

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

DTIC FILE COPY

4

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
1a. REPORT SECURITY CLASSIFICATION Unclassified		1b. RESTRICTIVE MARKINGS --		3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release, distribution unlimited	
2a. AUTHOR 3 06 1988 ULE		2b. MONITORING ORGANIZATION REPORT NUMBER(S) --		5. MONITORING ORGANIZATION REPORT NUMBER(S) --	
6a. NAME OF PERFORMING ORGANIZATION Southern Illinois University School of Medicine		6b. OFFICE SYMBOL (If applicable) --		7a. NAME OF MONITORING ORGANIZATION Cognitive Science Program Office of Naval Research (Code 1142CS)	
6c. ADDRESS (City, State, and ZIP Code) Department of Medical Education P.O. Box 19230 Springfield, IL 62794-9230		7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy St. Arlington, VA 22217-5000		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-87-G-0165	
8a. NAME OF FUNDING / SPONSORING ORGANIZATION --		8b. OFFICE SYMBOL (If applicable) --		10. SOURCE OF FUNDING NUMBERS	
8c. ADDRESS (City, State, and ZIP Code) --		PROGRAM ELEMENT NO. 61153N		PROJECT NO. RR04206	
		TASK NO. RR4206-01		WORK UNIT ACCESSION NO. 4422546-01	
11. TITLE (Include Security Classification) Acquisition, Understanding, and Application of Biomedical Science Knowledge					
12. PERSONAL AUTHOR(S) Paul J. Feltoovich					
13a. TYPE OF REPORT Final		13b. TIME COVERED FROM 87MAY01 TO 88APR30		14. DATE OF REPORT (Year, Month, Day) 1988, December	
15. PAGE COUNT 41					
16. SUPPLEMENTARY NOTATION A related paper appears in the proceedings of the Tenth Annual Conference of the Cognitive Science Society, Hillsdale, NJ: Erlbaum, and is attached.					
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)			
FIELD	GROUP	SUB-GROUP			
05	08				
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
This is a final report on an equipment grant, to establish an artificial intelligence laboratory at Southern Illinois University School of Medicine. The nature of the established laboratory is described. Systems developed during the term of the grant to support research are described, including a hypertext system for studying instruction of complex, ill-structured subject matter; simulations for studying the application of complex conceptual knowledge; and a microworld for cardiovascular biomedicine. The latter supports further research on the development of misconceptions in complex domains, the stability of these misconceptions, and mechanisms by which faulty beliefs can be changed. computer systems; artificial intelligence; medical computer applications;					
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS				21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL Susan M. Chipman				22b. TELEPHONE (Include Area Code) 202/696-4318	
				22c. OFFICE SYMBOL 1142CS	

DD Form 1473, JUN 86

Previous editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE

Unclassified

89 2 6 017

FINAL REPORT

Acquisition, Understanding, and Application
of Biomedical Science Knowledge

Grant No. N00014-87-G-0165

1 May 87 - 30 April 88

Paul J. Feltovich (Principal Investigator)

Richard L. Coulson

Southern Illinois University School of Medicine

This is the final report for ONR Grant No. N00014-87-G-0165, NR4422546-01, Cognitive Science Program. Approved for public release; distribution unlimited. Reproduction in whole or part is permitted for any purpose of the United States Government.

DISCLAIMER NOTICE

**THIS DOCUMENT IS BEST QUALITY
PRACTICABLE. THE COPY FURNISHED
TO DTIC CONTAINED A SIGNIFICANT
NUMBER OF PAGES WHICH DO NOT
REPRODUCE LEGIBLY.**

Background

Although the effective starting date of Grant No. N00014-87-G-0165 to Southern Illinois University School of Medicine was 1 May 87, the process of purchasing and receiving the equipment from Xerox Artificial Intelligence Systems required a substantial amount of time, and the laboratory was not installed until 22 September 1988, leaving seven months for work within the one-year period of the Grant. Within this period of time, the presence of the laboratory has enabled considerable progress in several areas including: the development of intelligent computer assisted instruction systems and systems to serve as experimental environments for other Office of Naval Research sponsored programs of research; the development of significant and growing working relationships between the Artificial Intelligence Laboratory of the School of Medicine and its parent campus (Southern Illinois University) at Carbondale, Illinois, particularly with the Department of Computer Science; the availability a sophisticated Artificial Intelligence training site for students of both the School of Medicine and the greater university; and recruitment of an AI specialist to the Department of Medical Education of the School of Medicine (who also has formal ties to the Department of Computer Science). All of these consequences of the Grant from the Office of Naval Research are described in this Final Report.

A total of five Xerox 1186 AI Workstations were purchased under the grant, as well as file and print service to support these machines (Attachment 1 contains a listing of equipment purchased under the grant). Three of the workstations are housed in the School of Medicine in Springfield, in a 400 or square foot facility built for this purpose by the School. Two others are housed in the laboratory of collaborator Richard L. Coulson (Professor, Physiology) on the parent campus of Southern Illinois University in Carbondale. The primary reason for splitting the laboratory is that our

- 1 -



☒	
☐	
☐	
ty Codes	
Dist	
Special	
A-1	23

research in bio-medical conceptual understanding involves the basic biomedical sciences as well as their relationship to clinical aspects of medical care, and major parts of the medical school training involving the basic sciences takes place in Carbondale. A significant side-benefit of splitting the lab, with a major hub in Carbondale, has been the establishment of an outlet to the parent University campus, with its greater access to graduate students, to all departments of a major university, and so forth. This strategy has worked well, with considerable benefit to the goals and activities of the laboratory, as will be discussed as appropriate throughout the report.

AI Systems Developed Under the Grant: Cardioworld

Cardioworld is designed to be a rich investigation environment in which medical students can explore and learn about many aspects of cardiovascular biomedical science, cardiovascular medicine, and useful relationships between these. It also functions as a research environment for other programs of research. The development of Cardioworld builds off of a program of research on medical students' acquisition, understanding, and application of biomedical science knowledge, and on the stability of belief systems that has been conducted for five years, within the last year under contracts from the Office of Naval Research (Contracts N00014-88-K0286, N00014-88-K-0077). Much has been learned in this research about difficulties students have in learning and understanding difficult biomedical concepts, and about how they develop and maintain misconceptions. Much of this work has been in the cardiovascular area of medicine. One aim of the Cardioworld system is to take advantage of this knowledge, and modern artificial intelligence tools, to promote greater competence in cardiovascular medicine among students. However, the primary objective is to use the system to learn more about complex conceptual understanding and the stability of conceptual belief.

Because our research has had an emphasis on basic biomedical science (as well as its applications), some major aspects of the design of the Cardioworld system have reflected this focus. In particular, the main computer screen that functions as the students' top-level interaction with the whole system is a network representation of biomedical concepts important to understanding cardiovascular medicine. This integrated visual display serves many purposes. For example, all modes of interaction with the system (some under development and described below, and others to be built in the future) available to the student are attached as menus to concept nodes in this visually displayed network. The screen is also used more dynamically within some of the systems that can be called from these menus (as in the hypertext system described below). Whatever the student is doing within Cardioworld, the main network display of biomedical concepts (hereafter, referred to as the Top Level display) is always visible, to help students keep track of where they are working within the relevant biomedical science and to show how his/her current target area of conceptual focus relates to others.

From this main screen and the menus associated with conceptual nodes, the student is to be able to enter: tutors designed to promote better understanding of concepts and to investigate students misconceptions and their stability; micro-worlds for exploring the functionality of concepts; modules that teach conceptual-clinical relationships; exercises involving clinical "presenting complaints" that exemplify patterns of underlying biomedical science concepts; assessment modules focusing on conceptual understanding; and clinical patient simulations of cases that embody sets of cardiovascular concepts. Some of these are currently being designed or will be designed (in many instances, independently of ONR funding), while substantial progress has been made on others during the term of the grant. The latter are described below.

The Cardioworld Explorer System for Teaching Basic Science--Clinical Ties

This Cardioworld Explorer is a hypertext system designed to enable medical students to explore the complex relationships between cardiovascular biomedical science concepts and clinical cardiovascular medicine. Before medical professionals can utilize knowledge of the basic sciences to help them in working with clinical cases, they must know and recognize what areas of biomedical science are pertinent to cases. This is made more difficult in medicine because medical clinical cases represent an ill-structured domain. By this we mean that many concepts (interacting contextually) are pertinent in a typical case of knowledge application, and the patterns of combination of concepts are inconsistent across case applications of the same nominal type. This is true in many areas of practice (e.g., electronic trouble-shooting) where a body of basic science concepts is pertinent to instances of application. Hence, it is difficult, for example, for a practitioner to know what aspects of cardiovascular biomedical science are pertinent in a clinical case, because cases that present clinically in a similar fashion may involve different aspects of cardiovascular biomedical science. The contextual interactions of cardiovascular concepts across clinical presentations are the focus of the cardiovascular hypertext system. The system is implemented using a customized version of the Xerox product NoteCards, working in the Xerox Loops environment.

The foundation of the Cardioworld Explorer system is a library of clinical cardiovascular "mini-cases." Each mini-case is a short description of a clinical presentation--a case of cardiovascular disease, or a part of a case, contained on a NoteCard. Each mini-case is coded on a multi-dimensional vector containing the cardiovascular biomedical concepts (e.g., Starling's law, impedance) that are pertinent to interpreting the mini case and key

clinical descriptors (e.g., dizziness, shortness of breath) that are represented in the presentation of the mini-case.

A student can use the system in several ways. The student may simply browse the library of mini-cases. Whenever a mini-case is activated (selected for viewing) by the student, the mini-case clinical presentation itself is displayed, and simultaneously nodes in the Top Level display (the network display showing the set of cardiovascular basic science concepts and their interrelationships) are highlighted with two different shadings to show the basic science concepts that are pertinent to interpretations of the minicase. The two different shadings reflect two levels of relevance of the concepts to the case--highly pertinent and a lesser level of pertinence. At the same time, other screen windows are displayed in which explanations are given for why and how the highlighted concepts are relevant to the mini-case.

Because numerous concepts are likely to be relevant to any case and cases with similar presentations may involve different sets of concepts, no simple rules can be enumerated (or learned by students) to capture the relationships between concepts and cases. By enabling students to browse (in the ultimate implementation) hundreds of clinical vignettes, while concurrently seeing patterns of relevant biomedical concepts (as well as discussions of why they are relevant), a sense of the imperfect correspondences that do exist can be acquired by the student. Two additional modes of interaction are available to enable students to investigate nuances of these relationships more directly. Both of these are engaged from the NoteCard representing a mini-case.

From any mini-case (NoteCard) a student can search the case library for: other cases that involve similar basic science; other cases that involve similar clinical presentations; cases involving any arbitrary set of basic science concepts specified by the student; or cases involving any arbitrary set of clinical features specified by the student.

When a student has inspected a particular mini-case and investigated the basic science concepts that would be useful in its interpretation, the student may inspect other cases that engage a similar profile of cardiovascular concepts (each of which, in turn, provides guidance on why the concepts are applicable and how they are involved in producing the clinical signs and symptoms of the case). This is done by engaging a search procedure that looks in the library for cases having similar applicable basic science concepts, using several levels of matching--from an identical profile of concepts to lesser degrees of overlap. Inspection of the cases resulting from the search enables the student to learn the kinds of variability that can result in clinical situations where the same cardiovascular concepts are relevant. In addition, the student can specify an arbitrary profile of cardiovascular concepts, perhaps systematically varying those that are pertinent to the target (entry) case, and engaging searches that assemble mini-cases that match these constellations. This enables more user-directed interrogation of the clinical variation that can result in the presence of similar applicable biomedical concepts.

The student can also work in the other direction, focusing on clinical presentations and their variability and investigating concepts that are pertinent to their interpretation. The student can do this from any mini-case, initiating searches for cases with similar clinical presentation (again, at several levels of match) or by specifying arbitrary variations of clinical presentation to initiate the search. This enables the student to investigate the range of conceptual variability across similar case presentations.

The entire computational architecture for those aspects of the Cardio-world Explorer system that have been described so far is in place. Development work during the term of the grant has been done using a case library of thirty mini-cases, and further progress after the grant period will be directed at building up the mini-case library. An authoring system is in

place to make this easy for medical informants. Another extension under development involves the capability for a student to combine mini-cases (e.g., parts of total clinical cases) to explore changes in applicability of cardiovascular concepts as a case segment occurs in the changing context of others. This poses much more complicated design challenges than what has been done already, as it requires an active physiological model to underly the system. Work has progressed on such a physiological model (described in the section below on the "Impedance Tutor") and future efforts will involve integrating this model with the Cardioworld Explorer.

An additional discussion of the theoretical basis for the cardiovascular hypertext system, involving the goals and difficulties of advanced knowledge acquisition and the development of cognitive flexibility required for functioning effectively in ill-structured domains, is included as Attachment 2.

Clinical Case Simulations

The focus of this project is to implement a knowledge based tutoring system for preclinical medical students on the diagnosis and treatment of clinical patients. The system involves patient simulations that a student can diagnose and treat. As part of Cardioworld, such simulations are menu options that students can choose in order to interact with cardiovascular clinical cases. The tutoring system is able to gather knowledge from experts ("learn" in interaction with experts) and use this knowledge to guide students in their approach to medical problems presented as patient simulations.

The system provides a complete, accurate patient simulation, with which students can practice patient care and patient problem understanding. In practicing clinical investigation and treatment with simulated patients, students can judge their own personal strengths and weaknesses in the skills and knowledge necessary to understand and manage such patients and determine

exactly what they need to study or review. The system also provides an interface for experts which can dynamically learn and interactively transfer the experts' knowledge to the system's knowledge base, which then provides students with on-line help when they encounter difficulties in dealing with a patient problem.

The transfer of expert knowledge and the assembly of a complete knowledge base is a very important but difficult task. Thus, besides an interface for interactively inputting (production) rules provided by experts, the system also contains a learning mechanism to acquire expert rules dynamically, as experts interact with the patient problems. In this way, the knowledge base can be assembled dynamically and efficiently.

The system has three interfaces, one for word processing, one for experts and one for students. The word processor's interface is designed for the input of large amounts of data, including patient response information to history, physical, and laboratory inquiries; the expert's interface is designed for testing and modifying the system as well as for knowledge acquisition and transfer; and the student's interface is designed for students' interaction with the simulated patients, when the necessary patient data are entered and knowledge has been transferred to the system from the expert interface. The whole system is driven by menus. Except for inputting data, the user generally needs only to choose an item from the menu shown on the screen, using the mouse to proceed.

This particular version of the tutor is instantiated to medical problem solving, based on the process of clinical reasoning. It is, however, a generic tutor amenable to any problem solving discipline involving open-ended problems and the creation and testing of hypotheses. The most significant feature of this tutor is that as experts work with the problem simulations, the tutoring system learns the procedures and rationale of the experts, adding

to the knowledge base which can then be employed by the system when tutoring a novice. The next phase of this project, which will be conducted independently of ONR support, is to expand the problem simulation base, accumulate expert experience, and develop a natural language interface for inquiries to the "patient."

The Cardiovascular Impedance Microworld (Tutor)

A major focus of our research on students' understanding (misunderstanding) of cardiovascular concepts and on the stability of misconceptions that students acquire has involved the particularly important concept of cardiovascular impedance. This has to do with factors of the cardiovascular system which oppose the flow of blood and how the interactions of these produce their joint effects. Students have great difficulty understanding impedance, and they acquire significant inappropriate beliefs. (An understanding of this concept is important to the management of many clinical cardiovascular conditions, especially hypertension.)

Much of the difficulty students experience stems from aspects of the cardiovascular system which are different from ostensibly similar flow systems (e.g., household plumbing) with which they have more experience and which are conceptually more tractable. These aspects include the continuously changing (pulsatile) pressure produced by the heart (as opposed to a more constant pressure in plumbing) and the stretchiness of the blood vessels (as opposed to more rigid pipes). These factors introduce kinds of opposition to blood flow (compliant and inertial reactance) that are not (to any significant degree) present in the more familiar systems, that are highly abstract, and that combine to produce their effects in complex and nonintuitive ways. These combined effects constitute the complex concept of impedance.

Another project under development is an impedance tutor, designed to help students understand this important concept better, and to serve as an experimental environment for our investigations of the stability of misconceptions and mechanisms of belief change. The project takes advantage of our considerable laboratory work that has revealed the major impediments to understanding of this concept and the types of misconceptions students develop. At the present stage of development the system constitutes more a microworld than a tutor, but will serve as the underlying basis for extensions of the system into intelligent tutoring.

A microworld is a computer model of a system to be understood that represents components of the system and the lawful relationships among them, so that a student can explore, manipulate parameters, conduct experiments, etc., in this system and observe the consequences of actions and the accuracy of predictions. In this regard, the Impedance Microworld rests on a mathematical model of the hemodynamics of the cardiovascular system. This model takes account of most factors of the cardiovascular system that influence the circulation (e.g., pressure pulsations, flows, wall thicknesses of vessels, density and viscosity of blood) and can propagate the consequences of changes to any of these factors (locally in a section of the circulation, as in a change in resistance there, or more globally, as in a change in heart rate) throughout the circulation. This underlying model, capturing the "lawful" regularities among numerous "parameters" and underlying elements of the circulation, supports the exploration of these by the student.

The interface to the microworld is a schematic of the cardiovascular system that is displayed as a window graphic. All entry-level interactions with the system are via menu selections attached to components of this screen or by direct buttoning of screen-active components of the circulation (e.g., a particular vein).

Impedance is a complex concept, involving subconcepts of varying complexity, ranging from relatively simple anatomical concepts to intricate interactions among abstract concepts. The system contains a set of exploration tools that enable the student to explore dynamically all aspects of impedance--from the most basic to the most combinatorially complex--in interaction with the physiological cardiovascular model.

The exploration tools correspond to a structural analyses of the concept of impedance that we have created and that has been used to guide our laboratory work on students' understanding and difficulties with the concept. It captures fairly well the different kinds of difficulties students experience in understanding; hence, in our computer work it has been "turned around" as an organizational structure for the set of tools of exploration. The scheme partitions the concept of impedance (opposition to blood flow) into four levels of conceptual complexity (see Fig. 1).

Level 1. Components of the concept at this level are primitive elements and are typically concrete entities with physically realizable properties. Examples in Fig. 1 for Opposition are the radius of a vessel and the viscosity of blood.

Level 2. The second level involves first order relationships among properties of primitives, and conceptual components which emerge as a result of these relationships. At this level, for example, vascular compliance exists as a function of vessel wall thickness, radius, and elasticity.

Level 3. This level is fundamentally different in that it requires a mapping or transformation of conceptual elements from Level 2 into a new, abstract "space." This space provides a kind of comparison or "common denominator" environment which enables commalities or comparisons of Level 2 elements to be made which cannot otherwise be made. With regard to Fig. 1,

Fig. 1

**STRUCTURE OF A CONCEPT:
"OPPOSITION TO BLOOD FLOW -- IMPEDANCE"**

LEVEL 1: Primitives

Radius (r)
Length (l)
Viscosity (η)

Radius
Length
Elasticity (ϵ)
Wall Thickness (s)

Radius
Length
Density (ρ)

LEVEL 2: First Order Relationships

Resistance

$$R = \frac{8 \cdot \eta \cdot l}{\pi \cdot r^4}$$

Compliance

$$C = \frac{\pi \cdot r^3}{\epsilon \cdot s \cdot l}$$

Inertance

$$L = \frac{\rho \cdot l}{\pi \cdot r^2}$$

LEVEL 3: Space Transformation -- Frequency (f)

Resistance

$$R = \frac{8 \cdot \eta \cdot l}{\pi \cdot r^4}$$

Compliant Reactance


$$X_C = 1/2 \cdot \pi \cdot f \cdot C$$

Inertial Reactance

$$X_L = 2 \cdot \pi \cdot f \cdot L$$

LEVEL 4: Relationships in Transformational Space; Vectorial

Impedance (Z)



$$(X_L + X_C) = 1 + (-4) = -3$$

$$X_R = R = 4$$

$$Z = \sqrt{(X_L + X_C)^2 + R^2} = 5$$

resistance, compliance, and inertance are not comparable as oppositional factors at Level 2. However, when these concept components are all mapped to a space involving the dimension of frequency (of the beating of the heart), new conceptual components emerge which are directly comparable. That is, transformed elements of compliant reactance and inertial reactance emerge which act equivalently to resistance as oppositional factors to the flow of blood. But, they must be thought of in relation to frequencies of a pulsatile, beating heart.

Level 4. In the comparator space of Level 3, new conceptual components emerge which can be compared and related. Yet, once comparable, they still combine in complex ways at Level 4. In particular, the reactances and resistance of Level 3 combine at Level 4 to form the overall concept of interest in the analysis. This is the concept of impedance or the total opposition to the flow of blood in the cardiovascular system. The concept of impedance requires the reactances and resistance to be considered in vectorial, rather than scalar combination.

The tools or modes of exploration associated with this concept analysis that are available to the student and the ways they can be utilized within the system are described below.

The student can manipulate any of the parameters in Fig. 1 which are then propagated by the cardiovascular physiological model, enabling the student to observe the consequences of the manipulations on the circulation. A number of modules are available as tools to aid students in these explorations, and these are described below. The modules function in several ways within the system: 1) they are engaged by the system itself as part of explanation and demonstration of the consequences of a manipulation; 2) they can be used by the student as tools for conducting advanced (or more "sophisticated") manipulations of the physiological model; or 3) they can be engaged as separate

exploration domains, in which a student can explore basic concepts of cardiovascular physiology or mathematical/analytic tools pertinent to understanding these.

- Anatomical dependence module

One such module involves the exploration of resistance, compliance, and inertance; the basic vascular parameters. The interface for this module is a screen display showing the numerical equations for these parameters as a function of their more basic elements (see Fig. 1) and a screen graphic illustration of a segment of cylindrical vascular conduit. The equations contain the symbols for more basic elements which can be manipulated, and the anatomical graphic shows visually the physical effects on a segment of vessel of any manipulations. If the radius of the vessel is increased, for example, the illustration will change to show a conduit with larger radius but with wall thickness decreased by an amount commensurate with conservation of mass of the wall of the vascular element. The same change in radius will be propagated to the other parameters (e.g., compliance or inertance) which also have radius or wall thickness as elements. Since all three parameters, resistance, compliance, and inertance have radius as an element, the values of all change with an alteration in radius, and the student can investigate these interactions. The anatomical module can be used separately or in conjunction with the "Critical Issues Module," which may receive inputs from or deliver inputs to the anatomy module.

- Critical issues module

In the interface for this module, viscosity, elasticity, and density (which are critical issues for resistance, compliance and inertance respectively) are isolated as dependent variables in a graphical presentation, along with their related independent variables which control them. For example, resistance depends upon the "critical issue" viscosity which, in the vascular

system, is itself dependent upon the velocity of blood flow. By exploring a parameter (e.g., resistance) as a function of its related critical issues, it can be seen, for example, that an alteration in radius which would effect the resistance may also affect the rate of flow of blood and, hence, the velocity of flow, which, in turn, is an independent variable controlling viscosity which may change and cause further change in resistance. Similarly, compliance depends upon a critical issue elasticity. Elasticity, as a stress-strain relationship of the vessel wall, is sensitive to radius via vessel circumference and wall thickness. The critical issue for inertance is density of the blood which varies with the water, cellular, and protein content of the blood. Factors which affect the density of the blood may also effect the viscosity, which in turn will feed back upon the resistance parameter by virtue of the viscosity element.

- Frequency analysis module

This module allows interactive exploration of the principal of decomposition of complex wave forms into component sinusoidal single frequencies. This module is particularly important because our research has shown that the ramifications of oscillation and pulsation in cardiovascular pressure and flow are poorly understood by students, and lead to many of their misconceptions.

The interface for this module is a screen graphic that enables the student to specify arbitrary functions of time for decomposition analysis. The user can "draw" an arbitrary wave form with the mouse or choose from a variety of default wave-forms contained in a library (e.g., musical chord or a pressure pulse). The wave-form specified is then frequency analysed using a fast Fourier transform, and the component sinusoids are displayed as functions of time, showing their respective amplitudes and phase time shifts. If provided

with separate pressure and flow wave-forms, this module also computes a complex impedance function using a pressure and flow frequency decomposition ratio.

This module can be used independently on user specified (or library selected) wave-forms or in conjunction with the physiological model. In this latter mode, the module can either receive blood pressure and flow wave-forms of the cardiovascular circulation from the physiological model for analysis, or the student can specify these as input to the physiological model and observe the consequences. This module, together with the anatomical dependence and critical issues modules, can exchange attributes among themselves or with a fourth module.

- Vector analysis module

The interface for this module is a screen graphic representing a vector space in which students can explore how the parameters of compliance and inertance are transformed into the impedance space of reactance by the action of wave form component frequencies. For this purpose the frequency information can be obtained from the frequency analysis module or specified by the user. The display shows how the reactances of compliant and inertial origin, having opposite numerical signs, sum together in opposing fashion on the ordinal axis of the graphical display (see Fig. 1). It further illustrates the notion of vector space summation by presenting the resistance value (obtained, for example, from the anatomical dependence module) as the abscissal value on the same display. Graphical display of the resultant impedance, obtained from combining in vector space the mathematically imaginary compliant and inertial reactance data with the mathematically real resistance data, is illustrated while the module performs the underlying numerical computations. This module, in conjunction with the physiological model, will synthesize the vector space

impedance calculation for all wave form component operational frequencies in the physiological model. In addition, it reconstructs the complex impedance function, which can be shown to be identical to that computed by the frequency analysis module using the pressure and flow frequency decomposition ratio. The model can also be used independently to explore the mathematics of vector space operation.

These four modules operate as side-effects of the operation of the physiological circulation model, but they also serve as exploratory tools for investigating the operation of the circulation model itself as well as for understanding the functions that the modules themselves perform. The modules can be disengaged from the physiological model so that a user can interact with them in a separate instructional mode. In separate exploratory (instructional) mode, attributes derived from module operation are not required to maintain internal consistency with the full operational physiological model. However, attribute consistency within the module itself is maintained in this instructional mode.

With the functional model of physiology in place, along with interactive tools for exploring and manipulating all aspects of the concept of impedance, the ground work is in place for extension of the microworld to interactive tutoring. Under design are two major new "modules" for the system. The first is an assessment/diagnostic module, which, by leading the student through a set of exercises involving manipulations of parameters and predictions of effects, can be used to diagnose misconceptions regarding impedance the student may hold. Development of the diagnostic module will benefit greatly from our prior research on types of student misconceptions. The second module is the interactive tutor itself. This module will interact with the assessment/diagnostic module to generate tailored instructional exercises and interactions for the student.

The Cardiovascular microworld and tutoring systems will have instructional benefit for students. However, our main interest in support of our research programs is that the systems will provide a computerized environment for conducting experiments on students misconceptions, the stability of these misconceptions under instructional challenges, and the interactive means by which mistaken beliefs can be changed. This computer environment will complement other laboratory studies of these same topics.

Institutional Development

In addition to the projects that are under development as a direct consequence of the presence of the laboratory, a number of additional developments have accrued, some of which will have substantial impact on the future productivity of the AI laboratory. (These are side-benefits of the lab, which have developed independently of direct ONR support. They are described to give a sense of extended benefits that have accrued from the presence of the AI laboratory.)

- The medical school and the parent university

Aspects of institutional development include: 1) projects conducted by School of Medicine faculty other than core members of the laboratory (independently of ONR funding), 2) enhanced curricular offerings, 3) a training site for medical students interested in AI, and 4) development of working relationships with the Department of Computer Science.

Projects are under development (with no cost to ONR) by faculty of the School of Medicine, other than core members of the laboratory. These include a tutor for important concepts of biostatistics and a curricular map (using the Xerox product NoteCards) for mapping the hundreds of concepts and their interrelationships involved in the pharmacology curriculum.

As a direct result of the presence of the AI laboratory, courses in the School of Medicine have been given in "Artificial Intelligence in Medicine" and "Common LISP", neither of which had been available before.

Medical students are increasingly attaching themselves to the laboratory, to engage in projects or to take part in the work of the laboratory. A particular example is a medical student who was associated with the lab throughout his final year of medical school and who was a major collaborator in the development of the cardiovascular hypertext system, during the term of the laboratory grant. This student is planning to pursue his work in artificial intelligence during his residency.

In addition to its effects on the School of Medicine, the laboratory has facilitated expanded working relationships between the School of Medicine and the Department of Computer Science of Southern Illinois University. The workstations in the Carbondale laboratory are available for use by selected students of the Department of Computer Science, and several have subsequently become engaged in the work of the lab. In addition, working collaborations are developing between faculty of the School of Medicine associated with the AI laboratory and faculty of the Department of Computer Science.

- Recruitment

The laboratory has been instrumental in attracting a new faculty member to the Department of Medical Education of the School of Medicine. This person is a computer scientist and artificial intelligence specialist, and comes to the School of Medicine from Xerox Artificial Intelligence Systems. His specialty within AI is truth and belief maintenance systems, and his research will complement the research of the lab involving students' acquisition and maintenance of misconceptions, and mechanisms for changing faulty belief.

- Potential for expanded programs

As a result of the communications and cooperation that the award has fostered between the Medical Education and Physiology Departments of the School of Medicine and the University Computer Science Department, preliminary discussion has begun on the development of a graduate Ph.D. program in the interdisciplinary areas of artificial intelligence and cognitive science. If fully developed, this program would provide a continuous infusion of new young talent into the various programs of the AI laboratory, as well as provide a basis for "commodity" (graduate and post-doctoral student) exchange with other established labs nationwide. This program is being explored as a joint program involving the cooperative efforts of the Departments of Physiology and Medical Education from the School of Medicine and the Departments of Computer Science, Linguistics and Psychology from the College of Liberal Arts.

Relationships to other ONR and DOD Supported Research

The research conducted in the artificial intelligence laboratory supports the goals of the other ONR research contracts acquired since the acquisition of the laboratory. The cardiovascular hypertext system to some extent, and the impedance tutor more importantly and directly, are involved in the research conducted under Contract No. N00014-88-K-0077. This research is investigating factors that contribute to the acquisition and stability of misconceptions, and modes of instruction for changing erroneous beliefs. The hypertext system, the impedance tutor and the clinical simulations all support goals of Contract No. N00014-88-K-0286. Work conducted under this contract is investigating areas of biomedical understanding and medical practice that are particularly difficult for the Independent Duty Submarine Corpsman to master, and modes of training and instruction that can help the corpsman with the difficult aspects of his job.

7319m/1-21
12/22/88

QUANTITY	DESCRIPTION	AMOUNT
1	8037C-300 RECONDITIONED COMBINATION FILE AND PRINT SERVER	26,910.00
1	009R80-307 300 MB SPARE DISK PACK	1,350.00
4	1186-103 AI WORKSTATION	46,432.00
4	073S80-509 NETWORK COMPONENTS- TRANSCIEVER	800.00
4	152S24-014 DROP CABLE FOR 1186	700.00
1	G56 1024KB MEMORY KIT FOR SERVER	1,470.00
1	T86 512KB SERVER MEMORY EXP. KIT	350.00
1	E30 RS232 COMMUNICATION KIT	150.00
1	F88 INTERACTIVE TERMINAL SERVICE	315.00
2	77D IBM PC EMULATORS	1,500.00
2	TBD IBM PC EMULATION ENABLING S/W	320.00
4	012R80538 LOOPS SOFTWARE FOR 1186	13,000.00
1	1186-103 AI WORKSTATION	11,718.00
1	4045-203 LASAR PRINTER	4,050.00
1	4045-403 LASAR PRINTER MEMORY EXPANSION	400.00
1	152S23990 ETHERNET CABLE	460.00
2	113P90648 ETHERNET TERMINATOR	40.00
1	073S80509 TRANSCIEVER	200.00
1	152S24014 DROP CABLE	175.00
1	012R80531 XEROX QUINTUS PROLOG	2,000.00
1	NOTECARDS FREE FORM DATA BASE	2,000.00
2	LISP USERS PACKAGE	500.00
1	300 FT. COMPUTER CABLE	185.22
4	HAYES 1200/2400 BAUD MODEM AND CABLE	1,984.00
1	XEROX MODERN FONT FAMILY	210.00
1	XEROX CLASIC MATH FONT FAMILY	210.00
1	61420 XEROX HELVETICA FONT FAMILY	695.00

1 F04 MODERN FONT FAMILY

~~490.00~~ ← 490.00

1 F03 MATH CLASSIC FONT FAMILY

245.00

1 F02 PRINTWHEEL FONT FAMILY

195.00

TOTAL EQUIPMENT PURCHASED

119,054.22

ATTACHMENT II

**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

In: Tenth Annual Conference of the Cognitive Science Society. Hillsdale, NJ: Erlbaum, 1988.

**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 view"* 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 profferred 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*, recombining 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 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.

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 Handv 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.

Acknowledgements

The research reported in this paper was supported by grants from the Army Research Institute (MDA903-86-K-0443), the Office of Naval Research (N00014-88-K-0286, N00014-87-G-0165, N00014-88-K-0077), the Josiah Macy Foundation (B852001), and the Office of Educational Research and Improvement (OEG 0087-C1001). The authors would like to acknowledge the contributions of Jane Adami and Joan Feltovich to various aspects of this research.