Tech., Inst., Cognition and Learning, Vol. 5, pp. 91-97
Reprints available directly from the publisher
Photocopying permitted by license only

Introduction to Knowledge Representation, Construction Methods, Associated Theories and Implications for Advanced Tutoring/Learning Systems

JOSEPH M. SCANDURA

MERGE Research Institute, USA

Knowledge representations (KRs), methods for constructing them and associated theories provide the infrastructure on which all advanced tutoring/learning systems are based. There are, however, important differences of opinion on how to represent knowledge, not to mention how to construct such representations and/or the kinds of theories to use in building tutoring and other learning systems.

Scientific theories typically begin with frameworks and/or taxonomies. Examples range from the Greek classic distinction between water, fire, earth, air and ideas through the plant and animal kingdoms in biology, the periodic table in chemistry to Bloom's taxonomy and Gagne's conditions of learning in education. As science matures, emphasis turns increasingly toward fundamental mechanisms: Galileo's experimental method, Newtonian theory, the role of genes, proteins and the double helix in modern biochemistry. In education, the closest we have come are production systems, Logo, LISP, Prolog, XML and even AuthorWare to model teaching and learning processes.

Technology, Instruction, Cognition & Learning (TICL) poses a dilemma in this respect. Top-down approaches center on theoretical frameworks, models and/or taxonomies. Experimental research typically involves informally defined variables (e.g., Simple vs. Complex, Verbal vs. Spatial, Procedural vs. Declarative, Expert vs. Neophyte, Facts vs. Concepts vs. Rules vs. High Level Meta-Skills, Problem Typologies, Kinds of Cognitive Models).

These approaches have been popular because they generally are easy to understand, and they provide motivation and general guidelines for educational development. They have important limitations, however, when used to develop adaptive and other instructional systems. The multiple and varied decisions programmers inevitably make during implementation often result in systems that have only indirect relationship to the original concepts. As noted at a TICL meeting by Norbert Seel some years ago, it is often difficult to distinguish one framework from another by looking at the systems themselves.

Bottom up approaches often have the opposite problem—ensuring that the instructional systems created accurately reflect real world learning and/or instructional processes. Characteristics of the formalisms (or programming languages) used often impose undo constraints on the learning systems developed. It has been difficult to build learning and/or tutoring systems based solely on such formalisms. Complications involved in modeling the learner's actual (e.g., biologically based) thought processes also have been substantial. Benefits in both cases have been debatable—among other reasons because the use of such systems typically requires significant changes in the way education is conducted (cf., recent study by the U.S. Department of Education covering the 2004-2005 school year & response at the Carnegie Learning web site).

Comparison of alternative approaches, including reassessment of fundamental assumptions, seems long overdue. Indeed, future progress in TICL depends on a deeper understanding of foundational issues. Similarities and differences, strengths and limitations of various KRs, construction methods and associated theory need to be documented, compared and evaluated as they pertain to instructional systems. Future progress in TICL will be both constrained and hard to assess until clear consensus is reached on these issues.

This special publication on Knowledge Representation was planned with these concerns in mind. The focus is not just on alternative KRs, methods for constructing them or associated theories, but on how and how well each satisfies important requirements for building advanced tutoring and/or learning systems. To accomplish these goals, leading investigators representing a wide range of approaches to foundational issues were invited to contribute, initially to a series of symposia at annual TICL meetings over period of two years. Among those who contributed are: David Jonassen, Doug Lenat, Ken Koedinger, Wally Wulfeck, John Durnin, Richard Schmid, Russell Almond, Gilbert Paquette, Ron Stevens, John F. Sowa, Joseph M. Scandura, Bob Mislevy, Russell Almond, Stellan Ohlsson, Tanja Mitrovic, Val Shute, Ron Stevens, Neil Heffernan and Janet Kolodner.

All major KR formalisms used in advanced tutoring/learning systems were represented: Production Systems, Constraints, Relational Networks, Conceptual Graphs, Procedures/Flow Charts/Directed Graphs, Hierarchies/Abstract Syntax Trees, Distributed Connectionist/Artificial Neural Net models, Case Based Reasoning and Logic Oriented languages.

These presentations were followed by numerous informal communications. Jonassen, Heffernan, Ohlsson/Mitrovic, Paquette and Scandura prepared position papers, with the understanding that they would serve as the basis for subsequent published commentary and dialogue. Sowa also contributed a short overview of concept mapping.

While given wide latitude, each author was asked to focus on:

A. *Knowledge Representation*: Explicit description of the knowledge representation formalism, including examples and relationships to other KRs to make the formalism clear. To what extent is the KR executable on a computer? How does the KR represent declarative knowledge, procedural knowledge, higher order knowledge (e.g., meta-knowledge, strategies, logical reasoning) and/or levels of expertise?

B. *Method(s) for Constructing KR*: An explicit description of the preferred method(s) for constructing such KRs, including examples and relationships to other KRs. How would you characterize the method: Knowledge engineering? Cognitive/structural task analysis? Relational analysis? Other? To what extent is the method: Informal? Systematic? Automated?

C. *Associated Theory*: Summary of key features of theory(ies) associated and/or derived from knowledge representation.

1. Specify learning, problem solving, motivation and automation mechanisms.
2. Specify methods used for knowledge assessment (individual differences).

D. *Instructional Systems*: Implications of the chosen knowledge representation, construction method(s), theory and/or instructional systems. Evaluate practical relevance. Scale/complexity/costs.

E. Summary Sections: Summarize key ideas, relationships to other KRs, construction methods, limitations and benefits, distinguishing features of KR and associated construction methods. What are the current challenges and opportunities?

The position papers fall into two basic categories. Some contributors, specifically Ohlsson/Mitrovic, Paquette and Scandura, attempt relatively complete coverage of deep infrastructure and constitute Issue Number 2 in Volume 5. These articles focus, respectively, on variants on production systems (constraints), relational networks and hierarchies (abstract syntax trees), along with associated theories and implications for instructional systems. To facilitate comparison of fundamental, yet often subtle issues, each author has listed the most directly relevant pages in the following Table:

	Ohlsson & Mitrovic Production Systems	Paquette Relational Networks	Scandura Abstract Syntax Trees
A. Knowledge Representation	102-12, 117-25	135-39	173-5, 181-94
B. Method(s) for Constructing KR	103-6	139-53	175, 194-216
C. Associated Theory		105-39	
1. Learning et al mechanisms	107-11		175-7, 216-31
2. Knowledge assessment methods	104-7		177-8, 231-41, 192
D. Instructional Systems	112-5	153-61	178-80, 241-51
E. Summary Sections	125-8	161-5	170-3, 193, 214-9, 225, 230-1, 239-41, 251-68

Ohlsson and Mitrovic introduce constraints in an attempt to address in what they see as limitations in ITS based on production systems, a biologically inspired representation starting from the bottom up. They argue that their constraint-based approach offers greater cognitive fidelity (of underlying mechanisms) and efficiency than various versions of ACT theory traditionally used to guide Intelligent Tutoring Systems (ITS) development and delivery. They question the relevance of buggy productions, for example, arguing that (even) CMU researchers are implicitly ignoring them in recently developed ITS. Even if one accepts Ohlsson and Mitrovic analysis, however, important questions remain: Do other production system (e.g., ACT) proponents agree? More generally, do ITS systems based on production systems, constraint based or otherwise, address all essentials educators consider important?

Paquette begins with high-level assumptions that have guided most educational theorizing for decades. With these as a starting point, he systematically introduces an extended set of (mostly) relational formalisms needed to build functioning instructional systems. The resulting formalism is both sophisticated and complex. These formalisms have been used successfully in building web-based systems. Among the major questions here are: How do these formalism fit with alternative top-down approaches – such as those proposed by Jonasssen? Sowa? More fundamentally, is the resulting edifice is the best and/or most efficient way to address deep structure theoretical concerns and/or to implement the complexities inherent in operational instructional systems?

Scandura starts with hierarchical ASTs and a fundamentally different set of theoretical assumptions. He attempts to show how ASTs and these assumptions naturally lead to a comprehensive, yet parsimonious and rigorous, instructional (SLT) theory whose basic assumptions, like ACT in ITS, have been subjected to extensive experimental test and used to develop adaptive and configurable tutors. Aside from highly structured domains, however, can the proposed method of Structural Analysis be used successfully to build advanced tutoring systems more generally? Do ASTs accommodate instructional phenomena as completely as Paquette's more pragmatic approach? Does SLT address cognitive essentials as well as theories associated with ITS (e.g., ACT, constraint-based modeling)? If so, what are the relative limitations and/or benefits with respect to building instructional systems?

A major goal of ensuing "deep structure" dialog will be to help clarify such issues.

The other articles are more focused on application, high-level conceptualization and an overview of concept mapping. In this sense, they provide instructive counterpoints, collectively constituting Volume 5, Issue Number 3.

Heffernan et al have been developing an ASSISTment System, implicitly based on production systems, to help provide fine-grained transfer assistance to learners. They argue that tests commonly used to assess progress frequently involve multiple kinds of knowledge (e.g., a geometry problem may also require a combination of algebra, number sense and measurement) – each being important. Theirs is a web-based system designed to assess math skills and to provide tutoring ASSISTments when students give wrong answers. Knowledge Components (KC), roughly corresponding to productions (including buggy ones) were associated with test items – with transfer presumed with problems requiring similar KCs. They found that the finest grain models were not necessarily the best. A key question: Could future research benefit from more systematic analysis, such as those detailed in the deep structure expositions above?

Jonassen takes a top-down, partly philosophical approach to knowledge representation. He questions the sufficiency of formal systems (ontologies), preferring Case-Based Reasoning (CBR), with focus on kinds of knowledge needed for thinking, explaining, understanding and remembering. He argues that usable knowledge is based on human experiences conveyed as stories. His method starts by showing domain experts a problem, asking them to recall solutions to similar problems, and indexing and making cases available to learners. He adds that complex, ill-structured problems often have multiple solutions, and concludes with a description of two projects building case libraries of technology integration stories. Two important questions: How exactly do these CBR methods differ from those used with deep infrastructure (formal) KRs detailed in Issue 2? Could the latter potentially add rigor without destroying CBR pragmatics?

Sowa's short summary outlines concept mapping. He argues that knowledge representation involves analyzing a given body of knowledge, identifying relevant concepts, relations and assumptions, and the translation of results into notation that can be processed by computer. Sowa identifies four notations for use at various stages of knowledge acquisition, analysis and representation: (1) informal Concept Maps, (2) semi-formalized Topic Maps, (3) formal Conceptual Graphs, and (4) formal, but highly readable Common Logic Controlled English (CLCE). The question remains: What implications do these ideas have for cognitive theory and/or instruction?

These three articles share much in common with contemporary TICL research on cognitive modeling. Accordingly, this special, triple issue of TICL concludes with a special issue on Cognitive Modeling compiled by Val Shute.

Val Shute and I served as co-chairs in the original planning for a series of symposia at TICL 2006 and 2007. I was primarily responsible for Knowledge Representation et al and Val, for Cognitive Modeling. Issue Number 4 in Volume 5 deals specifically with the latter topic, and was compiled and edited by Val Shute. Her introduction appears at the beginning of Issue 4 herein.

Although each has its own unique characteristics, most of the articles in Issues 3 and 4 generally take a top-down approach to representation, cognition and instruction – with a substantial gap between motivating theory and working systems. It will be interesting to draw parallels, essential differences and hopefully light that will positively impact and help direct future TICL research.

With these articles as points of reference, the authors and other contributors to the TICL symposia on knowledge representation and cognitive modeling have all been invited to engage in wide ranging commentary and dialog to be published in a subsequent issue of TICL. In addition to cross comparisons, this dialog is likely to follow two tracts: One focus will be on deep infrastructure. The other will focus primarily on conceptual, experimental and application issues in Cognitive Modeling.

Dialog on deep infrastructure has already begun as this issue goes to press, the first being a comment by Koedinger and a reply by Ohlsson on fundamental assumptions in (production system) ACT theory that most would otherwise miss. Don't miss what is coming!