; these connections are not drawn. When components are interrelated, there is a tendency to use just one linkage scheme, thereby underrepresenting the richness of interconnection in the system and promoting narrow, doctrinaire viewpoints (see the problem of single representations).

7. Passive transmission of knowledge. Knowledge is preemptively encoded under a scheme determined by external authority (e.g., a textbook) or a scheme which facilitates delivery and use. Knowledge is "handed" to the learner. The preemptive encoding is passively received by the learner, and useful benefits that result from personalized knowledge representations, derivable from active exploration and involvement in the subject area, do not develop. When active, participatory learning is encouraged, adequate support for the management of increased indeterminacy and cognitive load is not provided (e.g., mentor guidance, memory aids, etc.).

The next section will outline our theoretical approach to *remedying* the problems of advanced knowledge acquisition that are caused by these reductive biases.

Cognitive Flexibility Theory:
Themes of Advanced Knowledge Acquisition

Where has our research on the *problems* of advanced knowledge acquisition led us? To an overall theoretical orientation that in many ways derives its fundamental themes from the specific nature of those learning problems, as the problems relate to the characteristics of ill-structured domains and the special goals of advanced knowledge acquisition (i.e., mastery of conceptual complexity and knowledge application/transfer).

In this section we provide a brief discussion of our most fundamental, theoretically motivated remedies for the problems of advanced knowledge acquisition. The following themes constitute different facets of what we call *cognitive flexibility* [12]. The themes are, in a sense, *conditions* for developing mastery of complexity and knowledge transferability. Each of the headlined theoretical commitments has received some form of implementation, either in our experiments or in our theory-based computer hypertext systems (including one prototype that implements the theory's principles of advanced knowledge acquisition in cardiovascular medicine, the *Cardioworld Explorer*). Given the extreme limitations of space, the themes are discussed schematically and in the abstract; detailed development of theoretical rationales, examples of our concrete instantiations of the themes (in the biomedical domain and others that we have studied), and patterns of empirical support for our claims can be found in our cited papers.

1. Avoidance of Oversimplification and Overregularization. Because of the strong bias towards oversimplification that we have observed, it is clear that advanced knowledge acquisition must place a high premium on making salient those ways that knowledge is not as simple and orderly as it might first seem in introductory treatments. Where the problem is so often a presumption of simplicity and regularity, the remedy is to take special measures to *demonstrate complexities and irregularities*. It is important to lay bare the limitations of initial, first pass understandings, to highlight exceptions, to show how the superficially similar is

dissimilar and how superficial unities are broken. Where conceptual error frequently occurs from atomistic decomposition of complexly interacting information, followed by misguided attempts at "additive" reassembly of the decomposed elements, the remedy is to take pains to *highlight component interactions*, to clearly demonstrate the intricate patterns of *conceptual combination*.

This is a very general theme, encompassing many of the others that follow in this list. Cognitive flexibility involves the *selective* use of knowledge to *adaptively fit* the needs of understanding and decision making in a particular situation; the potential for maximally adaptive *knowledge assembly* depends on having available as full a representation of complexity to draw upon as possible.

2. Multiple Representations. Single representations (e.g., a single schema, organizational logic, line of argument, prototype, analogy, etc.) will miss important facets of complex concepts. Cognitive flexibility is dependent upon having a diversified repertoire of ways of thinking about a conceptual topic. Knowledge that will have to be used in many ways has to be learned, represented, and tried out (in application) in many ways.

The use of multiple representations is important at different levels. For example, we have found multiple analogies to be very useful in understanding complex individual concepts ([10]; see the example below of force production by muscle fibers; see also [5, 14]). However, the importance of multiple representations may be even more important for larger units of analysis. For example, we have found that students' understandings of the entire domain of biomedical knowledge is adversely affected by the tendency to use just one way of modeling the various phenomena they encounter, one that comes from the *metaphor of the machine*. This one "lens" leads them to take for granted certain issues related to the nature of explanations, the structure of mental models of functional systems, and so on. These students develop understandings that do not capture important aspects of the biomedical domain (e.g., inherently organic processes). Their understandings would be more complete if they were to augment the selective view that results from their mechanistic bias with other understandings that selectively emerge from the unique aspects of other cognitive "lenses," for example, from *organicist* metaphors [7].

The need for multiple representations applies not only to complex concepts, but to cases as well. In an ill-structured domain, cases (examples, occurrences, events--occasions of use of conceptual knowledge) tend to be complex and highly variable, one to the next. The complexity of cases requires that they be represented from multiple theoretical/conceptual perspectives--if cases are treated narrowly by characterizing them using a too limited subset of their relevant perspectives, the ability to process future cases will be limited. First, there will be an assumption that cases are simpler than they in fact are, and attempts to deal with new cases will prematurely conclude after they are only *partially* analyzed. Second, there will be insufficient preparedness to deal with the specific patterns of interaction of theoretical/conceptual perspectives within cases. Third, to the extent that performance in future cases will require reasoning from sets of precedent cases (which is always a greater need in ill-structured domains), the likelihood of having case representations available in prior knowledge which are maximally apt in their relation to some new case is lessened to the extent that cases are narrowly represented in memory. This is especially so when there is substantial across-case dissimilarity; the relative novelty of a new case in an ill-structured domain will require more elaborate efforts to find appropriate precedents--the wider the variety that is available, the better the chances of finding a fit.

An Example of Multiple Representations: Integrated Multiple Analogies for Complex Concepts. As we have said, our studies of medical students have indicated that one of the most serious contributors to the problems of advanced knowledge acquisition is the use of a single knowledge representation. Complex concepts can rarely be adequately represented using a single schema, theoretical perspective, line of exposition, and so on. Nevertheless, in practice, complex concepts frequently are represented in some single fashion, with substantial consequences.

Our remedy has been to approach learning in all of the domains that we have studied with the goal of promoting *multiple representations* (e.g., multiple precedent cases for a new case; multiple organizational schemes for representing the same content material in our computer hypertexts; etc.). Here we will briefly

consider just the case of analogy. We have discovered a large number of misconceptions that result from the overextended application of analogies [10]. To combat the negative effects of a powerful and seductive single analogy, we employ *sets of integrated multiple analogies*. Whenever a source concept in an analogy is missing important aspects of a target concept, or the source concept is in some way misleading about the target concept, we introduce *another* analogy to counteract those specific negative effects of the earlier analogy.

So, where we find that misconceptions about the nature of force production by muscle fibers often develop because of a common analogy to the operation of rowing crews (sarcomere "arms" and oars both generate force by a kind of "pulling"), other analogies are introduced to mitigate the negative effects of the limited rowing crew analogy [10]. An analogy to turnbuckles corrects misleading notions about the nature of relative movement and the gross structures within the muscle. An analogy to "finger handcuffs" covers important information missing in the rowing crew analogy about limits of fiber length (the elastin covering on muscle fiber bundles constricts at long lengths, stopping extension in a manner similar to the cross-hatched finger cuffs when you try to pull a finger out of each end). And so on. A composite imaging technique that helps the user to integrate the multiple analogies, so that the correct aspects of each analogy can be selectively instantiated in relevant contexts of use of the target concept, has also been developed. The procedure facilitates the learning of a concept (through the pedagogical benefits of analogy), while maintaining the integrity of the concept's complexities (by using multiple analogies to cover the concept's multifacetedness and to vitiate the force of incorrect aspects of any single analogy). (Also see [2].)

Theory-based hypertext systems to implement the themes of advanced knowledge acquisition in ill-structured domains: The importance of revisiting and rearranging in the development of multiple representations. Much of the work on computer hypertext systems has been driven by the power of the technology, rather than by a coherent view of the cognitive psychology of nonlinear and multidimensional learning and instruction. In contrast, our hypertext approaches have a basis in cognitive theory--they derive from the themes of Cognitive Flexibility Theory. And their realm of operation is specified; they are especially targeted at advanced knowledge acquisition in ill-structured domains. (There is no point in imposing the extra cognitive load of nonlinearity and multidimensionality if the domain being studied is simple and well-structured, or if the goals of learning are the more easily attainable ones of introductory treatments.) We will briefly characterize our approach to implementing Cognitive Flexibility Theory in computer hypertext systems.

Our hypertext systems build multiple representations in a manner that can be understood using a metaphor of landscape exploration. Deep understanding of a complex landscape will not be obtained from a single traversal. Similarly for a *conceptual* landscape. Rather, *the landscape must be criss-crossed in many directions* to master its complexity and to avoid having the fullness of the domain attenuated [12, 15]. The same sites in a landscape (the same cases or concepts in a knowledge domain) should be *revisited* from different directions, thought about from different perspectives, and so on. There is a limit to how much understanding of a complex entity can be achieved in a single treatment, in a single context, for a single purpose. By *repeating* the presentation of the same complex case or concept information in *new contexts*, additional aspects of the multifacetedness of these "landscape sites" are brought out, enabling the kind of rich representations necessary in a complex and ill-structured domain. Thus, cognitive flexibility is fostered by a flexible approach to learning and instruction. The same content material is covered in different ways, at different times, in order to demonstrate the potential flexibility of use inherent in that content ([11, 12]).

3. Centrality of Cases. The more ill-structured the domain, the poorer the guidance for knowledge *application* that "top-down" structures will generally provide. That is, the way abstract concepts (theories, general principles, etc.) should be used to facilitate understanding and to dictate action in naturally occurring cases becomes increasingly indeterminate in ill-structured domains. The application of knowledge to cases in an ill-structured domain (i.e., a domain in which cases are individually multidimensional, and irregularly related one to the next) cannot be prescribed in advance by general principles. This is because, in ill-structured domains, there is great variability from case to case regarding which conceptual elements will be relevant and in what pattern of combination. In an ill-structured domain, general principles will not capture

enough of the structured dynamics of cases; increased flexibility in responding to highly diverse new cases comes increasingly from reliance on reasoning from precedent cases.

Thus, examples/cases cannot be assigned the ancillary status of merely illustrating abstract principles (and then being discardable); the cases are key--examples are necessary, and not just nice [7, 11, 12].

4. Conceptual Knowledge as Knowledge-in-Use. Not only is it more difficult to count on top down prescriptions for performance in new cases in an ill-structured domain (i.e., abstract concepts/theories inadequately determine responses to new cases), but there is also considerable indeterminateness in defining conditions for *accessing* conceptual structures in the first place, to engage the guidance the conceptual structures *do* offer. It is not that abstract knowledge has no role in ill-structured domains, but that its role is highly intertwined with that of case-centered reasoning. Put another way, in an ill-structured domain there will be greatly increased variability across cases in the way the same concept is used or applied. Thus it is harder to get from features of cases to the concepts that might need to be applied to those cases. And it is harder to apply a concept, once accessed, if it has many different kinds of uses across cases--concepts must be *tailored* to their application contexts. The Wittgensteinian dictum that meaning is determined by use clearly applies in ill-structured domains. If a concept's meaning in use cannot be determined universally across cases (as in an ill-structured domain), then one must pay much more attention to the details of how the concept is used--*knowledge in practice*, rather than in the abstract [11, 12, 15].

In medical training, this issue of *variability and combination in concept instantiation* has an obvious implication for the traditional difficulty of integrating the biomedical basic science parts of the curriculum with the clinical parts. Physicians' practice would be improved if in problematic situations they could apply the interacting basic biomedical science concepts that underlie the clinical situation that is posing the problem. However, it is very difficult for medical students to learn how to get to the basic science concepts from clinical presenting features, partly because of the great variability across clinical cases in the way those concepts get instantiated. A key feature of our *Cardioworld Explorer* hypertext is that it permits the learner to selectively examine the full range of uses of any selected basic science concept (or any selected combination of concepts) across cases with differing clinical features, teaching the patterns of concept application and thus facilitating *access* to conceptual information in clinical contexts (as well as fostering an understanding of the different ways that a given concept has to be tailored to be clinically relevant).

Again, in an ill-structured domain the meaning of a concept is intimately connected to its patterns of use. When the uses (instances, cases) of the same concept have a complex and irregular distribution (i.e., the domain is ill-structured), adequate prepackaged prescriptions for proper activation of the concept cannot be provided (i.e., *concept instantiation is non-routine*). Instead, greater weight (than in a well-structured domain) must be given to activating concepts in a new case by examination of family resemblances across the features of *past cases* that have been called (labeled as instances of) that concept.

5. Schema Assembly (from Rigidity to Flexibility). In an ill-structured domain, emphasis must be shifted from *retrieval* of intact, rigid, precompiled knowledge structures, to *assembly* of knowledge from different conceptual and precedent case sources to adaptively fit the situation at hand [9, 12]. This follows, again, from characteristics of ill-structured domains. Since ill-structuredness implies kinds of complexity and irregularity that militate against the use of knowledge structures that assume routinizability across cases, the role of intact schema retrieval must be diminished--greater across-case differences cause a necessary decline in the ability of any large, single precompilation to fit a wide variety of cases. In complex and ill-structured domains, one cannot have a prepackaged schema for everything! As ill-structuredness increases, the use of *rigid* knowledge structures (i.e., the same precompiled knowledge structure used for many cases) must be replaced by *flexible*, recombinable knowledge structures. For any particular case, many *small* precompiled knowledge structures will need to be used. And there will be relatively little repetition of patterns across case-specific assemblies of these smaller pieces of precompiled knowledge. Accordingly, in knowledge acquisition for cognitive flexibility, the "storage of fixed knowledge is devalued in favor of the *mobilization of potential knowledge*" ([12], p. 181). (See also [8].)

6. Noncompartmentalization of Concepts and Cases (Multiple Interconnectedness). Because of the complex and irregular way that abstract conceptual features weave through cases/examples in ill-structured domains, knowledge cannot be neatly compartmentalized. In order to enable the situation-dependent, adaptive schema assembly from disparate knowledge sources that characterizes cognitive flexibility, those multiple sources must be highly interconnected. Concepts cannot be treated as separate "chapters." Retroactive assembly of independently taught, and noninterrelated, constituent conceptual aspects too often fails. Also, although cases have to be focused on separately, so that the complexity of case structure is conveyed, they should not be taught in *just* that way--connections across cases must also be established. Rather than relegating concepts or cases to separate compartments, chapters, and so on, our systems strive for *multiple interconnectedness (of cases and concepts) along multiple conceptual and clinical dimensions.*

Our approach to fostering multiple interconnectedness of knowledge representations in our hypertexts is to code case segments with a multidimensional vector indicating the relevance of a variety of thematic/conceptual dimensions to that case segment [11]. (Positive values in the vector also point to commentary, providing expert guidance about the nature of the conceptual dimension's instantiation in that particular case segment; this helps with the problem of teaching conceptual knowledge-in-use discussed earlier). Then, as the hypertext program guides the learner in criss-crossing the domain's "landscape," by exploring patterns of overlap in the vectors for different case segments, knowledge representations are built up in which parts of cases are connected with many parts of other cases, along many conceptual/theoretical dimensions of case-segment similarity. In that way, many alternative paths are established to get from one part of the overall knowledge base to any other part of the knowledge base that aspects of some future case may signal as relevant. Thus, the potential for flexible, situation-adaptive schema assembly is fostered (along with such other virtues as the establishment of multiple routes for memory access to any node in the system).

So, for example, in the *Cardioworld Explorer* segments of clinical cases are encoded with a vector of clinical and basic biomedical science themes that are relevant to each segment. The system can then establish connections between a segment of one case and segments of many other cases, along the various (conceptual and clinical) thematic dimensions represented in the vector. In case-based instruction, it is often true that there are important, instructive relationships between an aspect of one case and aspects of others. Such relationships are rarely brought out. Our hypertext systems capture these many lessons that are missed in strict case-by-case (or problem-by-problem) instruction. In an ill-structured domain, facilitating retrieval of multiple (partial) precedents is important, because understanding what to do in a given case context will usually require reference to more than any single prototype--the case in question will be "kind of like this earlier one, kind of like that one," and so on. Also, understanding of the case in question will require that various concepts be brought to bear and integrated; this, too, is facilitated by the multiple conceptual coding scheme employed in our systems.

There are several other benefits of the multiple-conceptual coding of multiple case segments. A power/efficiency advantage is that it allows the hypertexts to automatically generate large numbers of lessons (many "landscape criss-crossings"). If, for example, each of 20 cases is divided into an average of 10 case segments, each with a value on 15 relevant thematic dimensions, there is a many-fold increase in the number of possible automatizable instructional comparisons and contrasts that results from having 200 case segments (instead of 20 full cases) intertwined by relationships in the 15-slot vector.

Also, the use of case segments prevents the subsumption to a "common denominator" that occurs when larger structural units are used: an interesting local element of a case will tend to get lost if it has features that are not present in other parts of the case (when the monolithic case is the structural unit). Using small case segments (minicases) helps to retain the *plurality* of situations.

There is another virtue of the division into case segments and the multidimensional coding of the segments that relates to keeping case understanding from being overly simplified. In an ill-structured knowledge domain, by definition, there is sufficient variability across cases (due in part to the interaction of the many

factors that make up complex cases) that the set of cases that might be nominally grouped together under some schema or classification will be greatly variable in their particulars. A case, instead of being represented as one kind of thing, conveying one kind of "lesson," is instead clearly shown to the learner to to be *many things*. Cases of the same nominal type have different segments or scenes that are demonstrated not to be the same, and each of the segments is shown to have multiple significances. Therefore, the common temptation to nest cases uniquely under a single superordinate conceptual category will be resisted, making it less likely that the complex relationships among cases in a domain will be artificially regularized. In an ill-structured domain, cases are related to many different concepts of the domain, and it promotes dysfunctional simplification to hierarchically nest or "slot" cases under single conceptual categories (e.g., "The following cases are examples of X [only]"). When there is considerable across-case variability, as there will be in an ill-structured domain, cognitive flexibility requires that case information be coded conceptually for the many different kinds of use that new situations may require.

The thematic coding scheme and the landscape criss-crossing system of instruction result in a weblike multiple interconnectedness on multiple dimensions that is not subject to the limitations of instruction that is characterized by a single organizational slant. Instead of a single text with a single organizational scheme and a single sequencing of comparisons and contrasts, our hypertexts allow the same information to be automatically reconfigured according to a huge number of possible organizational schemes, determined by using subsets of the multiple thematic coding space--our hypertexts enable the *virtually limitless automatic generation of new text configurations*. Because of the richness of ill-structured domains such as biomedical science, each of these text configurations teaches some case- (experience-) grounded lessons that would not have been taught (or easily seen if taught) from another text's organizational perspective. Such additional experiences and perspectives are always helpful in a complex domain--a physician never learns all that it would be helpful to learn (which is why additional experience is always valued in a physician). Hypertext systems like the *Cardioworld Explorer* systematically *consolidate the process of acquiring experience*.

Yet another virtue of the multiple interconnectedness along multiple dimensions of the representations that our systems build has to do with the problem of reciprocal misconception compounding that we have observed in our studies of medical students and physicians [6, 7]. Misconceptions bolster each other and combine to form seductively entrenched *networks of misconception*. Our approach helps to forestall the development of misconception networks by developing a kind of *positive reciprocation*. Because *correctly conceived representations with a high degree of multiple interconnectedness are established*, the fresh entry of fallacious knowledge at any node in the weblike network will fire off so many connections that it would be likely to activate some *misconception-disabling correct knowledge*. Before you can go too far wrong, you are likely to touch something that sets you right.

7. Active Participation, Tutorial Guidance, and Adjunct Support for the Management of Complexity. In an ill-structured domain, knowledge cannot just be handed to the learner. A priori codifications of knowledge are likely to misrepresent. (That is part of what ill-structuredness means.) Hence the importance, increasingly widely recognized today, of active learner involvement in knowledge acquisition, accompanied by opportunistic guidance by expert mentors (which can be incorporated in a computer program--it does not have to be live, one-to-one guidance). Furthermore, aids must be provided to help the learner manage the added complexity that comes with ill-structure. Our hypertext programs allow learners to explore complex conceptual landscapes in many directions, with expert guidance and various kinds of cognitive support (e.g., integrated visual displays). When there are limits to the *explicit* transmission of knowledge, learners will need special kinds of help in figuring things out for themselves. (see [1, 4, 12].)

Recapitulation: A Shift from Single to Multiple Representations and from Generic Schema Retrieval to Situation-Specific Knowledge Assembly

In general, we argue that the goals of advanced knowledge acquisition in complex and ill-structured domains can best be attained (and the problems we have identified avoided) by the development of mental

representations that support *cognitive flexibility*. Central to the cultivation of cognitive flexibility are approaches to learning, instruction, and knowledge representation that: (a) allow an important role for *multiple representations*; (b) view learning as the multidirectional and multiperspectival "criss-crossing" of cases and concepts that make up complex domains' "landscapes" (with resulting interconnectedness along multiple dimensions); and (c) foster the ability to assemble diverse knowledge sources to adaptively fit the needs of a particular knowledge application situation (rather than the search for a precompiled schema that fits the situation). We suggest that theory-based computer hypertext systems can implement the goals and strategies of Cognitive Flexibility Theory, engendering multiple cognitive representations that capture the real-world complexities of the kinds of cases to which abstract conceptual knowledge must be applied.

Acknowledgement

References

1. Barrows, H. S., & Tamblyn, R. M. (1980). Problem-based learning. New York: Springer.
2. Burstein, M. (1983). Concept formation by incremental analogical reasoning and debugging. Proceedings of the International Machine Learning Workshop, Champaign, IL.
3. Champagne, A. B., Gunstone, R. F., & Klopfer, L. E. (1985). Effecting changes in cognitive structures among physics students. In L. H. West & A. L. Pines (Eds.), Cognitive structure and conceptual change. Orlando, FL: Academic Press.
4. Collins, A., Brown, J. S., & Newman, S. E. (in press). Cognitive apprenticeship: Teaching the craft of reading, writing, and mathematics. In L. B. Resnick (Ed.), Cognition and instruction: Issues and agenda. Hillsdale, NJ: Erlbaum.
5. Collins, A., & Gentner, D. (in press). How people construct mental models. In S. Quinn & D. Holland (Eds.), Cultural models in thought and language. Cambridge: Cambridge University Press.
6. Coulson, R. L., Feltovich, P. J., & Spiro, R. J. (1986). Foundations of a misunderstanding of the ultrastructural basis of myocardial failure: A reciprocating network of oversimplifications (Tech. Rep. No. 1). Springfield: Southern Illinois University School of Medicine, Conceptual Knowledge Research Project.
7. Feltovich, P. J., Spiro, R. J., & Coulson, R. L. (in press). The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions. In D. Evans & V. Patel (Eds.), The cognitive sciences in medicine. Cambridge, MA: M.I.T. Press.
8. Schank, R. C. (1982). Dynamic memory. Cambridge: Cambridge University Press.
9. Spiro, R. J. (1980). Accommodative reconstruction in prose recall. Journal of Verbal Learning and Verbal Behavior, 19, 84-95.
10. Spiro, R. J., Feltovich, P. J., & Coulson, R. L. (in press). Multiple analogies for complex concepts: Antidotes for analogy-induced misconception in advanced knowledge acquisition. In S. Vosniadou & A. Ortony (Eds.), Similarity and analogical reasoning. Cambridge: Cambridge University Press.
11. Spiro, R. J., & Jehng, J. C. (in press). Random access instruction: Theory and technology for the nonlinear and multidimensional traversal of complex subject matter. In D. Nix & R. J. Spiro (Eds.), The Handy Project: New directions in multimedia instruction. Hillsdale, NJ: Erlbaum.
12. Spiro, R. J., Vispoel, W. L., Schmitz, J., Samarapungavan, A., & Boerger, A. (1987). Knowledge acquisition for application: Cognitive flexibility and transfer in complex content domains. In B. C. Britton & S. Glynn (Eds.), Executive control processes. Hillsdale, NJ: Erlbaum.
13. White, B. Y. (1984). Designing computer games to help physics students understand Newton's Law of Motion. Cognition and Instruction, 1, 69-108.
14. White, B. Y., & Frederiksen, J. R. (1987, November). Causal model progressions as a foundation for intelligent learning environments (BBN Report No. 6686). Cambridge, MA: Bolt Beranek and Newman, Inc.
15. Wittgenstein, L. (1953). Philosophical investigations. New York: Macmillan.