Beyond the Repertory Grid: New Approaches to Constructivist Knowledge Acquisition Tool Development

Jeffrey M. Bradshaw
Computer Science, Research and Technology, Boeing Computer Services
Seattle, Washington 98124 USA

Kenneth M. Ford
Institute for Human & Machine Cognition, University of West Florida,
Pensacola, Florida 32514 USA

Jack R. Adams-Webber
Department of Psychology, Brock University, St. Catharines, Ontario, Canada L2S 3A1

John H. Boose
Computer Science, Research and Technology, Boeing Computer Services
Seattle, Washington 98124 USA

Abstract

Personal construct theory provides both a plausible theoretical foundation for knowledge acquisition and a practical approach to modeling. Yet, only a fraction of the ideas latent in this theory have been tapped. Recently, several researchers have been taking a second look at the theory, to discover new ways that it can shed light on the foundations and practice of knowledge acquisition. These efforts have led to the development of three "second-generation" constructivist knowledge acquisition systems: DDUCKS, ICONKAT, and KSSn/KRS. These tools extend repertory grid techniques in various ways and integrate them with tools springing from complementary perspectives. New understandings of relationships between personal construct theory, assimilation theory, logic, semantic networks, and decision analysis have formed the underpinnings of these systems. Theoretical progress has fostered practical development in system architecture, analysis and induction techniques, and group use of knowledge acquisition tools.

To appear in a special knowledge acquisition issue of the *International Journal of Intelligent Systems,* January and February, 1993; also in K. Ford and J.M. Bradshaw (Eds.), *Knowledge Acquisition as Modeling*. New York: John Wiley, in press.

1. Introduction: Personal Construct Theory and Knowledge Acquisition

Personal construct theory (see Refs. 1-10) has provided both a plausible theoretical foundation and an effective practical approach to knowledge acquisition in a variety of settings. In particular, efforts to apply *repertory grid* techniques to knowledge acquisition have met with a great deal of success. In fact, personal construct theory and repertory grids have become so widely known and used that people often equate them. Yet despite the high level of research activity, only a fraction of the ideas latent in personal construct theory have been tapped. Recently, several researchers have been taking a second look at the theory, to discover new ways that it can shed light on the foundations and practice of knowledge acquisition. These efforts have led to the development of a new generation of personal-construct-based knowledge acquisition tools. While not discarding repertory grid techniques and representations, these tools extend them in various ways and integrate them with tools springing from other complementary perspectives.

In this paper we review past contributions of personal construct theory and summarize new directions. Section one discusses personal construct theory and repertory grids. Section two discusses new understandings of relationships with complementary theoretical perspectives underpinning constructivist approaches to knowledge acquisition. Section three examines three "second generation" knowledge acquisition tools that have benefited from these theoretical developments.

1.1. Constructive Alternativism and the Foundational Role of Distinctions

Personal construct theory is based on the research of George Kelly.[1-3] The theory's fundamental postulate and its eleven corollaries were derived from a single epistemological premise, that of *constructive alternativism*. According to this principle, 'reality' does not reveal itself to us directly, but rather is subject to as many different constructions as we are able to invent.[4-7] Thus, any event is open to a variety of different interpretations. This does not imply, however, that one interpretation of that event is as good as any other. On the contrary, different ways of construing the same event can be evaluated in term of their relative predictive utility.[8] That is, some interpretations of an event may prove more useful than others for anticipating similar events in the future.

1.2. Personal Constructs

Kelly defined his notion of *personal construct* as follows (Ref. 1, p. 61):

> In its minimum context a construct is a way in which at least two elements are similar and contrast with a third.

Thus, a construct simultaneously differentiates and integrates. To *construe* is both to abstract from past events, and to provide a reference axis for anticipating future events based on that abstraction. The process of construal thus lays the ground for all subsequent logical and mathematical reasoning: (Ref. 1, p. 278):

> The statistics of probability are based upon the concept of replicated events. And, of course, they are also contrived to measure the predictability of further replications of the events. The two factors from which predictions are made are the number of replications already observed and the amount of similarity which can be abstracted among the replications. The latter factor involves some complicated logical problems—for example, representative sampling—and, in practice, it is the one which usually makes predictions go awry. Since the abstractive judgment of what it is that has been replicated is the basis for measuring the amount of similarity, we find that the concept-formation task which precedes the statistical manipulation is basic to any conclusions one reaches by mathematical logic.

Though few would disagree with Kelly's observation, in practice designers of knowledge acquisition tools have given little attention to supporting the preliminary conceptual aspects of modeling that Kelly identifies as so crucial.

1.3. Persons-as-Scientists

Kelly's theory provides a rich characterization of the efforts of individuals to actively anticipate and control their environment. He draws explicit parallels between the processes that guide scientific research and those involved in everyday activities. His notion of *personal scientist* assumes that all people actively seek to predict and control events by forming relevant hypotheses, and then testing them against their experience.[10] In Kelly's own words (Ref. 1, p. 43), "the aspirations of the scientist are essentially the aspirations of all men." As Einstein put it (Ref. 11, p. 763), "The whole of science is nothing more than a refinement of everyday thinking."

In Kelly's view, humans model their environment and scientists model humans through a like process of *simulation* (Ref. 12, pp. 225-226):

> I think truth can be approached by simulation and by simulation only… Man gets at the truth of things… by erecting constructs to simulate it the best he can… [And scientists] devise machines to simulate—not man directly—but theories about man… the theories, in turn, are constructed to simulate the human processes they are supposed to explain. But the simulation does not stop there. The persons themselves are simulators. They attempt to simulate each other—too much, some say. They simulate their parents, their gods, a presumed rational way of life, and the expectations of others. In fact, a lot of people even make a big to-do about simulating themselves. This is known as 'trying to be yourself' and is often regarded as quite an accomplishment. Sometimes people simulate machines. This is sometimes called 'being objective.' [One scientist] has even programmed his people to behave like computers. Some psychologists undoubtedly will take this to mean that he has succeeded in getting people to behave psychologically.

In Kelly's view, a major goal of both individuals and social systems is anticipation. We simulate to improve the 'accuracy' of our anticipation of aspects of the future that are important to us. Action is a form of active anticipation that seeks to make desirable outcomes more likely.

1.4. Fundamental Postulate and Corollaries of the Theory

Kelly's fundamental postulate asserts that (Ref. 1, p. 46), "A person's processes are psychologically channelized by the ways in which he anticipates events." Hence, for Kelly, all our representational processes are essentially anticipatory.[4] He elaborated the logical implications of this proposition in terms of eleven corollaries, five of which are directly relevant to this paper.

Dichotomy Corollary: "A person's construction system is composed of a finite number of dichotomous constructs" (Ref. 1, p. 59). Kelly believed that the dichotomous structure of personal constructs is an essential feature of the way in which people organize information. For example, if a person simultaneously perceived an event to be equally pleasant and unpleasant in the same respect, then this distinction would be meaningless for that event.

Construction Corollary: "A person anticipates events by construing their replications" (Ref. 1, p. 50). Each person employs constructs to forecast events, and later to evaluate the predictive utility of those forecasts. Although the same event obviously never recurs, we use our personal constructs to organize perceived similarities and differences among events into coherent patterns or 'schemata.' Using these schemata as 'templates' we detect recurrent themes in our experience over time and feed these representations forward as expectations about the future (cf. Ref. 13).

Experience Corollary: "A person's construction system varies as he successively construes the replication of events" (Ref. 1, p. 72). With the passage of time, the perception of new events constitutes an ongoing validation process that serves to confirm or disconfirm some of an individual's anticipations. As a result, a person's constructs undergo continuous, progressive change. Kelly assumed that these changes in personal constructs are generally the result of predictive failures (cf. Ref. 14). As noted by Ford (Ref. 13, p. 190):

> We humans frequently anticipate the occurrence or non-occurrence of future events based on our willingness to project observed uniformities into the future. Thus, we continually glide from the past into the future with our previous experience preceding us—illuminating and organizing the manner in which subsequent events will be manifest to us.

The process through which people continuously anticipate events and test the efficacy of their constructions is termed the *experience cycle* (cf. Ref. 15). Kelly believed that this cycle of anticipation, investment, encounter, confirmation/disconfirmation and constructive revision represents a useful heuristic for conceptualizing human experience.

Range Corollary: "Each construct is convenient for the anticipation of a finite range of events only" (Ref. 1, p. 68). Each of a person's constructs has a *range of convenience,* which comprises "all those things to which the user would find its application useful." Accordingly, the range of convenience of a construct defines its extension in terms of a single aspect of a limited domain of events.[8] Not only individual constructs, but also, by implication, systems and subsystems of interrelated constructs have specific ranges of convenience. This suggests that some degree of functional differentiation among subsystems of constructs can enhance its overall range of convenience with respect to the variety of events that can be accommodated within its framework.[16]

Organization Corollary: "Each person characteristically evolves, for his own convenience in anticipating events, a construction system embracing ordinal relationships between constructs" (Ref. 1, p. 56). Constructs usually are deployed in conjunction with related constructs in interpreting and predicting events. Indeed, a necessary condition for organized thought is some degree of overlap between the constructs' ranges of convenience.[17] It is this overlap, or intersection, between the constructs' extensions that enables an individual to formulate "hypotheses." That is, in interpreting an event we essentially categorize it in terms of one or more constructs, and then by reviewing our personal systems of related constructs, we can derive predictive inferences from that initial categorization. For example, suppose that an individual's subordinate construct 'polite/rude' was subsumed by a superordinate construct 'considerate/inconsiderate.' This individual would expect considerate behavior from people who are polite. It is this predictive function of a person's construct system that provides the logical rationale for the Kellyan view that human beings are characterized by an *anticipatory stance*.

1.5. Repertory Grids

Kelly's *Role Construct Repertory Grid Test*[1] is essentially a method of eliciting constructs and analyzing relationships between them. It differs from conventional sorting tests, such as that devised by Vygotsky,[18] in that relationships between the categories are evaluated rather than the accuracy of the sorting. As Osgood, Suci, and Tannenbaum point out, Kelly's technique closely resembles Semantic Differential procedure.[19] In fact, some contemporary repertory grids are almost identical in form to Osgood's own instrument.

Elements (alternative events, states, or entities) and *constructs* (dimensions of similarity and difference between elements) are central to knowledge representation in repertory grids. The most basic form of repertory grid is a rectangular matrix with elements as columns and constructs as rows (see Figure 1). Each row-column intersect in the grid contains a rating showing how a person applied a given construct to a particular element. Kelly suggested several techniques for eliciting constructs from individuals,[20,21] however, a standard list of constructs relevant to a given context can be provided to respondents.[22,23]

Element Index:

1 - Normal
2 - Nonspecific Wall Abnormalities
3 - Nonspecific WMA/MVP
4 - Suspicion for Ischemia
5 - Abnormal Fx/Ischemia
6 - Abnormal Fx/Cardiomyopathy
7 - Abnormal Fx/Valve Disease
8 - Severe LV Dysfunction

Rate elements on a scale from 1 to 3; a '1' is at the LHP, a '3' is at the RHP			**ELEMENTS**							
			1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
C1	central crescent	displced or absent	1	1	2	3	3	2	2	3
C2	no blue fingers	blue finger	1	2	2	3	3	1	1	3
C3	concentric	asynchronous	1	2	2	3	3	1	1	3
C4	intact	interrupted	1	1	2	2	2	1	1	3
C5	improve/exer	worsen/exer	1	2	2	3	2	2	2	3
C6	symmetric	assymetric	1	1	2	3	3	1	2	3

Figure 1. Example of a repertory grid for diagnosis of heart wall motion abnormalities (adapted from figure in Ref. 24).

Constructs can be elicited by Kelly's *method of triads,*[1] that is, presenting elements three at a time and asking how any two of them are similar to each other and different from the third. For example, consider a physician who is diagnosing heart wall motion abnormalities (*severe-lv-dysfunction, abnorm-fx/valve-disease,* etc.). We might begin by asking in what way are two heart problems alike and thereby different from another one? The physician could reply that a *normal* heart condition and an *abnormal-fx cardiomypathy* condition are alike in that they exhibit *no blue fingers,* whereas if *abnormal-fx/ischemia* is the problem, a *blue finger* appears. *Blue finger/no blue fingers* are the poles of the construct.

There are practical, as well as theoretical, considerations that underlie the use of triads of elements in interviewing experts. Because people are so good at listing relevant distinctions, it is tempting to let them add constructs to the grid at will rather than using a structured interview process. While it is true that most experts can readily generate a set of terms to describe their domain, our experience is that unstructured methods typically produce less interesting terms than triad methods. Since triadic elicitation frames the task as one of distinguishing among elements, the expert generates a minimal set of discriminating dimensions, rather than a larger set of descriptive ones that may or may not be of practical use. In Ford et al.,[24] for example, the expert did not verbalize the key diagnostic factor (*blue finger* in Figure 1) until he completed a repertory grid. Experts often find using this tool leads to *discovery* rather than simple documentation of known facts and relationships. This is particularly true for experts who are on the leading edge of their profession.

With respect to generality, distinctions based on the presentation of two elements tend to be less robust. In Kelly's terms, they are relatively *impermeable*, that is, they are more specific to the two elements being considered and less likely to be applicable to new elements introduced later. Experts produce more *permeable* constructs when they consider three or more elements at a time.

Researchers have developed several forms of grids,[20,25-27] including implication grids, resistance-to-change grids, bipolar implication grids, dependency grids, exchange grids, and mode grids among others. They have also devised a variety of grid formats in which people, objects, events, situations, or other kinds of elements, are either categorized, rated, or rank-ordered on a set of constructs. Several analysis procedures have also been developed. We discuss some of these techniques in Section 3.2.

In addition to eliciting and analyzing knowledge, repertory grid techniques have other features that are useful in the knowledge acquisition process. First, they apply to a variety of problems. For instance, repertory grids can be viewed as a component of a database in entity-attribute form[28] with elements as entities, constructs as attributes, and allocations of elements to locations on construct dimensions as values. Secondly, representing knowledge in repertory grids can simplify the creation of interfaces to databases and spreadsheets.[28-31] Another advantage of grids is that they make it easy to inspect and analyze the organization and logic of expert knowledge. The spreadsheet-like visual metaphor amplifies the expert's ability to recognize and offer distinctions between the elements.[32] Recognition and completion of patterns in the data are facilitated by the structure and relative compactness of the matrix representation as compared to rules. Furthermore, representation by grids facilitates testing for ambiguity, redundancy, and incompleteness.[33]

2. Recent Theoretical Developments

Before our discussion of specific tools in section three, it is important to understand something of the background of their evolution. Each has been influenced by the particular theoretical and practical interests of their developers. Furthermore, because of extensive collaboration, and other professional interaction, a great deal of cross-fertilization across research groups has taken place. For example, work on ICONKAT at the University of West Florida has benefited from the collaboration of Novak, a major contributor to *assimilation theory*. *Concept maps,* originally developed by Novak,[34] and applied to educational settings, are a mediating representation that have been used successfully by Ford, Adams-Webber, and their colleagues[24] in several aspects of the knowledge acquisition process (see Section 3.1.2). Both the theory and the technique are constructivist. We discuss assimilation theory and concept maps in Section 2.1.

A further set of developments has served to clarify the underlying rationale of repertory methods. *Aquinas*[35] was a knowledge acquisition tool that provided an explicit hierarchical representation for elements and constructs. We briefly describe the approach in Section 2.2. At the University of Calgary, Gaines and Shaw[36-39] are working to produce a theoretical foundation for tools that organize constructs into hierarchies. They have defined an intensional logic of distinctions that is compatible with the KL-ONE family of semantic network representations.[40,41] In addition, the research group associated with the University of West Florida[32,42,43] have developed a logic of confirmation that incorporates the basic tenets of personal construct psychology directly into the logic as grounds for the determination of relevance, thus strengthening the logic and extending personal construct theory. In this approach, the degree of confirmation is characterized by epistemic probabilities arrived at by measuring the overlap (partial entailment) between the constructs' extensions represented as binary bit strings. We discuss these developments in Section 2.3.

Finally, at The Boeing Company, Bradshaw and Boose[44] have combined repertory grids with decision-analytic representations called *influence diagrams* and *possibility tables*.[30,45] The new understandings resulting from work combining repertory grids and influence diagrams provide a basis for principled reasoning under uncertainty and explicit representation of preferences for problems involving significant risk, high stakes, or complex tradeoffs. Work combining repertory grids and possibility tables extends constructivist techniques to exploratory design and configuration problems.[46] We discuss influence diagrams and possibility tables in Section 2.4.

2.1. Assimilation Theory and Concept Maps

Ausubel's assimilation theory is a cognitive learning theory that has been widely applied to education.[47,48] Like Kelly's personal construct theory, it is based on a constructivist model of human cognitive processes. Specifically, it describes how concepts are acquired and organized within a learner's cognitive structure.

Ausubel argues that learning is synonymous with a change in the meaning of experience. His fundamental premise seems deceptively simple (Ref. 48, p. 159):

> Meaningful learning results when new information is acquired by deliberate effort on the part of the learner to link the new information with relevant, preexisting concepts or propositions in the learner's own cognitive structure.

In short, meaningful learning involves the assimilation of new concepts and propositions into existing cognitive structures. In Ausubel's model, cognitive structure can be described as a hierarchically organized collection of concepts representing one's knowledge and experience.[34] Concepts are perceived regularities in events or objects, designated by a label.[24] Assimilation theory stresses that meaningful learning requires that the learner's cognitive structure contain anchoring concepts to which new material can be related or linked. For this reason, Ausubel argued that "the most important single factor influencing learning is what the learner already knows. Ascertain this and teach him accordingly."

The concept map is assimilation theory's major methodological tool for ascertaining what is already known. In educational settings, concept mapping techniques have aided people of every age to examine many fields of knowledge. Much of the assimilation theoretic research to date has involved and exploited concept mapping.[49] In addition, concept maps are of increasing interest to those engaged in the process of knowledge acquisition for the construction of knowledge-based systems.[24,50] Essentially, concept maps provide context-dependent representations of a specific domain of knowledge within a set of concepts. They are constructed so that the interrelationships among the included concepts are evident. In fact, concept maps have been shown to help students "learn how to learn" by making explicit their personally constructed knowledge and providing a structure for linking in new information. As a mediating representation, concept maps offer a flexible framework for eliciting, representing, and communicating the emerging domain model. In this way, they are well suited to the view of knowledge acquisition as a constructive modeling process in which the knowledge engineer and domain expert collaboratively build a domain model.

Concept maps structure a set of concepts into a hierarchical framework. More general, inclusive concepts are found at the highest levels, with progressively more specific and less inclusive concepts arranged below them. In this way, concept maps display Ausubel's notion of subsumption, namely that new information is often relative to and subsumable under more inclusive concepts. All concepts at any given level in the hierarchy will tend to have a similar degree of generality. Figure 2 shows a portion of a concept map produced by an expert in nuclear cardiology (by convention links run top-to-bottom unless marked with an arrowhead).

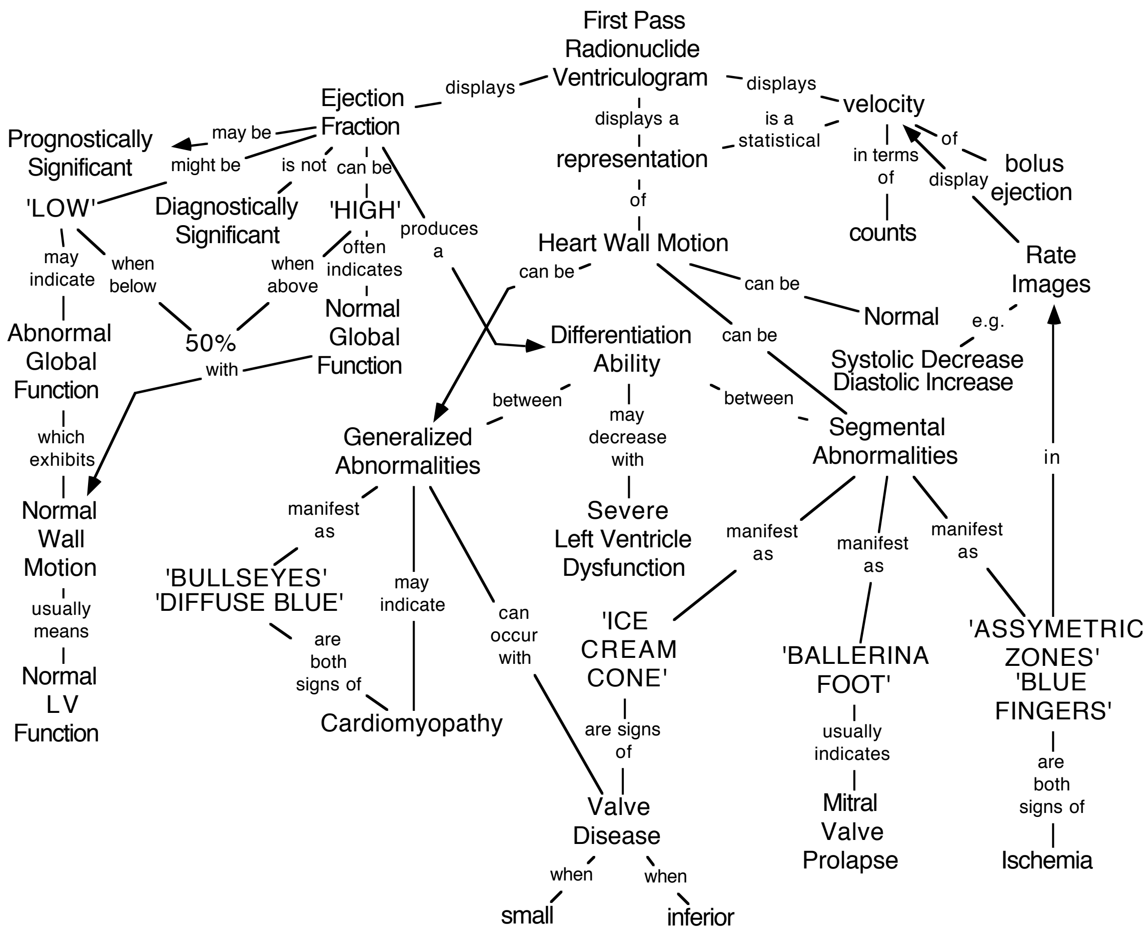

Figure 2. A portion of a concept map from the domain of nuclear cardiology (adapted from figure in Ref. 24).

Relationships between concepts in a map represent propositions. Propositions form semantic units by linking together two or more concepts. In its most rudimentary form, a concept map contains just two concepts connected by a linking word to form a proposition. For example, "John is tall" would represent a simple map forming a valid proposition about the concepts "John" and "tall". A concept acquires additional meaning as more propositions include it. Thus, "John is tall", "John is a person", "John eats" and so on, all expand the meaning of the concept "John". In this sense, we can think of concept maps as representing meaning in a framework of embedded propositions. Much of the expressive power of concept maps comes from the fact that the user is free to employ an unlimited set of linking words to show how meanings have been developed. When concepts and linking words are carefully chosen, these maps are powerful tools for representing and communicating nuances of meaning.

In the ICONKAT system, concept maps are an important mediating representation used to provide a hierarchically ordered, conceptual overview of the domain model arising from the collaborative efforts of the expert and knowledge engineer. The concept maps provide "knowledge landscapes" (essentially topographical maps) of the domain that inform the knowledge engineer about the potential and appropriate use of other methods, such as repertory grids. For example, they may indicate where there is enough knowledge (at given level abstraction) to warrant the use of a repertory grid. More importantly, in the ICONKAT environment, concept maps comprise the organizational structure for entire domain model. It is into this semantic structure that other mediating representations (e.g., repertory grids, video, text, etc.) can be linked. Finally, the ICONKAT approach to explanation relies on the aforementioned organizational structure of the domain model represented as a hierarchical collection of concept maps. In ICONKAT, users construct their own explanations while navigating their way through the linkages among clusters of related mediating representations constituting the model.

Although ICONKAT provides the most extensive set of tools to assist with concept mapping, KSSn/KRS and DDUCKS also have benefited from the use of concept-map-like structures.

2.2. Hierarchical Knowledge and Repertory Grids

The need to represent hierarchies of elements and constructs at varying levels was recognized by Boose[51,52] in his work on the Expertise Transfer System (ETS). For instance, one of the first experts interviewed tried to build a jet engine diagnostic aid. Parts and systems were listed as potential elements (problem areas), and diagnostic symptoms were generated through triadic comparison. Unfortunately, elements such as "spark plug" and "electrical system" both appeared in the same grid. This caused difficulties when using grid elicitation techniques and some analysis tools. Even in cases where grid information was at a similar level of abstraction, there was a limit to how much information could be comfortably represented in a single rating grid. One application included a 38-by-35 grid, but it was hard for the expert to manage and comprehend that much information at once. A method was needed to decompose large grids into manageable, related subgrids.

In later versions, ETS used a *laddering* technique developed by Hinkle[26] that asked “how” and “why” questions to elicit constructs at different levels of abstraction. Then, in their work on *Aquinas*, Boose and Bradshaw[35,53] developed a scheme in which the underlying representation of grids was no longer a two-dimensional matrix, but rather a network of linked frames in a multidimensional space. A particular frame represented information about a concept at the intersection of a particular case (problem area), an expert (or other knowledge source) who knew something about that case, an element that represented an alternative that was a possible solution for the problem, and a particular distinction that was relevant to the selection of the alternative. Experts used a “map view” to navigate through the knowledge base (Figure 3). By selecting combinations of nodes in the case, expert, element, and construct hierarchies, they could specify a portion of the knowledge base to be displayed as a repertory grid. Structural changes made in the network views were immediately reflected in the grids, and vice versa. Various dialogues, in conjunction with laddering and statistical analysis techniques, helped experts decide how to decompose and structure the hierarchies.

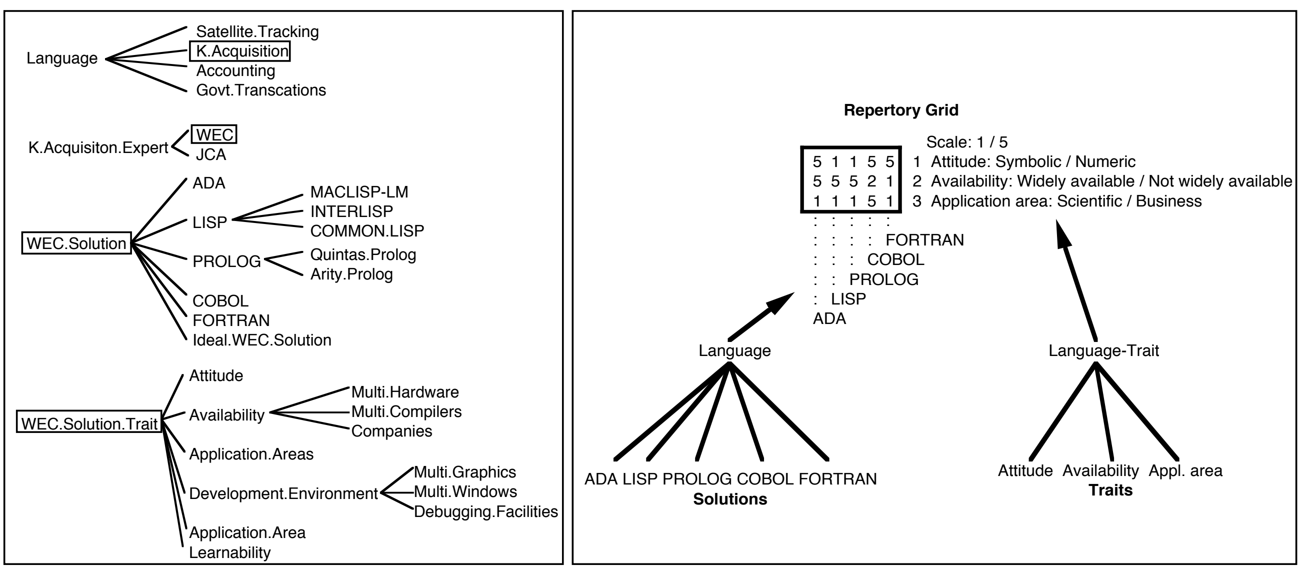

Figure 3. Experts select nodes from each hierarchy to show portions of the knowledge base in grid format (adapted from figure in Ref. 35).

Gaines and Shaw’s[54] intensional logic provide a theoretical foundation for generalizing and refining some of the ideas previously demonstrated in *Aquinas*. As described in the next section, their theoretical framework also provides a conceptual bridge between constructivist knowledge acquisition tools and the semantic network representations for which relationships between concepts are a major focus.

2.3. Recent Developments in the Logic of Personal Construct Theory and Repertory Grids

An important set of studies has helped to clarify the logical rational underlying some constructivist knowledge acquisition tools. Gaines and Shaw have defined an intensional logic of distinctions[54] that is compatible with the KL-ONE family of semantic network representations.[40,41] We describe these developments in Section 2.3.1. In a related vein, Ford and colleagues[32,42,43] have proposed a theory of confirmation that incorporates the basic tenets of personal construct psychology directly into the logic as a basis for the determination of relevance, thus strengthening the logic and extending personal construct theory. Section 2.3.2 gives a brief summary of these developments.

2.3.1. Intensional Logic and Semantic Networks

As a foundation for KSSn/KRS, Gaines and Shaw[54] show how distinctions may interrelate. They take the relations of subsumption and disjunction to be minimally sufficient to define an intensional logic of distinctions, from which more complex relations may be derived. Subsumption between computational concepts corresponds to the "is-a" relation common in semantic network representations, while disjunction corresponds to the definition of disjoint concepts (Ref. 54, p. 9):

> A concept is defined to be *that mental entity imputed to a distinction making agent as enabling it to make a particular distinction*....A construct is defined formally to be a triple of two disjoint distinctions mutually subsumed by a third… and psychologically as the triple of concepts assumed to underlie the distinctions.

Figure 4(A) illustrates how the notion of similarity is captured through the shared concept *age*, while the notion of contrast is captured through the disjunctive arc separating the subsumed concepts *young* and *old*. The minimal subsuming concept, *age,* illustrates the idea that a construct is convenient for anticipating a finite range of events. Since concepts in the graph are nodes, and constructs are arcs, Gaines and Shaw conclude that concept and construct are graph-theoretically dual relations.

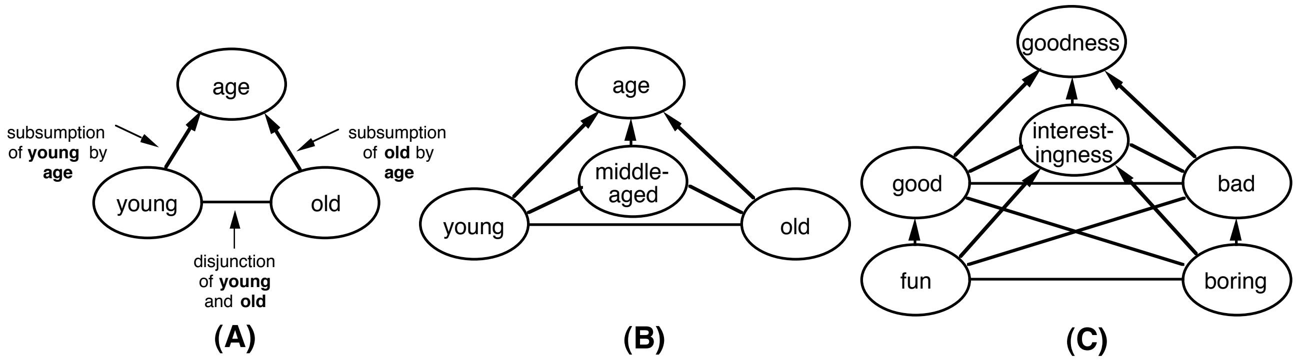

Figure 4. The structure of a construct (A); a construct with three alternative values (B); and two constructs in an ordinal relationship (C) (adapted from figure in Ref. 54).

Figure 4(B) shows a triple of disjoint concepts (young—middle-aged—old) that could be seen as alternative values along an age dimension. The possibility of putting events along numeric scales may be represented by extending this structure in various ways. Ordinal relations between constructs may be derived from the ordinal relation of subsumption, as shown in Figure 4(C). Here, the construct characterized by the triple of goodness: good–bad subsumes the construct interestingness: fun–boring.

Applying this conceptual framework to events requires a description of how the distinctions relate to the things distinguished or represented (i.e., *elements*). Figure 5 represents Ferio and Jules as elements (or *individuals*) placed at particular points. Individuals are necessarily leaf nodes in the graph representation. We can characterize the repertory grid as a matrix of concepts, individuals, and constraints.

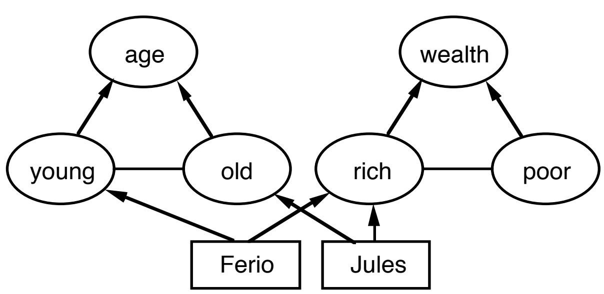

Figure 5. Ferio and Jules as elements (individuals) in relation to the constructs of age and wealth (adapted from figure in Ref. 54).

Gaines and Shaw's work has laid a foundation for implementing constructivist representations that correspond to a formal semantics for semantic nets. This will enable a better exchange of ideas between researchers. Work in this area has also produced a "visual language" for semantic networks. We discuss the implementation of this language in KSSn/KRS in Section 3.1.3 below.

2.3.2. The Logic of Confirmation and Personal Construct Theory

The process of deriving construct relationships from repertory grid data relates to the problem of induction—the production of universal generalizations based on a finite number of evidences. The inductive probability of an argument depends on the strength of the evidence that the premises provide for the conclusion. Closely related to efforts aimed at developing an adequate logic of induction, are those focused on elaborating what is known as the logic of confirmation. The central problem in the study of the logic of confirmation has long been the problem of *relevance*. In this context, a useful theory of relevance is one that plausibly elucidates the method implicit in judgments of confirmation as performed by actual humans—not by some imaginary fully rational being. The shifting relevance of aspects of ordinary situations would be cause for gloom in AI and cognitive science, were it not for the successful human exemplar.

Consider the paradox of the raven.[55] Most observers find it disturbing that the existence of a white handkerchief can be formally shown to confirm the hypothesis "all ravens are black." This result offends our intuitions which hold the existence of a white handkerchief irrelevant to a hypothesis about ravens. Ford and Adams-Webber[32,43] have elaborated a constructivist approach to classificatory confirmation that justifies this natural intuition. Despite their formal equivalence in terms of symbolic logic, the proposition that "all ravens are black" and "all non-black things are non-ravens" are pragmatically very different. That is, they are not confirmed and disconfirmed by the same evidence.

In an attempt to address the paradox of the raven, Von Wright[56] noted that generalizations have an associated range of relevance, and consequently, only things within a generalization's range of relevance may constitute confirming or disconfirming evidence. All other things (i.e., things outside the range of relevance) are irrelevant. Furthermore, when a generalization's range of relevance is not specified (which is typically the case), it is taken to be the "natural range of relevance," meaning "the class of things that fall under the antecedent term."[56] However, this seems a little like question begging; a serious practical problem remains, how is this "natural range of relevance" to be operationally defined?

The class of items deemed as falling under the antecedent term will vary from person to person and over time. From a Kellyan point of view, the problem of epistemic confirmation (i.e., non-demonstrative inference in the service of individual fixation of belief) is at the most primitive level, fundamentally psychological in nature, and will not submit to a purely syntactic (i.e., *a priori* analytic) approach. Ford and Adams-Webber have suggested that Kelly's range corollary (discussed in Section 1.4) can lend Von Wright's theoretical notion of relevance a basis for realization.[32,42,43] According to personal construct theory, hypotheses are based on the overlap or intersection of the constructs' ranges of convenience. Moreover, the repertory grid provides a method of operationally defining Von Wright's range of relevance for a given hypothesis in the universe defined by the grid elements. Specifically, Ford and Adams-Webber have elaborated a theory of non-quantitative (or classificatory) confirmation[32,42,43] that incorporates the fundamental tenets of personal construct psychology directly into the logic as a foundation for the determination of relevance, thus strengthening the logic, and extending personal construct psychology. This work on a classificatory theory of confirmation provided the foundation for the subsequent development of a quantitative logic of confirmation (discussed below) consistent with personal construct psychology.

Logic is traditionally presented as if there is a great conceptual chasm between considerations of deductive and inductive logic. In fact, inductive logic is frequently regarded as a contradiction in terms, or at best as a poor sibling of deductive logic. However, when operating from within the logical framework of entailment, the processes of induction and deduction may be intimately related. Deductive logic can be characterized by the idea of complete logical entailment, while inductive logic can be described by a relation of partial entailment. In this limited sense, deductive reasoning may be considered a special case of inductive reasoning. This situation is portrayed diagrammatically in Figure 6.

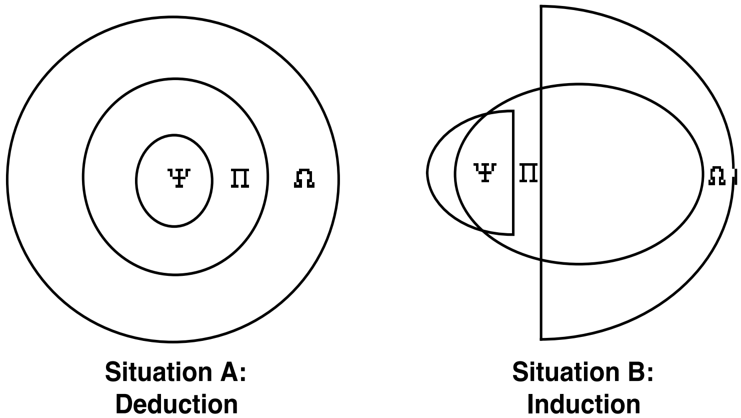

Figure 6. Pictorial representation of deduction and induction as entailment (adapted from figure in Ref. 57).

In Figure 6, the case of deductive logic is illustrated by situation A, which represents the universal propositions “all Ψs are Π,” “all Πs are Ω,” and “all Ψs are Ω.” Situation B represents the case of inductive logic where the amount of overlap or partial entailment is measured by degree of confirmation. Note that situation B illustrates the propositions “most Ψs are Π,” “most Πs are Ω,” and “no Ψs are Ω.” Thus, transitivity does not hold under this probabilistic interpretation of induction.

Although diagrams such as those in Figure 6 are intuitively helpful, the task of measuring partial entailment remains problematic. Any method founded on purely numerical mechanisms cannot provide the foundation for a probabilistic logic with truth-functional connectives.[58] Likewise, systems that assign numerical truth values to propositions or nonmaterial conditionals cannot provide a truth-functional probability logic without additional information about the relationships between the atomic components of the antecedent and consequent.

We consider a hypothesis such as, “all Ψs are Π” to be a nonmaterial conditional of the form, “Whatever χ might be, if χ is a Ψ then χ is Π.” We think of this conditional as affirming a bundle of individual conditionals:[59] “If χ_1 is Ψ, χ_1 is Π;” “If χ_2 is Ψ, χ_2 is Π;” and so on. The probability of such a proposition is based on a sample space of points in a universe w, corresponding to situations in which the proposition will be either true or false. Thus, we represent such nonmaterial conditionals as bit strings consisting of 1’s and 0’s. For example, $i(\Psi) = (i(\Psi)_1, i(\Psi)_2, \ldots, i(\Psi)_n)$ denotes the binary bit-string representing the occurrence or nonoccurrence of Ψ. Thus, we have:

$$i(\psi)_i = \begin{cases} 1 & \text{if } \chi_i \text{ is a } \psi \\ 0 & \text{otherwise} \end{cases}$$

In other words,

$$\begin{aligned}
&\text{if } \chi_1 \text{ is a } \Psi \text{ then } i(\Psi)_1 = 1, \\
&\text{if } \chi_2 \text{ is a } \Psi \text{ then } i(\Psi)_2 = 1, \\
&\text{if } \chi_3 \text{ is not a } \Psi \text{ then } i(\Psi)_3 = 0, \\
&\quad\vdots \\
&\text{if } \chi_n \text{ is a } \Psi \text{ then } i(\Psi)_n = 1, \text{ otherwise } i(\Psi)_n = 0.
\end{aligned}$$

Likewise, the symbol i(Π) is referred to as the “incidence of Π,” and denotes an ordered subset of w consisting of all points in which Π occurs or is true.

The bit strings described above provide the needed information for a truth functional probability logic applicable to the measurement of partial entailment. Bundy[58] has proposed an incidence calculus in which a set theoretic function is associated with each logical connective from propositional logic. An extended version of this calculus has been applied to the measurement of partial entailment.[32,42,43] In this approach, the degree of confirmation is characterized by epistemic probabilities arrived at by measuring the overlap (partial entailment) between the constructs’ extensions represented as binary bit strings. ICONKAT and its predecessor *Nicod* both employ this method to automatically generate rules from repertory grid data. DDUCKS and KRS/KSSn use somewhat different approaches to derive rules from repertory grid data, as briefly described in section 3.2.

2.4. Decision Analysis, Influence Diagrams, and Possibility Tables

Many knowledge-based systems are prescriptive in nature. They aim not only to describe some actual or potential state of affairs, but also to recommend specific actions. Recommendations made by such systems depend on: the alternatives available, information about consequences associated with the alternatives, and preferences among these consequences. Unfortunately, knowledge-based systems typically treat preferences implicitly and heuristically, making no provision for value structures differing from those built into the system. In this section we discuss approaches combining decision analysis with constructivist methodologies to overcome these limitations.

2.4.1. The Need for Explicit Preference Models in Knowledge-Based Systems

In their discussion of preferences, Langlotz, Shortliffe, and Fagan[60] cite an example rule from MYCIN. This heuristic captures a physician's knowledge that tetracycline therapy should be avoided for children because it may cause dental staining.

> If:
>
> 1) The therapy under consideration is tetracycline
>
> 2) The age (in years) of the patient is less than 8
>
> Then:
>
> There is strongly suggestive evidence (.8) that tetracycline is not an appropriate therapy for use against the organism.

Clancey[61] gives a possible chain of four support rules for this heuristic. The first three inferences have to do with how one event relates to the occurrence of the next. The fourth, however, is a compiled plan of action based on the inference chain. Langlotz et al.[60] make the point that no matter how finely we break down a chain of reasoning, one rule in the chain will always recommend action based on the situation. Action recommendations always presuppose a set of preferences, either stated or implied, that cannot be derived from the logic of evidence.

```
tetracycline in youngster
    => chelation of the drug in growing bones
            => teeth discoloration
                    => undesirable body change
                            => don't administer tetracycline
```

Because of the nature of heuristics, it is difficult to represent explicitly and flexibly the unique circumstances and tradeoffs that may justify an exception to the heuristic. What if, for example, we found one or more of the following to be true:

- the infecting organism were resistant to all drugs except tetracycline?
- the only undesirable bodily change that tetracycline caused was minor intestinal distress?
- the probability of staining due to tetracycline for a particular patient was only 1 in 100? 1 in 1000?

When tradeoffs are embedded implicitly within heuristics, it becomes impractical to ask, let alone answer, such questions. For example, we could modify the knowledge base by adding additional premises to the rule above:

3) The organism can be treated by something other than tetracycline

4) There is evidence that tetracycline will cause significant intestinal distress

5) The probability of dental staining due to tetracycline for the patient is less than .01

But since the strength of our recommendation may vary depending on the circumstances present in a given situation, we would need to add a separate rule for each action and each particular combination of evidence. Representing the knowledge in this form makes it impossible to vary the parameters of preference tradeoffs (e.g., risk of dental staining versus effectiveness of tetracycline versus cost of treatment) smoothly in response to differences in situation and preferences between patients. While it is possible to muster empirical arguments for the truth or falsity of some *evidential* claim, the judgments of *utility* that guide recommendations and action (given that evidence) are inherently subjective: some patients are more willing to take risks than others; some are more concerned about treatment effectiveness; some are more able or willing than others to pay for expensive alternatives that minimize risk. The greater the stakes of the decision, the more serious are the consequences of implicit, inflexible representations of preferences. For this reason, several researchers have discussed the need to include such explicit preference models in the knowledge engineering process.[60,62-66]

Bradshaw and Boose[44,45] argued that decision analysis could be effectively combined with constructivist methods to model complex problems. Early on, they realized that heuristic approaches to uncertainty and preferences were inadequate for high-stakes decision-making. After evaluating alternative approaches, they settled on a method that combines repertory grids with decision-analytic representations called *influence diagrams*.

2.4.2. Combining Repertory Grids and Influence Diagrams

Influence diagrams have been an important advance in the representation of decision problems[67,68] and recent developments have extended their usefulness as a structuring and communication device between participants in the decision.[64,69-73] Influence diagrams can be directly solved to obtain recommended actions in a way that is consistent with probability and utility theory. In addition, several analysis techniques (e.g., sensitivity analysis, value of information, value of control) can be used to gain insight into the problem being represented.

Figure 7 shows a screen snapshot of a DDUCKS virtual notebook containing an influence diagram. The diagram represents a generic medical decision making template.[74] The problem is to determine the best treatment alternative for a cancer patient, taking treatment risks and other diagnostic uncertainties into account. The treatment strategy is composed of two decisions (Test, Treatment), represented by square nodes on the diagram. Round nodes represent treatment uncertainties (Results, Therapeutic Effect, Side-Effect), diagnostic uncertainties (Patient Demographics, Observable Symptom, Hypothetical Disorder, Physiological Need), and Cost. The eight-sided node labeled "Value" has been designated as the criterion to maximize in evaluating the model to determine the best treatment strategy.

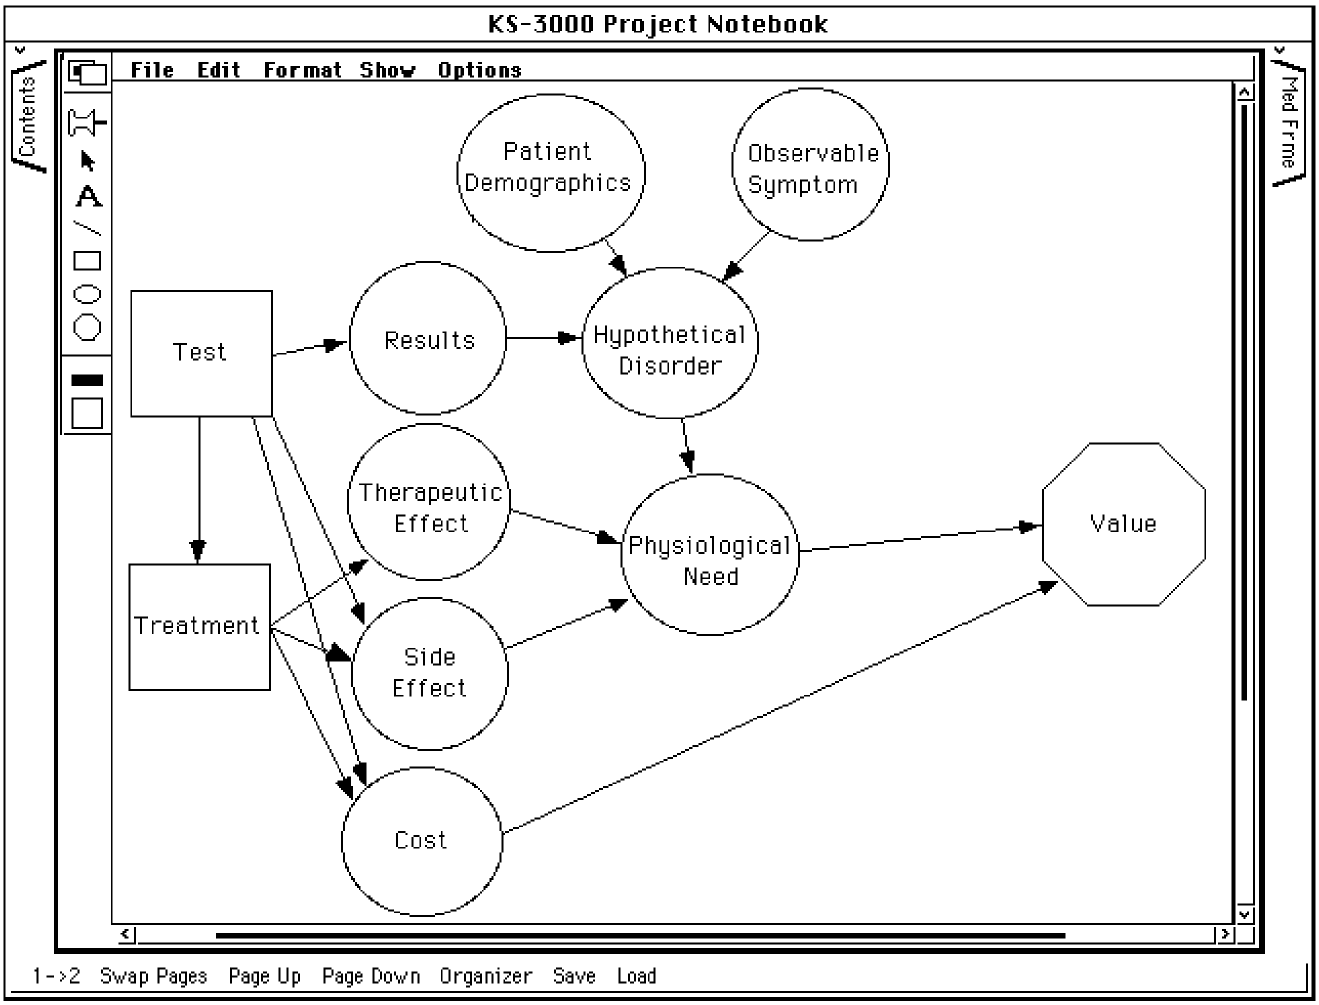

Figure 7. A DDUCKS virtual notebook containing an influence diagram for a generic medical decision making template.

Unfortunately, the creation of valid influence diagram models can require a relatively high level of sophistication in the theory and practice of decision analysis. Influence diagram-based tools contain some of the algorithms of decision analysis practice, but cannot embody the experience and intuition of decision analysis professionals in formulating and appraising decision models.

Constructivist knowledge acquisition methods can be used to overcome some of these problems. For example, Gaines[75-77] has elaborated aspects of constructivist theory that bear on the role of preferences in personal decision making. He proposes a hierarchical model that posits two fundamental processes operating as a person models the world:

- flow of *information* as surprise about events ("news of a difference"[78]) upward through the hierarchy when lower levels cannot account for events, and

- flow of *preferences* downward as lower-level predictive models accounting for events are created to be consistent with higher-level ones. The flow of preference can ultimately result in action as higher levels attempt to influence the anticipated future.

Bradshaw and Boose have attempted to make operational certain aspects of Gaines' model by making a distinction between *information grids* and *preference grids*. Information grids represent beliefs about events, qualities, or states of the world and conditional probability relationships to other events, qualities, or states (as in the grid in Figure 1). In preference grids, the elements represent alternatives for a decision and constructs represent distinctions of utility that are used to select the best alternatives (as in the possibility table example below). For details on how repertory grids and influence diagrams may be combined, see Bradshaw, Covington, Russo and Boose.[29,30]

2.4.3. Combining Repertory Grids and Possibility Tables for Synthesis Problems

In addition to developing methods for representing and reasoning with explicit preferences, Boose and Bradshaw sought for ways to overcome other tool limitations for certain classes of decision-making problems. There is a traditional distinction in the literature between *analysis* and *synthesis* problems.[79,80] Analysis problems are generally defined as those for which the alternatives can be enumerated comfortably (e.g., simple classification or diagnosis). On the other hand, problems involving synthesis (e.g., configuration, scheduling, planning, and design) are subject to combinatorial explosion, typically involving far too many possibilities to list. Synthesis problems are often solved by *constructing* (rather than merely *selecting* between) alternatives. These alternatives are constructed so as to be consistent with hard constraints and "good enough" with respect to soft constraints.

Many papers in the literature have disparaged classification models, suggesting that they are inherently inferior to simulation models. However, Clancey[81] has eloquently argued for the necessity and irreducibility of classification models, and researchers such as Gaines[36] have demonstrated elegant approaches to resource allocation problems based entirely on classification. Bradshaw et al.[82] suggested that decision analysis and constructivist methodologies could be combined to treat synthesis problems as a sequence of decisions subject to local and global constraints. Instead of using influence diagrams, the approach involved connecting repertory grids to representations called *possibility tables*, which have been used manually in configuration and design problems for many years.[83-85] This approach is implemented in DDUCKS as *Canard*.[86]

Within *Canard,* hierarchical possibility tables are used to structure information about complex alternatives, outcomes, or plans (Figure 8). Columns in the possibility table represent components, functions, or issues that relate to the artifact being designed. Within each column, the various possibilities listed identify a set of options being considered. Above each column, a set of distinguishing criteria appears. Comments about alternatives, columns, possibilities, criteria, ratings, preferences, and constraints are portrayed in text annotation panes within the possibility table and repertory grid views. Annotations define, justify, or assert something about a particular element in the table or grid. To construct a design alternative, a designer selects a set of possibilities from one or more columns, defining a path through the table. The system can also suggest new alternatives by permuting the constraint space. The names of alternatives appear in the leftmost column.

Possibility tables can be associated with repertory grids to enter information about criteria affecting the choice of options. A column of possibilities are represented as elements in the grid, while the criteria are shown as constructs. *Canard* can use repertory grid techniques to elicit and structure information about design possibilities and gather the constraints and criteria that guide a designer in the selection of these possibilities. Repertory-grid-based analysis tools help designers determine the adequacy of the constraints and objectives and focus their attention on descriptions needing further refinement. (See Ref. 86 for additional information about *Canard's* constraint elicitation, propagation, and refinement techniques.)

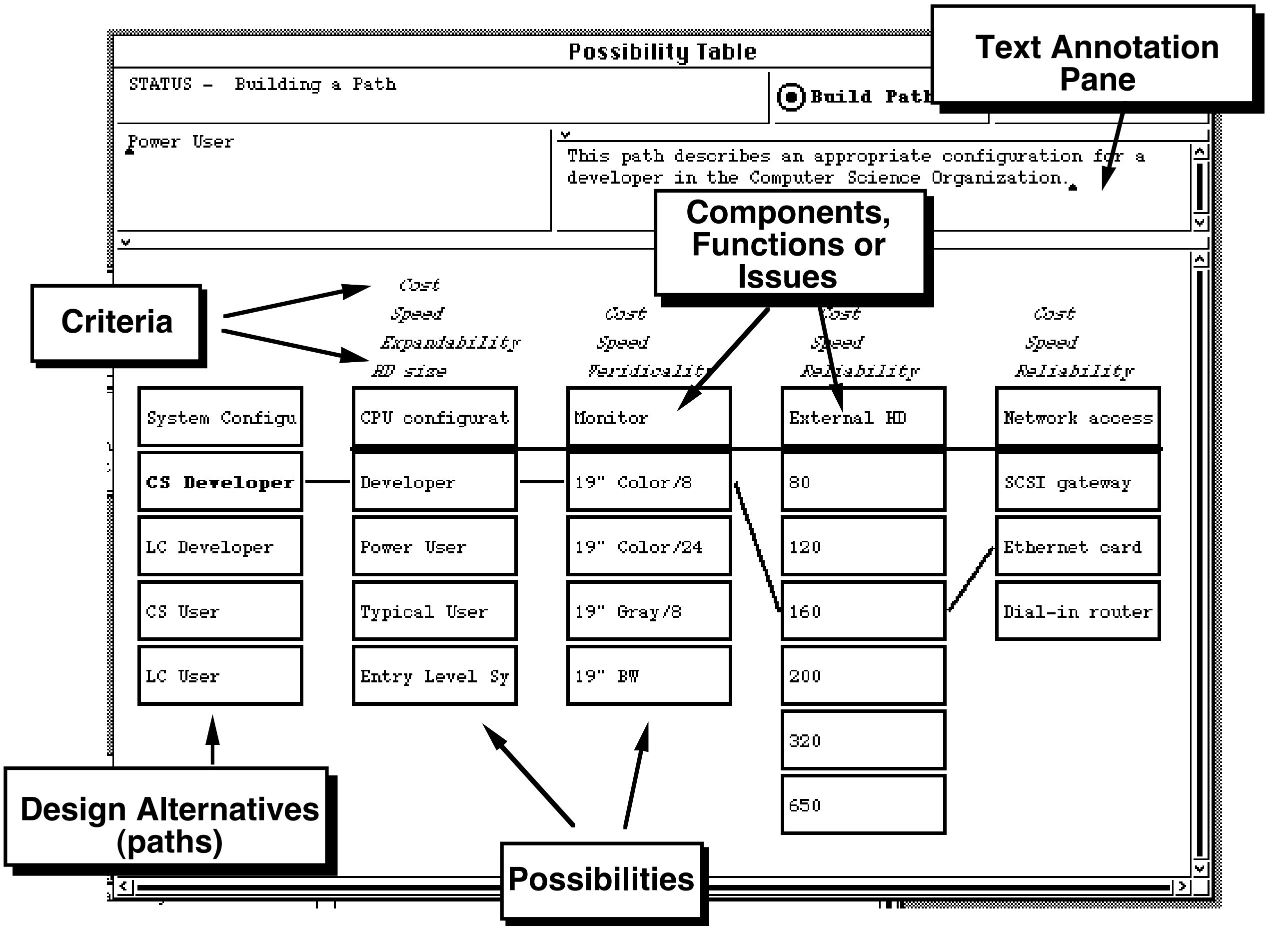

Figure 8. A *Canard* possibility table for configuration of a computer system.

We have seen how decision analysis techniques can be combined with constructivist representations and techniques to model information and preferences in an explicit and rigorous manner. This is especially valuable for problems involving a high degree of uncertainty, significant risk, high stakes, or complex tradeoffs.

3. A New Generation of Constructivist Knowledge Acquisition Tools

Over the past several years, many tools incorporating repertory grids have been applied to knowledge acquisition. These include *Aquinas*,[35,53] DART,[46] ETS,[51,52] Flexigrid (see [87]), FMS Aid,[88] KITTEN,[89] Kriton,[90] KSS0,[76,77] *Nicod*,[32] PCS,[91] and PLANET.[92-94] In addition, Wahl[95] reported success in applications of repertory grid techniques, and validation in a statistics domain was discussed by Gammack and Young.[96]

In the remainder of the paper, we focus our attention on three "second-generation" constructivist knowledge acquisition systems: DDUCKS,[74,97,98] ICONKAT,[24,99] and KSSn/KRS5.[54,100] These systems form an interesting cross-section of the state-of-the-art for two reasons:

1. They embody the extensive theoretical work integrating the complementary perspectives discussed in Section 2.
2. They have evolved through use of the tools in a variety of contexts and overcoming the limitations discovered their application.

In discussing these tools, we pass over much of the development history leading up to these efforts. More comprehensive historical accounts have been written by Boose et al.,[101] Ford et al.,[102] and Gaines.[103]

We begin with a discussion of general architectural and user-interface features of the tools (section 3.1). Following this, we describe new developments in analysis and induction techniques (3.2), multiple expert analysis, and group use of tools (3.3).

3.1. General Architectural and User-Interface Features

There are interesting similarities and differences in the architectures and user-interface of the three systems. We give a description of each tool below.

3.1.1. DDUCKS

DDUCKS (Decision and Design Utilities for Comprehensive Knowledge Support) is an "open architecture" constructivist knowledge modeling environment. Researchers are exploring how individual knowledge modeling and decision support tools can work cooperatively with one another and with commercial applications such as spreadsheets, databases, or hypermedia software.[104] Applications include enterprise modeling and integration,[105] group decision support in an electronic-meeting room environment,[106] and bone-marrow transplant patient follow-up care.[74]

It is useful to think of DDUCKS in terms of four layers of functionality: workbench, shell, application, and consultation. Starting with any layer in the system, a user can produce a set of tools, models, and ontologies that can be used to help in configuration of a more specialized system at the layer below.[98,107] This approach was inspired in many respects by the success of the PROTEGE architecture in facilitating reuse and configuration of methods, tasks, and mechanisms for diverse applications.[108,109]

DDUCKS is based on a three-schemata approach to knowledge representation that distinguishes between external, conceptual, and internal schemata.[107,110] As an implementation of the external schemata, we emphasize the use of *mediating representations* that serve as a means of communication between expert and knowledge engineer. *Intermediate representations* implement the conceptual schema, and help bridge the gap between the mediating representations and a particular implementation formalism.

Intermediate representation. The intermediate representation (i.e., concept model) consists of entities, relationships, and situations as the primary concepts, and domains, properties, and constraints as secondary concepts.[98] DDUCKS uses CODE4 as the underlying conceptual representation.[111-113] The concept model in CODE provides facilities for semantic unification of information that may be simultaneously portrayed from a number of perspectives (e.g., repertory grids, concept maps, possibility tables, influence diagrams).[98] It sits between the views and the implementation formalism, translating the user's actions into changes in knowledge and database structures. The general taxonomy for conceptual modeling has been derived from Tauzovich and Skuce,[114] with extensions for dynamic and epistemic aspects of the model. CODE4 provides a rich paradigm for the definition of knowledge-level concepts. A collection of integrated tools supports the important and frequently overlooked aspects of conceptual, ontological, and terminological analysis.[115] We are developing extensions to the representation to allow the system to make use of additional inferencing and representation facilities similar to those found in Sowa's[116,117] conceptual graphs and Gaines' KRS,[54,99] which interpret taxonomic and entity-relationship structures in terms of typed formal logics. A first order logic system and a simple natural language system allow various types of syntactic and semantic checks to be performed, if desired. A comprehensive lexicon allows references to concepts to be automatically maintained and quickly accessed. We emphasize the importance of comprehensive lexical support so that terminology can be carefully chosen and subsequently managed. Concept libraries and default inferencing mechanisms can be augmented by users employing graphical views and an integrated scripting and query language. A translator is currently under development[98] to allow conversion of knowledge represented into KIF[162] syntax by means of Gruber's Ontolingua.[164]

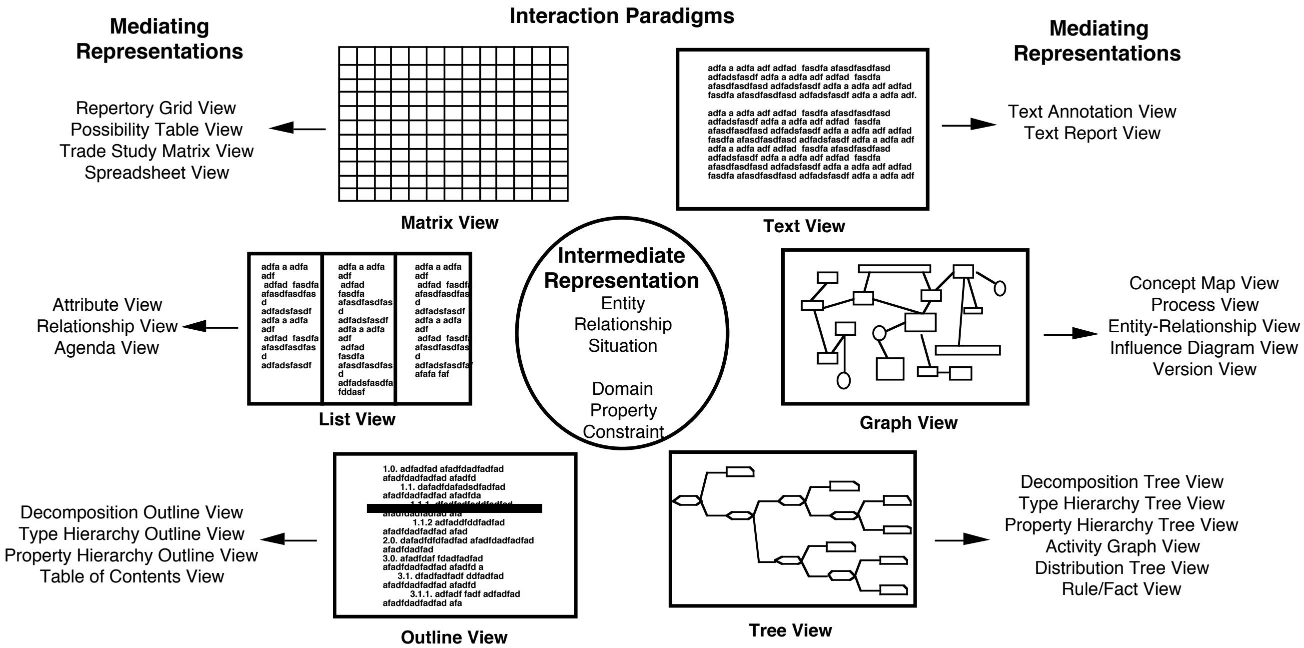

Figure 9. The intermediate representation in DDUCKS, surrounded by examples of generic interaction paradigms, and mediating representations.

Interaction Paradigms. User-interface management systems (UIMS) are becoming an essential part of interactive tool development and end-user tailoring.[118] We are extending the capabilities of a Smalltalk-80-based direct-manipulation user-interface builder to build a DDUCKS UIMS, called *Geoducks*. *Geoducks* relies on the Smalltalk-80 MVC (model-view-controller) approach for managing consistency among views.[119-121] The MVC approach provides a way to factor out the data in an underlying model from the data in dependent views, so that changes to the model in one view are immediately reflected in all related views.

The six views surrounding the intermediate representation (see Figure 9) correspond to the generic user-interface interaction paradigms that are implemented as abstract "pluggable" view classes.[120-122] These views are generic in the sense that they define the graphical form for the representation, but the form has no underlying semantics. Within DDUCKS, various configurations of these interaction paradigms can be called up in *sketchpad mode* to record free-form graphical and textual information. For example, individuals and groups can capture back-of-the-envelope drawings, agendas, issues, action items, requirements, concept-map-like structures, and other information pertinent to their task. While not part of the formal model, users can link elements created in sketchpad mode to elements in other views in hypertext fashion.

Mediating representations. Specialization facilities for concepts in a modeling ontology, in conjunction with declarative *filtering agents (interpreters* and *expressors)* allow users to tailor generic interaction paradigms for modeling purposes.[98] By combining one or more interaction paradigms with a semantics and problem-solving method defined in the filtering agents, a methodology-specific or application-specific mediating representation may be created. Users define mappings between symbols and actions in the interaction paradigms, and operations on logical entities, relationships, and properties in the intermediate representation. As shown in Figure 9, the same interaction paradigm may be used to display and operate on different aspects of the concept model, and the similar aspects of the model may be edited using different interaction paradigms. For example, influence diagrams combine a graph view with the concepts of decision, chance, and value nodes and the problem-solving method of maximization of expected utility across decision alternatives. Trade study matrices (a methodology-specific kind of repertory grid) are built out of a matrix view, the concepts of alternatives, criteria, and ratings, and a heuristic classification problem-solving method. Process views combine a graph view with the a formal definition of activities and relationships between them. Type definition views allow the users to extend the built-in ontology.[98] Configured with semantic information, these mediating representations operate in *modeling mode*, portraying different perspectives on the formal concept model in the intermediate representation. By virtue of the Smalltalk-80 model-view-controller paradigm, consistency is maintained among the model views.

Virtual notebook. The volume and diversity of information that can be represented in DDUCKS drives a requirement for ways to manage, organize, and link that information. A *virtual notebook* facility helps collaborating individuals collect and organize the diverse materials associated with a particular knowledge acquisition project. It also helps manage changes between different versions and views of the model as it evolves. A new notebook is typically opened in “double-page” mode, displaying a page on the right and one on the left as in a paper notebook. The left page typically contains a table of contents view listing the set of pages available in the notebook. The right page might contain a representation for some portion of the knowledge base. Users move from page to page by selecting a “tab” on the side of the notebook or selecting an item in the table of contents view. Alternatively, the user can query the notebook to bring up pages meeting user-defined criteria. Figure 7 shows a virtual notebook in single-page mode.

Using *templates*, groups can tailor the contents of the boiler-plate virtual notebook to be consistent with their preferences for accessing, viewing, and using the information. For example, a knowledge engineering team's blank notebook can come preconfigured with information about organizational standards (e.g., concept and method libraries, reporting forms) and procedures (e.g., required steps in a project plan), just as a real notebook could be preloaded with labeled dividers and forms. Besides its obvious use in managing information about the model, the virtual notebook supports the team as a simple computer-supported meeting facilitation tool and as a form of group memory.

3.1.2. ICONKAT

ICONKAT (Integrated Constructivist Knowledge Acquisition Tool) is a knowledge acquisition and representation system under evolutionary development at the University of West Florida. It incorporates principles and techniques from both personal construct theory and assimilation theory. ICONKAT provides extensive interactive assistance to the domain expert and knowledge engineer in cooperatively modeling expertise. Like DDUCKS, ICONKAT is based on a three-schemata knowledge representation approach. The conceptual domain model is constructed within the framework provided by ICONKAT's mediating and intermediate representations—providing direct support for model creation, documentation, maintenance, knowledge base generation, and the resulting expert system's explanation facility.

ICONKAT's collaborative modeling environment exploits the expressiveness of concept maps to assist users in hierarchically organizing the various mediating representations (e.g., other concept maps, repertory grids, images, audio, video, documents) into browseable hypermedia domain models (see Figure 11). Interestingly, concept maps play a twin role in this process. First, concept maps are one of the principal means by which the expert and knowledge engineer represent knowledge about the domain. In particular, concept maps have proven effective in eliciting and representing what the participants see as the knowledge landscape or topology at a given level of abstraction. Second, concept maps furnish a rich organizational framework that can serve as the interface to the domain model. Thus, while the expert and knowledge engineer collaborate in using concept maps to model the former's problem-solving knowledge, they are also, in essence, building the structure of the interface that subsequent users will employ to explore the model.

The knowledge acquisition process and its ramifications do not culminate with deployment of the system, but rather extend throughout its useful life. Accordingly, in addition to a flexible modeling environment, ICONKAT has been designed with the complete knowledge-based system life cycle in mind. In particular, ICONKAT supports a new explanation paradigm, in which, the domain model that emerges from the knowledge acquisition process is subsequently exported from the development environment to the delivery environment, where it serves as the foundation of the explanation capability for the deployed system.

ICONKAT was used in the design and construction of NUCES: Nuclear Cardiology Expert System.[123] This is a large-scale expert system for the diagnosis of first pass cardiac functional images, a noninvasive radionuclide technique used to evaluate heart wall motion abnormalities.

Mediating representations. ICONKAT's mediating representations are designed to promote communication and understanding between the human participants in the knowledge acquisition process. A good mediating representation fosters the constructive modeling processes (e.g., meaning making and meaning sharing) by empowering domain experts and knowledge engineers to cooperatively build models of expert knowledge. Furthermore, mediating representations may facilitate explanation (see discussion below) by enabling the system's eventual users to explore the conceptual domain model without resorting to low-level representations (e.g., C code, lisp, rules).

ICONKAT's principal mediating representations are the concept map and the repertory grid. It uses these complementary mediating representations synergistically. In ICONKAT, concept maps depict the conceptual relationships of the domain as constructed during the knowledge acquisition process. For example, the concept map in Figure 2 expresses relationships among ejection fraction (a critical numerical value), other manifestations of heart wall image abnormalities (e.g., "blue fingers"), specific heart diseases (e.g., ischemia), and human physiology. The relevant disease states appear at the lowest levels of the map, and were incorporated into the expert's repertory grid (Figure 1). Note that the map includes the domain expert's personally constructed expertise in the form of visual analogies that he employs as markers for perceived image abnormalities (e.g., "blue fingers," "bull's-eyes" and "ice cream cones"). These markers are the basis upon which the expert differentiates the various disease states, and were included as constructs in a repertory grid. In addition to concept maps and repertory grids, ICONKAT supports the use of a variety of other mediating representations, such as images, audio, Quicktime movies, and documents.

Intermediate representation. ICONKAT's intermediate representations perform an important integrative function and are an area of ongoing evolution, testing and revision. Although mediating representations have enhanced the richness and subtlety with which the human participants in the knowledge acquisition process can model the domain, the need for integrative intermediate representations has become increasingly apparent. For example, much of the work on ICONKAT's intermediate representations has focused on how they might enhance the relationship between repertory grids and concept maps. In addition, the kind of representations that ICONKAT provides as the basis of its modeling environment (i.e., informal, graphical and textual mediating representations) are designed for the benefit of humans, while implementation formalisms are focused on computational issues—causing a substantial semantic gap. ICONKAT's intermediate representations are designed to partially bridge this gap, thus enabling feedback, analysis, and verification throughout the entire process of system development. ICONKAT's intermediate representations (sometimes referred to as the 'glue') consist of a collection of modeling primitives implemented as abstract data types in C++.

Support for Explanation.

A machine that incorporates expert judgment in a given domain is more likely to find acceptance by those seeking its advice if it can explain its recommendations. Accordingly, an explanation capability should enable a user to get a complete, understandable answer to any sort of relevant question about the knowledge explicitly and *implicitly* embodied in a system's implementation formalism (e.g., rules, frames, or whatever). Unfortunately, the capacity of most current expert systems to explain their findings (i.e., conclusions) is limited to inadequate, causal descriptions of the behavior of the performance environment's reasoning mechanism. One key to the design of explanation subsystems that are capable of deeper and less mechanistic accounts is to recognize that the development of an explanation facility is a fundamental aspect of the knowledge acquisition process.[99]

Instead of arduously constructing a model of problem solving expertise, and then throwing it away (upon translation into the syntax of the performance environment), ICONKAT's explanation paradigm allows users to exploit the model formed during the knowledge acquisition process. As depicted in Figure 10, the model resulting from the knowledge elicitation process is exported from the development environment to the delivery environment. It serves there as the foundation of the explanation capability for the deployed system.

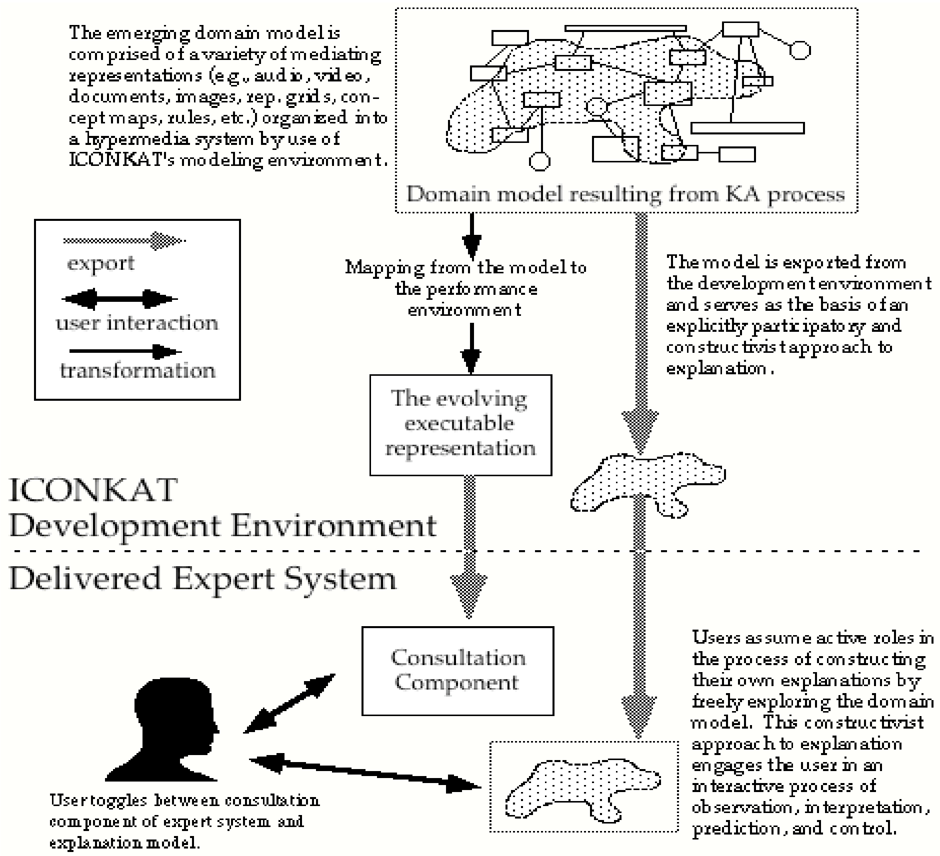

Figure 10. Transition of the domain model from development to the basis for explanation of the delivered system (adapted from figure in Ref. 24).

A session from NUCES (a medical expert system built in the ICONKAT environment) illustrates the ICONKAT approach to explanation (see Figure 11). When a user requests explanation, the performance environment is interrupted, and the user is switched into the context-sensitive explanation subsystem and conveyed to an appropriate location within the multidimensional space representing the model. From there, the user can assume an active role in the process of constructing his or her own explanation by freely exploring the conceptual model and browsing among a wealth of supporting objects (e.g., audio, video, documents, images, repertory grids, concept maps, rules, etc.). Users end their browsing as soon as they are confident that they have constructed an adequate explanation from the available information. This constructivist approach to explanation engages the user in an interactive process of observation, interpretation, prediction, and control.

< INSERT FIGURE 11 ABOUT HERE >

Figure 11. A NUCES session illustrating the notion of participatory explanation.

The navigation problem, an important concern in hypermedia systems, is largely ameliorated by use of the use concept maps as a guide to traversing the logical linkages among clusters of related objects (see the ‘Concept Map’ window in Figure 11). Concept maps provide an elegant, easily understood interface to the domain model. A system of concept maps is interrelated by generalization and specialization relationships among concepts, which lead to a hierarchical organization. The explanation subsystem provides a window that shows the hierarchical ordering of the various maps, highlights the current location of the user in the hierarchy, and permits movement to any other map by clicking on the desired map in the hierarchy (see the window ‘Concept Map Hierarchy’ in Figure 11).

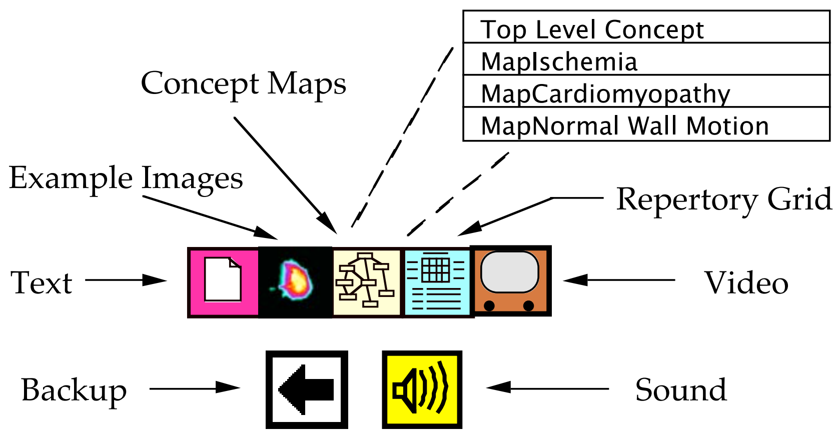

Figure 12. Close-up of the icons found at each node of the concept maps.

Depending on the location of the user in the domain model, he or she has different options to explore. At each node, the user can select from a menu of icons as shown in Figure 12. These correspond to *text* (a textual document), *images*, a popup menu of *concept maps*, *repertory grids* or *video* (implemented using Quicktime) related to the topic of the selected node. These icons will appear in various combinations depending on what information is available for a given concept. The ‘Concept Map’ window in Figure 11 shows how the concepts (nodes) are populated with the icon menus illustrated in Figure 12. At any time, the user can backtrack by clicking on the ‘back-arrow’ icon, as shown in the ‘Concept Map’ window. This scheme provides the user great flexibility in navigating through related concepts, as well as, guideposts in moving among the various sources of information available for a specific concept.

3.1.3. KSSn/KRS

KSSn (Knowledge Support System) is an ongoing experiment in the development of knowledge acquisition tools which incorporates aspects of personal construct psychology. KSSn[100] is designed around a knowledge representation server (KRS) implemented in C++, providing services based on those of KL-ONE/CLASSIC[40,41] augmented with inverse roles, data types for integers, reals, strings and dates, and with rule representation that allows one rule to be declared an exception to others. The server supports the operations of intensional logic, and one of the modules attached to it is a graphic knowledge editor supporting the associated visual language.[54] While KSSn provides export facilities to expert system shells, the KL-ONE/CLASSIC inference capabilities of the server allow the system to be used as a complete problem solving environment. For example, KSSn has been used on a room allocation problem[36] derived from an ESPRIT project[124] that was placed in the public domain as part of Project Sisyphus.[125]

Figure 13 shows the architecture of KSSn as a family of modules attached to the knowledge representation server, KRS. The description of the system is taken directly from Gaines.[36]

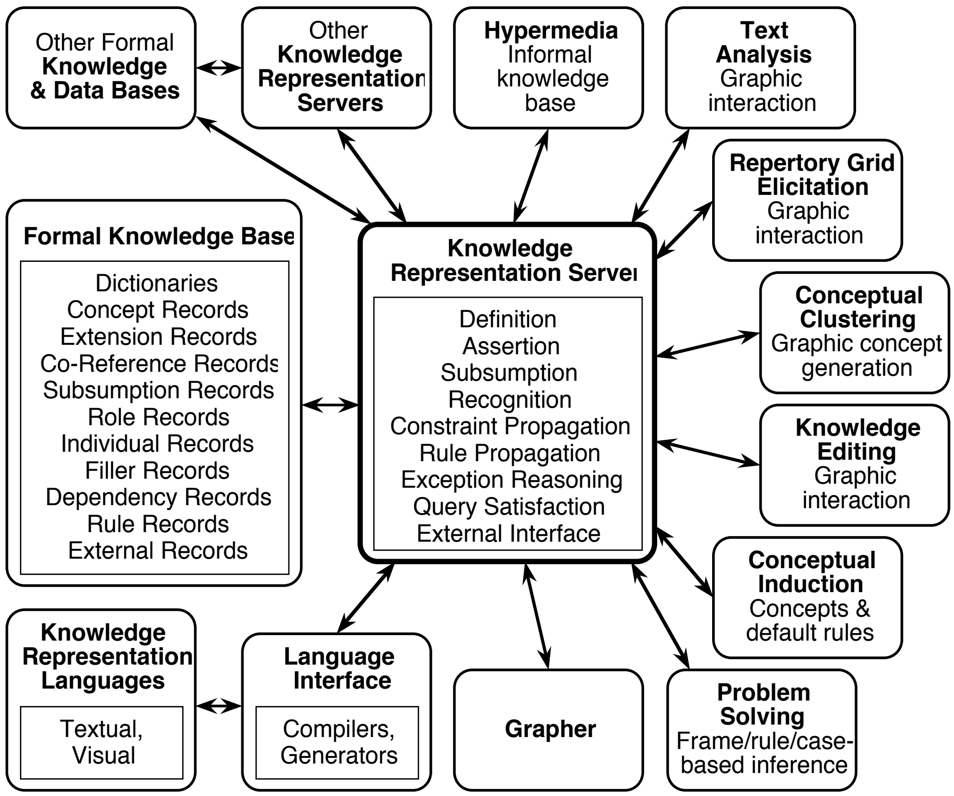

Figure 13. Architecture of KRS (taken from Gaines[36]).

The modules are (reading clockwise from the top left):

- Interface modules to other knowledge bases and servers, including databases.
- A hypermedia module allowing informal knowledge structures in text and images to be captured, accessed and linked. The linkage structure itself is held as a knowledge base. Domain-specific tools may be developed in HyperCard and existing knowledge acquisition tools in HyperCard may be integrated, for example Woodward's[126] Cognosys for the analysis of protocol data in textual form.
- A text analysis module allowing documents to be analyzed in terms of word usage, and associations between significant words to be graphed—based on TEXAN in KSS0. This enables protocols and technical documents to be used to initiate knowledge acquisition.
- A repertory grid expertise transfer module allowing graphic definition of concepts and graphic creation and editing of individuals—based on the elicitation screens of KSS0.
- A conceptual clustering module allowing interactive definition of new concepts—based on the hierarchical and spatial clustering from KSS0.

- A knowledge editing module allowing the interactive development and editing of knowledge structures through a visual language.
- A conceptual induction module creating rules about specified subsets of individuals and transforming them to a minimal set of concepts and default rules—based on the INDUCT algorithm.
- A problem solving module supporting frame, rule and case-based inference from the knowledge structures.
- A grapher laying out specified parts of the concept subsumption graph, concept structures and individual structures—based on an incremental layout algorithm that can be used interactively to support the production of clear visual knowledge structures.
- A language interface accepting and generating definitions and assertions in formal knowledge representation languages, both textual and visual.

The knowledge representation services of KRS, the central server module, correspond to those of CLASSIC,[41] augmented with inverse roles, data types for integers, reals, strings and dates, and with rule representation that allows one rule to be declared an exception to others. For the purposes of this paper KRS may be seen as providing a fast and principled implementation of a frame/rule knowledge representation and inference engine capable of operating with large knowledge bases.

KDRAW Visual Language. An important component of KSSn is the graphic knowledge editor, KDraw. This is a drawing tool designed for ease of use that provides a visual structure editor for semantic networks representing classes, objects and rules in KRS. Nosek and Roth[127] have demonstrated empirically that the visual presentation of knowledge structures as semantic nets leads to more effective human understanding than does textual presentation of the same structures. Gaines[54] has developed a formal visual language that corresponds exactly to the underlying algebraic semantics of KRS (see Figures 4 and 5). It has remarkably few visual primitives and is easily learned and understood.

The KDraw design defines the visual syntax and underlying semantics of a visual language for term subsumption knowledge representation languages in the KL-ONE family. It focuses on the use of the language to enter and edit knowledge visually, and on its application in a highly interactive graphic structure editor. However, the language is also well-suited to the display of knowledge structures, and the system includes a grapher using Watanabe's[128] heuristics.

The editor is modeled on Apple's MacDraw with additional features appropriate to the language such as arcs remaining attached to nodes when they are dragged. The syntax of possible node interconnections and constraint expressions is enforced—it is not possible to enter a graph that is syntactically incorrect. Cut-and-paste of graphs and subgraphs is supported, and pop-up menus allow nodes to be connected with the minimum of effort. Updates are efficient and graphs with over a thousand nodes can be manipulated interactively. Scroll, zoom and fit-to-size facilities allow large data structures to be navigated easily.

The grapher interface to the knowledge representation server allows the knowledge structures to be used deductively to solve problems and give advice. Other programs such as HyperCard can also access the server and provide additional functionality such as customizable end-user interfaces. Repertory grid data and induced rules, elicited and analyzed through the KSS0-style modules, may be exported to the grapher for visual analysis and editing.

3.2. Analysis and Induction Techniques

Statistical procedures implemented in general-purpose repertory grid tools such as PLANET[92,93] and OMNIGRID[129] have been used for many years to explore interesting relationships among elements and constructs.[130] Such analysis techniques have included information measures, nonparametric factor analysis, conventional factor analysis, principal component analysis, multidimensional scaling, and hierarchical cluster and linkage analyses, among others.[16,27] There also have been some attempts to construct precise mathematical models of the cognitive processes reflected in grid data and to use these models for both simulating the performance of hypothetical respondents and predicting the responses of real ones.[131]

KSS0 contain facilities for eliciting distinctions from text input or protocols.[76,77] Text input from a book or a set of protocols may be analyzed through a procedure which clusters associated words and renders them as knowledge structures in the KRS visual language. The text is fully indexed by all non-noise words grouped by their stems, and a coupling matrix of word associations is calculated using a simple distance-in-text measure. This leads to a schema from which the expert can select related elements and initial constructs with which to commence grid elicitation or semantic network construction. In the longer term the text analysis system could be extended with the more powerful parsing and semantic analysis techniques now being developed for knowledge bases.[132]

Induction techniques in ETS,[52] KSS0,[133] and *Nicod*[32] were originally developed to create rules from repertory grids for export to commercial expert system shells. A brief discussion of the logic of confirmation method by which both *Nicod* and ICONKAT derive rules from repertory grid data is given in Section 2.3.1. Induction techniques based on information theory have also become more widely available.[134] For example, in KSSn/KRS, Gaines[135] has developed INDUCT, a set of empirical induction techniques that derive potential implications between concepts based on noisy datasets or repertory grid information. The system is an extension of Cendrowska's[136] PRISM algorithm, augmented by the capability of generating more compact and intuitive rule sets that include explicit exceptions.[38] Facilities in KSSn/KRS hierarchical concept structures to be derived directly from these rules. In their work on *Axotl* and *Aquinas,* Bradshaw and Boose explored the induction of Bayesian graphical models from repertory grids and databases. This work is being continued as part of the medical application of DDUCKS.[137,138]

ICONKAT contains a novel approach supporting elicitation of superordinate constructs through the use of neural nets. GridNet uses the expert's preliminary repertory grid data as input to a self-organizing, multilevel, artificial neural net. The net uses back propagation to identify abstractions ("hidden features") taken from this nonlinear hierarchy. These artificial neural net abstractions are then fed back to the expert as element clusters for the elicitation of new superordinate constructs. In their nuclear cardiology application, Ford and his colleagues were successful in discovering high-level, pivotal constructs as a result of using GridNet.[24]

3.3. Multiple Expert Analysis and Group Use of Tools

Two of the most active areas of knowledge acquisition research concern multiple expert analysis and group use of tools. Theoretical as well as practical concerns also make this one of the more controversial areas. We present a summary of work in this area by each group separately below.

3.3.1. KSSn/KRS

Shaw[139] developed an approach to account for the psychological process not only of individual people but also for that of functional groups. Gaines and Shaw[93] used this point of view to develop several techniques to compare and contrast repertory grids obtained from different individuals. They employ *exchange grids* for the measurement of understanding and agreement between either two people or two occasions. Another procedure produces a set of *socionets* which indicates the links of similar construing within the group, and a *mode grid* showing the dimensions which are readily understood by the majority of the group. Information from analyses can be used to establish consensus about terminology and distinctions.[140]

Besides the analysis and use of knowledge from multiple experts, the integration of tools in KSSn/KRS has provided the means for the development of special-purpose 'groupware' applications, based in part on work from previous efforts.[91] One of these, *RepGridNet,* supports integration of repertory grids and socioanalysis tools with an electronic mail subsystem to facilitate the formation and management of 'special interest groups.'[141] A second effort has developed *GroupWriter*, a word processor that supports collaborative authoring of scientific papers.[142] *GroupWriter* provides a hypertext-style linkage structure to document components it is readily extended to provide hypertext links to modules within a knowledge acquisition system. For example, source documents and transcripts of interviews and protocols may be linked to the knowledge structures that have been developed based on them. This can be done without undermining the visual appearance of the original document. Documents generated as part of the knowledge engineering process can also be treated in this way—they are normal documents with full typographic formatting and graphics, but they are also tightly embedded in the knowledge structures being developed.

3.3.2. ICONKAT

There are as many opinions as there are experts
— Franklin Delano Roosevelt

Ford et al.[143,144] have elaborated a personally constructed and socially situated view of expertise that helps us understand the problems that often arise in conjunction with the elicitation and representation of expertise from multiple domain experts. From this perspective, knowledge can be viewed as functional but fallible representations not of reality writ large but of experience. In this sense, an expert is perceived to possess more functional representations than non-experts. For example, certain physicians are deemed to be 'experts' not necessarily because they possess more valid medical information than their colleagues, but rather because they are perceived to be experts (for a variety of reasons) by their medical constituency. The expert's representations or procedures need not be valid, in a rational-empirical sense, they need only be functional in helping the constituencies manage their uncertainty, just, for example, as all kinds of 'invalid' past medical practice (when seen from the vantage of current medical belief) have done. It follows that the expertise does not reside in the expert per se but in the expert-in-context. In brief, expertise is socially situated. Not only have we lost an external (reality) reference for expertise, but we have lost an individual reference as well. The minimum unit of analysis is not the individual expert, but rather is the expert in context with his or her constituency.

Eliciting, representing, and usefully coalescing the personally constructed and context dependent knowledge of several experts is a daunting task. Of course, it is not difficult to elicit all kinds of standard information (i.e., widely shared consensual beliefs) from several experts. However, it is much more difficult to elicit and represent the personally constructed experiential knowledge that accounts for each of them having a constituency (the minimum requirement for holding expert status in the first place). Further, much of this knowledge is social and/or political in nature and all of it is context dependent with respect to its usefulness. If, as we propose, expertise is personally constructed and context dependent, then any effort to employ multiple domain experts must also elicit and represent the various contexts in which they operate.

If on the other hand, the knowledge engineer assumes that there exists some 'gold standard' of knowledge and that experts each possess various parts of this existent knowledge (i.e., REALITY), they might be tempted to follow the naive strategy, "that if one expert is good, then two are better." However, this is usually a mistake. In contrast, it has been observed that in some cases, co-mingling the domain models of multiple experts tends to cause a "regression to the mean," and that the resulting system is 'less expert' than either individual.[143] In addition, Ford and Adams-Webber[143] have noted on several occasions that the more successful the knowledge acquisition process has been in modeling a particular expert's most relevant functional abstractions—the more difficult it will be to add another expert to the emerging and typically idiosyncratic domain model. In the somewhat unusual circumstance in which the knowledge engineer has multiple bona fide domain experts at his/her disposal, we posit that it is usually preferable to build a separate knowledge base for each expert rather than attempting to mingle their expertise in a single unified knowledge base.

This is not to say that there are no circumstances which warrant the use of multiple domain experts, but rather to counsel caution in their application. In fact, ICONKAT's collaborative modeling environment is being adapted for application in a collaborative learning project with children of several countries.[145] Collaborative learning is construed as an enterprise in which the learners, and perhaps their teachers, cooperatively build an explicit knowledge model. There are strong analogies between this collaborative learning project and the ICONKAT approach to knowledge acquisition with multiple domain experts.

3.3.3. DDUCKS

Aquinas incorporated some of the multiple expert analysis techniques developed by Shaw and Gaines and used them to guide negotiation among experts.[146,147] *Aquinas* also added the significant new feature of allowing consultation users to review the results from multiple experts simultaneously.[148,149] The reasoning engine used results from the experts to display dissenting opinions (i.e., the set of consultation results that was most different from the rest). These notions of "running the experts in parallel" (i.e., independent expert systems) and of presenting dissenting opinions seem to be useful in some situations. We plan over time to incorporate some of these features in DDUCKS. However, the cautions offered above and in Ford et al.[143,144] should be kept in mind.

In addition to the multiple expert analysis techniques discussed above, an effort is underway to adapt the DDUCKS environment for group use. Boose et al.[106] describe a comprehensive decision model for group decision support systems, based on an analysis of commercially-available electronic meeting support systems[150] and their experience with automated knowledge acquisition tools. The integrated model combines current brainstorming-oriented methods,[151] structured text argumentation,[152,153] repertory grids, possibility tables, and influence diagrams. Each component will address weaknesses in current group decision support systems for certain types of problems.

4. Conclusions: The Future of Personal Construct Theory and Knowledge Acquisition

In conclusion, it is interesting to characterize the evolution of architecture for systems developed by The Boeing Company, University of West Florida, and University of Calgary groups in terms of four stages:

1. *Era of the single-approach rapid-prototyping tools*. The earliest predecessors of these tools (i.e., ETS, *Nicod*, PLANET) were based on repertory grid interviewing techniques, representations, and analysis tools. They were generally used for rapid-prototyping of classification problems, following which, rules were exported to a disk file for use by a commercial expert system shell.

2. *Era of monolithic integration*. Integration became the theme as knowledge acquisition workbenches (e.g., *Aquinas* and KSS0) incorporated several additional tools and representations. Within *Aquinas*, export to external shells was de-emphasized as internal problem solving capabilities increased and became difficult to replicate in traditional shells.

3. *Era of decoupling and interapplication communication.* This stage represents the current state-of-the-art. Developers of ICONKAT, KSSn/KRS, and DDUCKS decouple components of their systems to allow integration with tools springing from complementary theoretical perspectives (e.g., concept maps, neural networks, influence diagrams, possibility tables, semantic networks) and to exploit emerging operating-system-based interapplication communication protocols. Two-way communication with commercial (e.g., HyperCard, Excel, Nexpert, databases) and internally-developed (e.g., DART, Babylon) applications is established.[104] Easy tailoring of the user-interface is made possible by the adoption of three-schemata knowledge representation approaches[107,110] and more general configurability allowing multiple problem-solving approaches is made possible through the use of PROTEGE II and KADS-like architectures.[108,109,154,155] A better understanding of expertise and modeling is informing the design of knowledge acquisition tools.[81,110,143,144,156,157] Knowledge acquisition tools are beginning to target wider applications such as information retrieval, education, personal development, group decision support, and design rationale support.

4. *Era of knowledge sharing, adaptivity, and intelligent interoperability*. Over time, we expect to see systems migrate to a state of "intelligent interoperability," that is "intelligent cooperation among systems to optimally achieve specified goals."[158] In such a scheme, nearly all communication between systems will be on a peer-to-peer basis, with transparent mechanisms preserving semantics across diverse representations and ontological frameworks.[159-165] Computing will become "ubiquitous" as microprocessors become an integral part of office and home surroundings and as inexpensive laptop and palmtop systems with wireless communication systems proliferate.[166-168] System components will become even more modular and finely-grained.[169,170] Radical tailorability, embeddability, and reuse of components will be possible through sophisticated object-management and end-user-oriented configuration environments.[171-176] A wider variety of diagrammatic and pictorial representations will be available.[177,178] A final step toward intelligent interoperability would be to embed agent programs within each cooperating system, which include primitives for communicating and sharing resources with other agents.[179-182] Sophisticated intelligent assistant and information retrieval systems capable of learning and using context-specific indices will be based on such capabilities.[183-185]

Over time, the convergence of advanced documentation and knowledge representation tools will culminate in the development of a wide variety of *knowledge media,* computational environments in which explicitly represented knowledge serves as a means of communication among people and their programs.[186,187] Adaptive document systems will tailor the content and presentation of multimedia produced to the interests, abilities, and situation of the user, fully integrating the functions of design documentation, operations manuals, and computer-based training tools.[188,189] They will transcend the artificial boundaries between documentation and modeling tools, freely incorporating not only static drawings and graphical animations as illustrations, but also simulations and full working models.

As new needs and opportunities arise, the roots of personal construct theory continue to nourish vigorous growth and development in knowledge acquisition research. We look forward to the new era of work ahead, and see a bright future for modeling tools based on constructivist theory.

Acknowledgments

We express our appreciation to Neil Agnew, Russell Almond, Miroslav Benda, Guy Boy, Kathleen Bradshaw, John Brennan, Alberto Cañas, John Coffey, Stan Covington, Jim Fulton, Pete and Cindy Holm, Sam Holtzman, Ron Howard, Earl Hunt, Jeremy Jones, Oscar and Sharon Kipersztok, Cathy Kitto, Joe Koszarek, Janusz Kowalik, Tim Lethbridge, Johnny Liseth, David Madigan, Jim Matheson, Allen Matsumoto, Thom Nguyen, Joseph Novak, Steve Poltrock, Dave Purdon, Bob Schneble, Doug Schuler, Kish Sharma, Dave Shema, Doug Skuce, Dan Small, Howard Stahl, Bruce Wilson, Jeff Yerkes, and Debra Zarley for their contributions and support. This work has benefited from discussions with numerous colleagues in the knowledge acquisition community over the years, in particular Brian Gaines, Mildred Shaw and Brian Woodward who provided materials used in this paper to describe their work.

References

1. G.A. Kelly, *The Psychology of Personal Constructs*, Norton, New York, 1955.
2. G.A. Kelly, “A mathematical approach to psychology,” in *Clinical Psychology and Personality: The Selected Papers of George Kelly,* B.A. Maher (Ed.), Wiley, New York, 1969, pp. 7-45.
3. G.A. Kelly, “A brief introduction to personal construct theory,” in *Perspectives in Personal Construct Theory,* D. Bannister (Ed.), Academic Press, London, 1970, pp. 1-29.
4. J.R. Adams-Webber, “Kelly’s pragmatic constructivism,” *Canadian Psychology*, **30**, 190-193 (1989).
5. J.R. Adams-Webber, “Personal construct theory and cognitive science,” *International Journal of Personal Construct Psychology*, **3**, 415-421 (1990).
6. N.M. Agnew and J.L. Brown, “Foundations for a theory of knowing: I. Construing reality,” *Canadian Psychology,* **30**, 152-167 (1989).

7. N.M. Agnew and J.L. Brown, “Foundations for a theory of knowing: II. Fallible but functional knowledge,” *Canadian Psychology*, **30**, 168-183 (1989).
8. J.C. Mancuso and B.N. Eimer, “Fitting things into sorts: The range corollary,” in *The Construing Person,* J.C. Mancuso and J.R. Adams-Webber (Eds.), Praeger, New York, 1982, pp. 130-151.
9. G.A. Kelly, “Ontological acceleration,” in *Clinical Psychology and Personality: The Selected Papers of George Kelly,* B. Maher (Ed.), Wiley, New York, 1966, p. 129.
10. T. Mischel, “Personal constructs, rules, and the logic of clinical activity,” *Psychological Review,* **71**, 180-192 (1964).
11. A. Einstein, “Physics and Reality,” in *J. Bartlett’s Familiar Quotations,* E.M. Beck (Ed.), Little, Brown and Company, Boston, 1936, p. 763.
12. G.A. Kelly, “Discussion: Aldous, the personable computer,” in *Computer Simulation of Personality,* S.S. Tomkins and S. Messick (Eds.), Wiley, New York, 1963, pp. 225-226.
13. K.M. Ford, “A constructivist view of the frame problem in artificial intelligence,” *Canadian Psychology*, **30**, 188-190 (1989).
14. J.R. Adams-Webber and J.C. Mancuso (Eds.), “The pragmatic logic of personal construct psychology,” in *Applications of Personal Construct Psychology,* Academic Press, New York, 1983, pp. 1-10.
15. R.A. Neimeyer, *The Development of Personal Construct Psychology*, University of Nebraska Press, Lincoln, 1985.
16. J.R. Adams-Webber, *Personal Construct Theory: Concepts and Applications*, Wiley, New York, 1979.
17. J.R. Adams-Webber, “Actual structure and potential chaos,” in *Perspectives in Personal Construct Theory,* D. Bannister (Ed.), Academic Press, London, 1970, pp. 30-45.
18. L.S. Vygotsky, *Language and Thought*, MIT Press, Cambridge, MA, 1962.
19. C.E. Osgood, G. Suci, and P.H. Tannenbaum, *The Measurement of Meaning*, University of Illinois Press, Urbana, 1957.
20. J.R. Adams-Webber, “Repertory grid technique,” in *Encyclopedia of Psychology,* R. Corsini (Ed.), John Wiley Interscience, New York, 1987, pp. 973-974.
21. B. Adams and J.R. Adams-Webber, “Testing a computational model of interpersonal judgment with personal constructs,” *International Journal of Personal Construct Psychology,* **5**(2), 145-157 (1992).
22. J.R. Adams-Webber, “Self-other contrast and the development of personal constructs,” *Canadian Journal of Behavioural Science,* **17**, 303-314 (1985).

23. J.R. Adams-Webber, "Construing self and others," in *Anticipating Personal Construct Theory,* F. Epting and A.W. Landfield (Eds.), University of Nebraska Press, Lincoln, 1985, pp. 58-69.
24. K.M. Ford, H. Stahl, J.R. Adams-Webber, A.J. Cañas, J. Novak, and J.C. Jones, "ICONKAT: An integrated constructivist knowledge acquisition tool," *Knowledge Acquisition,* **3**(2), 215-236 (1991).
25. F. Fransella and D. Bannister, *A Manual for Repertory Grid Technique*, Academic Press, London, 1977.
26. D.N. Hinkle, *The Change of Personal Constructs from the Viewpoint of a Theory of Implications*, Doctoral Dissertation, Ohio State University, 1965.
27. M.L.G. Shaw (Ed.), *Recent Advances in Personal Construct Technology*, Academic Press, New York, 1981.
28. P.P. Chen (Ed.), *Entity-Relationship Approach to Systems Analysis and Design*, North Holland, New York, 1980.
29. J.M. Bradshaw, S.P. Covington, P.J. Russo, and J.H. Boose, "Knowledge acquisition for intelligent decision systems: Integrating *Aquinas* and *Axotl* in DDUCKS," in *Uncertainty in Artificial Intelligence,* M. Henrion et al. (Eds.), Elsevier, Amsterdam, 1990, pp. 255-270.
30. J.M. Bradshaw, S.P. Covington, P.J. Russo, and J.H. Boose, "Knowledge acquisition techniques for decision analysis using *Axotl* and *Aquinas*," *Knowledge Acquisition,* **3**(1), 49-77 (1991).
31. A.J. Cañas, K.M. Ford, D.B. Crouch, and K.V. Vasani, "Personalizing databases with knowledge acquisition tools," *Proceedings of Fifth Annual Florida AI Research Symposium* (FLAIRS-92), Ft. Lauderdale, 1992, pp. 240-242.
32. K.M. Ford, F.E. Petry, J.R. Adams-Webber, and P.J. Chang, "An approach to knowledge acquisition based on the structure of personal construct systems," *IEEE Transactions on Knowledge and Data Engineering,* **3**, 78-87 (1991).
33. B.J. Cragun and H.J. Steudel, "A decision-table-based processor for checking completeness and consistency in rule-based expert systems," *International Journal of Man-Machine Studies,* **26**, 633-648 (1987).
34. J.D. Novak, *A Theory of Education*, Cornell University Press, Ithaca, NY, 1977.
35. J.H. Boose and J.M. Bradshaw, "Expertise transfer and complex problems: Using *Aquinas* as a knowledge-acquisition workbench for knowledge-based systems," *International Journal of Man-Machine Studies,* **26**, 3-28 (1987).
36. B.R. Gaines, "Organizational modeling and problem solving using an object-oriented knowledge representation server and visual language," *Proceedings of the Conference on Organizational Computing Systems*, 1991, pp. 80-94, ACM Press.

37. B.R. Gaines, “An interactive visual language for term subsumption visual languages,” *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, San Mateo, CA, 1991.
38. B.R. Gaines, “Integrating rules in term subsumption knowledge representation servers,” *Proceedings of the Ninth National Conference on Artificial Intelligence*, Menlo Park, CA, 1991.
39. B.R. Gaines, “Empirical investigation of knowledge representation servers: Design issues and applications experience with KRS,” *SIGART Bulletin,* **2**(3), 45-56 (1991).
40. R.J. Brachman and J. Schmolze, “An overview of the KL-ONE knowledge representation system,” *Cognitive Science,* **9**(2), 171-216 (1985).
41. A. Borgida, R.J. Brachman, D.L. McGuinness, and L.A. Resnick, “CLASSIC: A structural data model for objects,” *ACM SIGMOD International Conference on the Management of Data*, ACM Press, 1989.
42. K.M. Ford, *An Approach to the Automated Acquisition of Production Rules from Repertory Grid Data*, Ph.D. Dissertation, Tulane University, 1987.
43. K.M. Ford and J.R. Adams-Webber, “The structure of personal construct systems and the logic of confirmation,” *International Journal of Personal Construct Psychology,* **4**, 15-41 (1991).
44. J.M. Bradshaw and J.H. Boose, “Decision analysis techniques for knowledge acquisition: Combining information and preferences using *Aquinas* and *Axotl*,” in *Progress in Knowledge Acquisition for Knowledge-Based Systems,* J.H. Boose and B.R. Gaines (Eds.), Academic Press, London, 1990.
45. J.M. Bradshaw and J.H. Boose, “Decision analysis techniques for knowledge acquisition: Combining information and preferences using *Aquinas* and *Axotl*,” *International Journal of Man-Machine Studies,* **32**(2), 121-186 (1990).
46. J.H. Boose, J.M. Bradshaw, and D.B. Shema, “Knowledge-based design rationale capture: Automating engineering trade studies,” in *Knowledge-Aided Design,* M. Green (Ed.), Academic Press, London, 1992, pp. 197-214.
47. D.P. Ausubel, *The Psychology of Meaningful Verbal Learning*, Grune & Stratton, New York, 1963.
48. D.P. Ausubel, J.D. Novak, and H. Hanesian, *Educational Psychology: A Cognitive View*, Second Ed., 1978.
49. J.D. Novak and D.B. Gowin, *Learning How to Learn*, Cambridge University Press, New York, 1984.
50. D.E. Snyder, M.D. McNeese, B.S. Zaff, and M. Gomes, “Knowledge acquisition of tactical air-to-ground mission information using concept mapping,” *Proceedings of the AAAI Cognitive*

Aspects of Knowledge Acquisition Session of the Spring Symposium, Stanford, CA, March, 1992, pp. 228-234.

51. J.H. Boose, “Personal construct theory and the transfer of human expertise,” *Proceedings of the National Conference on Artificial Intelligence* (AAAI-84), Morgan Kaufmann, San Mateo, CA, 1984, pp. 27-33.
52. J.H. Boose, *Expertise Transfer for Expert System Design*, Elsevier, New York, 1986.
53. J.H. Boose, D.B. Shema, and J.M. Bradshaw, “Recent progress in *Aquinas:* A knowledge acquisition workbench,” *Knowledge Acquisition,* **1**(2), 185-214 (1989).
54. B.R. Gaines and M.L.G. Shaw, “Cognitive and logical foundations of knowledge acquisition,” *Proceedings of the Fifth Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Department of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1990, pp. 9.1-9.25.
55. C.G. Hempel (Ed.), “Studies in the logic of confirmation,” in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science,* Free Press, New York, 1965, pp. 3-51.
56. G.H. Von Wright, “The paradoxes of confirmation,” in *Aspects of Inductive Logic,* J. Hintikka and P. Suppes (Eds.), North-Holland, Amsterdam, 1966, pp. 208-218.
57. W.C. Salmon, “Confirmation,” *Scientific American,* **228**, 75-83 (1973).
58. A. Bundy, “Incidence calculus: A mechanism for probabilistic reasoning,” *Journal of Automated Reasoning,* **1**, 263-283 (1985).
59. W.V. Quine, *Methods of Logic*, Harvard University Press, Cambridge, MA, 1982.
60. C.P. Langlotz, E.H. Shortliffe, and L.M. Fagan, “Using decision theory to justify heuristics,” *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, 1986, pp. 215-219.
61. W.J. Clancey, “The epistemology of a rule-based expert system—A framework for explanation,” *Artificial Intelligence,* **20**, 215-251 (1983).
62. M. Henrion and D.R. Cooley, “An experimental comparison of knowledge engineering for expert systems and for decision analysis,” *Proceedings AAAI-87 Sixth National Conference on Artificial Intelligence*, Seattle, Washington, 1987, pp. 471-476.
63. M. Henrion, “Practical issues in constructing a Bayes’ belief network,” *Proceedings of the Third AAAI Workshop on Uncertainty*, Seattle, Washington, July 10-12, 1987, pp. 132-139.
64. M. Henrion, J.S. Breese, and E.J. Horvitz, “Decision analysis and expert systems,” *AI Magazine*, **12**(4), 64-91 (Winter, 1991).
65. S.H. Holtzman, *Intelligent Decision Systems*, Addison-Wesley, Reading, MA, 1989.

66. R.L. Keeney, *Value-driven Expert Systems For Decision Support*, University of Southern California, Systems Science Department, 1986.
67. A.C. Miller, M.W. Merkhofer, R.A. Howard, and T.R. Rice, *Development of Automated Aids for Decision Analysis*, Stanford Research Institute, 1976.
68. R.A. Howard and J.E. Matheson (Eds.), "Influence diagrams," in *Readings on the Principles and Applications of Decision Analysis*, Strategic Decisions Group, Menlo Park, CA, 1984, pp. 719-762.
69. D. Heckerman, *Probabilistic Similarity Networks*, ACM Doctoral Dissertation Award Series, MIT Press, Cambridge, MA, 1991.
70. D. Geiger and D. Heckerman, "Advances in probabilistic reasoning," in *Proceedings of the Seventh Uncertainty in Artificial Intelligence Conference,* B.D. D'Ambrosio, P. Smets, and P. Bonissone (Eds.), Morgan Kaufmann, San Mateo, CA, 1991, pp. 118-126.
71. R.A. Howard, "Knowledge maps," *Management Science,* **35**, 903-922 (1988).
72. J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, CA, 1988.
73. C. Wiecha and M. Henrion, *Demaps: A Graphical Decision Support System*, Carnegie-Mellon University, Department of Engineering and Public Policy, 1987.
74. J.M. Bradshaw, C.R. Chapman, and K.M. Sullivan, "An application of DDUCKS to bone-marrow transplant patient support," *Working Notes of the AAAI 1992 Artificial Intelligence in Medicine* Session of the Spring Symposium, Stanford, CA, 1992, pp. 9-10.
75. B.R. Gaines, "System identification, approximation, and complexity," *International Journal of General Systems,* **3**, 145-174 (1977).
76. B.R. Gaines, "An overview of knowledge acquisition and transfer," *International Journal of Man-Machine Studies,* **26**, 453-472 (1987).
77. B.R. Gaines, "How do experts acquire expertise?," *Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Reading, England, 1987.
78. G. Bateson, *Steps to an Ecology of Mind*, Ballantine Books, New York, 1972.
79. M.F. Rubinstein, *Patterns of Problem Solving*, Prentice-Hall, Englewood Cliffs, NJ, 1975.
80. J.A. Wise, "Decisions in design: Analyzing and aiding the art of synthesis," in *Behavioral Decision Making: Theory and Analysis,* G. Wright (Ed.), Plenum Press, New York, 1985, pp. 283-308.
81. W.J. Clancey, "The knowledge level reinterpreted: Modeling socio-technical systems," in special issue on knowledge acquisition of *International Journal of Intelligent Systems,* K.M. Ford and

J.M. Bradshaw (Eds.), 1993. Also to appear in *Knowledge Acquisition as Modeling,* K.M. Ford and J.M. Bradshaw (Eds.), Wiley , New York, 1993.

82. J.M. Bradshaw, J.H. Boose, S.P. Covington, and P.J. Russo, "How to do with grids what people say you can't," *Proceedings of the Third Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Department of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1988, pp. 4.1-4.20.
83. J.C. Jones, *Design Methods: Seeds of Human Futures*, Second Ed. Wiley, New York, 1981.
84. P. McNamee and J. Celona, *Decision Analysis for the Professional—With Supertree*, Scientific Press, Redwood City, CA, 1987.
85. F. Zwicky, *Discovery, Invention, and Research through the Morphological Approach*, Macmillan, New York, 1969.
86. D.B. Shema, J.M. Bradshaw, S.P. Covington, and J.H. Boose, "Design knowledge capture and alternatives generation using possibility tables in Canard," *Knowledge Acquisition,* **2**(4), 345-364 (1990).
87. W. Thompson and B. Thompson, "Creating expert systems from examples," *AI Expert,* **2**(1), 21-26 (1987).
88. C. Garg-Janardan and G. Salvendy, "A conceptual framework for knowledge elicitation," *International Journal of Man-Machine Studies,* **26**(4), 521-532 (1987).
89. M.L.G. Shaw and B.R. Gaines, "KITTEN: Knowledge Initiation and Transfer Tools for Experts and Novices," *International Journal of Man-Machine Studies,* **27**(3), 251-280 (1987).
90. J. Diederich, I. Ruhmann, and M. May, "KRITON: A knowledge acquisition tool for expert systems," *International Journal of Man-Machine Studies,* **26**(1), 29-40 (1987).
91. M.L.G. Shaw and E. Chang, "A participant construct system," *Proceedings of the North American Personal Construct Network Second Biennial Conference*, Calgary, Alberta, Canada, 1986.
92. M.L.G. Shaw, "Conversational heuristics for enhancing personal understanding of the world," in *General Systems Research: A Science, A Methodology, A Technology,* B.R. Gaines (Ed.), Society for General Systems Research, Kentucky, 1979, pp. 270-277.
93. B.R. Gaines and M.L.G. Shaw, "New directions in the analysis and interactive elicitation of personal construct systems," in *Recent Advances in Personal Construct Technology,* M.L.G. Shaw (Ed.), Academic Press, New York, 1981, pp. 147-182.
94. B.R. Gaines and M.L.G. Shaw, "Interactive elicitation of knowledge from experts," *Future Computing Systems,* **1**(2), 151-190 (1986).

95. D. Wahl, "An application of declarative modeling to aircraft fault isolation and diagnosis," in *Intelligent Simulation Environments,* Society for Computer Simulation, La Jolla, CA, 1986, pp. 25-28.

96. J.G. Gammack and R.M. Young, "Psychological techniques for eliciting expert knowledge," *R&D in Expert Systems: Proceedings of the 4th Expert System Conference*, Warwick, England, 1984.

97. J.M. Bradshaw, P. Holm, O. Kipersztok, and T. Nguyen, "*eQuality:* An application of *Axotl II* to process management," in *Current Developments in Knowledge Acquisition: EKAW-92,* T. Wetter et al. (Eds.), Springer-Verlag, Berlin/Heidelberg, 1992, pp. 425-444.

98. J.M. Bradshaw, P.D. Holm, J.H. Boose, D. Skuce, and T.C. Lethbridge, "Sharable ontologies as a basis for communication and collaboration in conceptual modeling," *Proceedings of the Seventh Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Dept. of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1992.

99. K.M. Ford, A.J. Cañas, and J.R. Adams-Webber, "Explanation as a knowledge acquisition issue," *Proceedings of Fourth Annual Florida AI Research Symposium* (FLAIRS-91), Cocoa Beach, 1991, pp. 185-191.

100. B.R. Gaines, "Empirical investigation of knowledge representation servers: Design issues and applications experience with KRS," in *AAAI Spring Symposium: Implemented Knowledge Representation and Reasoning Systems*, Stanford, March, 1991, pp. 87-101.

101. J.H. Boose, J.M. Bradshaw, C.M. Kitto, and D.B. Shema, "From ETS to *Aquinas*: Six years of knowledge acquisition tool development," in *Proceedings of the Fourth Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Dept. of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1989, pp. 5.1-5.17.

102. K.M. Ford, J.R. Adams-Webber, H. Stahl, and M. Bringmann, "Constructivist approaches to knowledge acquisition," in *Knowledge Acquisition: Current Practices and Trends*, K.L. McGraw and C.R. Westphal (Eds.), Ellis Horwood Ltd., New York, 1990, pp. 34-54.

103. B.R. Gaines, "Knowledge acquisition systems for rapid prototyping of expert systems," *INFOR,* **26**(4), 256-285 (1988).

104. J.M. Bradshaw, P. Holm, O. Kipersztok, T. Nguyen, P.J. Russo, and J.H. Boose, "Intelligent interoperability in DDUCKS," *Working Notes of the AAAI-91 Cooperation Among Heterogeneous Intelligent Systems Workshop*, Anaheim, CA, 1991.

105. J.M. Bradshaw, P. Holm, O. Kipersztok, and T. Nguyen, "DDUCKS: A framework for corporate knowledge representation in an enterprise integration architecture," *Proceedings of the Concurrent Engineering and CALS Conference*, Washington, D.C., 1992.
106. J.H. Boose, J.M. Bradshaw, J.L. Koszarek, and D.B. Shema, "Better group decisions: Using knowledge acquisition techniques to build richer decision models," *1992 Hawaii International Conference on Systems Sciences* (HICSS-92), Kauai, Hawaii, 1992, pp. 4-13.
107. J.M. Bradshaw, K.M. Ford, and J.R. Adams-Webber, "Knowledge representation for knowledge acquisition: A three-schemata approach," *Proceedings of the Sixth Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Department of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1991, pp. 4.1-4.20.
108. A. Puerta, S. Tu, and M. Musen, "Modeling tasks with mechanisms," in special issue on knowledge acquisition of *International Journal of Intelligent Systems,* K.M. Ford and J.M. Bradshaw (Eds.), 1993. Also to appear in K.M. Ford and J.M. Bradshaw (Eds.), *Knowledge Acquisition as Modeling,* Wiley, New York, 1993.
109. M.A. Musen, *Automated Generation of Model-Based Knowledge-Acquisition Tools*, Morgan Kaufmann, San Mateo, CA, 1989.
110. K.M. Ford, J.M. Bradshaw, J.R. Adams-Webber, and N.M. Agnew, "Knowledge acquisition as a constructive modeling activity," in special issue on knowledge acquisition of *International Journal of Intelligent Systems,* K.M. Ford and J.M. Bradshaw (Eds.), 1993. Also to appear in K.M. Ford and J.M. Bradshaw (Eds.), *Knowledge Acquisition as Modeling,* Wiley, New York, 1993.
111. T.C. Lethbridge and D. Skuce, "Integrating techniques for conceptual modeling," *Proceedings of the Seventh Banff Annual Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Department of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1992.
112. T.C. Lethbridge, "Creative knowledge acquisition: An analysis," *Proceedings of the Sixth Banff Annual Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Dept. of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1991.
113. D. Skuce, "A frame-like knowledge acquisition tool integrating abstract data types and logic," in *Principles of Semantic Networks,* J. Sowa (Ed.), Morgan Kaufmann, San Mateo, CA, 1991.
114. B. Tauzovich and D. Skuce, *A General Taxonomy for Conceptual Data Modeling*, Cognos Inc., Research Division, 3755 Riverside Dr., PO Box 9707, Ottawa, Ontario, Canada K1G 3Z4, 1990.
115. D. Skuce and I. Monarch, "Ontological issues in knowledge base design: Some problems and suggestions," *Proceedings of the Fifth Banff Knowledge Acquisition for Knowledge-Based*

Systems Workshop, SRDG Publications, University of Calgary, Department of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1990.

116. J.F. Sowa, *Conceptual Structures: Information Processing in Mind and Machine*, Addison-Wesley, Reading, MA, 1984.
117. J.F. Sowa (Ed.), “Towards the expressive power of natural languages,” in *Principles of Semantic Networks*, Morgan Kaufmann, San Mateo, CA, 1991, pp. 157-189.
118. D. Hix, “Generations of user-interface management systems,” *IEEE Software*, September, pp. 77-87 (1990).
119. A. Goldberg, “Information models, views, and controllers,” *Dr. Dobb’s Journal*, July, pp. 1-4 (1990).
120. G.E. Krasner and S.T. Pope, “A cookbook for using the model-view-controller user interface paradigm in Smalltalk-80,” *Journal of Object-Oriented Programming*, August-September, pp. 26-49 (1988).
121. S.S. Adams, “MetaMethods: The MVC paradigm,” *HOOPLA!*, **1**(4), 5-6, 13-21 (1988).
122. S.S. Adams, “MetaMethods: Active Values,” *HOOPLA!*, **1**(1), 3-6 (1988).
123. K.M. Ford, A.J. Cañas, J. Coffey, E.J. Andrews, N. Schad, and H. Stahl, “Interpreting functional images with NUCES: Nuclear Cardiology Expert System,” *Proceedings of Fifth Annual Florida AI Research Symposium* (FLAIRS-92), Ft. Lauderdale, 1992, pp. 85-90.
124. A. Voß, W. Karbach, U. Drouven, D. Lorek, and R. Schuckey, *Operationalization of a synthetic problem*, ESPRIT Basic Research Project P3178 REFLECT Task I.2.1 Report, 1990.
125. M. Linster (Ed.), *Sisyphus Working Papers Part 2: Models of Problem Solving*, EKAW-91, University of Strathclyde, Glasgow, 1991.
126. B. Woodward, “Knowledge engineering at the front-end: Defining the domain,” *Knowledge Acquisition,* **2**(1), 73-94 (1990).
127. J.T. Nosek and I. Roth, “A comparison of formal knowledge representations as communication tools: Predicate logic vs. semantic network,” *International Journal of Man-Machine Studies,* **33**, 227-239 (1990).
128. H. Watanabe, “Heuristic graph displayer for G-BASE,” *International Journal of Man-Machine Studies,* **30**(3), 287-302 (1989).
129. J. Mitterer and J.R. Adams-Webber, “OMNIGRID: A program for the construction, administration, and analysis of repertory grids,” in *Cognition and Personal Structure,* J.C. Mancuso and M.L.G. Shaw (Eds.), Praeger, New York, 1988, pp. 91-106.

130. K.W. Sewell, J.R. Adams-Webber, J. Mitterer, and R.L. Cromwell, “Computerized repertory grids: Review of the literature,” *International Journal of Personal Construct Psychology,* **5**, 1-23 (1992).
131. V.A. Lefebvre, V.D. Lefebvre, and J.R. Adams-Webber, “Modeling an experiment on construing self and others,” *Journal of Mathematical Psychology,* **30**, 317-330 (1986).
132. F. Gomez and C. Segami, “Knowledge acquisition from natural language for expert systems based on classification problem-solving methods,” *Knowledge Acquisition,* **2**(2), 107-128 (1990).
133. B.R. Gaines and M.L.G. Shaw, “Induction of inference rules for expert systems,” *Fuzzy Sets and Systems,* **8**(3), 315-328 (1986).
134. P. Smyth and R.M. Goodman, “An information-theoretic approach to rule induction from databases,” *IEEE Transactions on Knowledge and Data Engineering,* **4**(4), 301-316 (1992).
135. B.R. Gaines, “An ounce of knowledge is worth a ton of data: Quantitative studies of the trade-off between expertise and data based on statistically well-founded empirical induction,” *Proceedings of the Sixth International Workshop on Machine Learning*, Morgan Kaufmann, San Mateo, CA, 1989.
136. J. Cendrowska, “An algorithm for inducing modular rules,” *International Journal of Man-Machine Studies,* **27**(4), 349-370 (1987).
137. D. Madigan, A.E. Raftery, J.C. York, J.M. Bradshaw, and R.G. Almond, “Strategies for graphical model selection,” Submitted to the *Third Conference on AI and Statistics*, Ft. Lauderdale, FL, 1992.
138. D. Madigan, J.C. York, J.M. Bradshaw, and R.G. Almond, “Bayesian graphical models for predicting errors in databases,” Submitted to the *Third Conference on AI and Statistics*, Ft. Lauderdale, FL, 1992.
139. M.L.G. Shaw, “Communities of knowledge,” in *Anticipating Personal Construct Psychology,* F. Epting and A.W. Landfield (Eds.), University of Nebraska Press, Lincoln, 1985, pp. 25-35.
140. M.L.G. Shaw and B.R. Gaines, “Comparing conceptual structures: Consensus, conflict, correspondence and contrast,” *Knowledge Acquisition,* **1**(4), 341-364 (1989).
141. M.L.G. Shaw and B.R. Gaines, “Supporting personal networking through computer networking,” in *Human Factors in Computing Systems • Reaching Through Technology: CHI’91 Conference Proceedings,* S.P. Robertson, G.M. Olson, and J.S. Olson (Eds.), ACM Press, New York, 1991, pp. 437-438.
142. N. Malcolm and B.R. Gaines, *A word processor supporting group-writing*, University of Calgary, Department of Computer Science, Technical Report No. 91/420/04, 1991.

143. K.M. Ford and J.R. Adams-Webber, "Knowledge acquisition and constructivist epistemology," in *The Psychology of Expertise: Empirical Approaches to Knowledge Acquisition*, R.R. Hoffman (Ed.), Springer-Verlag, New York, 1992, pp. 121-136.
144. K.M. Ford and N.M. Agnew, "Expertise: Socially situated and personally constructed," in *Working Notes of the Cognitive Aspects of Knowledge Acquisition Session of the AAAI Spring Symposium*, Stanford, CA, 1992, pp. 80-87.
145. A.J. Cañas and K.M. Ford, "Intelligent support for collaborative modeling," *Proceedings of Sixth Annual Florida AI Research Symposium* (FLAIRS-92), Ft. Lauderdale, 1993, pp. 17-22.
146. D. Schuler, P. Russo, J. Boose, and J. Bradshaw, "Using personal construct techniques for collaborative evaluation," *International Journal of Man-Machine Studies,* **33**, 521-536 (1990).
147. J.H. Boose, "Rapid acquisition and combination of knowledge from multiple experts in the same domain," *Future Computing Systems Journal,* **1**, 191-216 (1986).
148. J.H. Boose and J.M. Bradshaw, "Expertise transfer and complex problems: Using *Aquinas* as a knowledge-acquisition workbench for knowledge-based systems," in *Knowledge Acquisition Tools for Expert Systems,* J.H. Boose and B.R. Gaines (Eds.), Academic Press, London, 1988, pp. 39-64.
149. J.H. Boose and J.M. Bradshaw, "*Aquinas:* A knowledge acquisition workbench," in *AI Tools and Techniques,* M. Yazdani and M. Richer (Eds.), Ablex, Norwood, NJ, 1989, pp. 151-182.
150. B.Q. Post, "Building the business case for group support technology," *Proceedings of the 1992 Hawaii International Conference on Systems Sciences*, Kauai, Hawaii, 1992.
151. J.R. Nunamaker Jr., D. Vogel, A. Heminger, and B. Martz, "Experiences at IBM with group support systems: A field study," *Decision Support Systems,* **5**(2), 183-196 (1989).
152. J. Conklin and M. Begeman, "gIBIS: A hypertext tool for team design deliberation," *Proceedings of Hypertext '87*, Chapel Hill, NC, 1987.
153. J. Conklin and M. Begeman, "gIBIS: A hypertext tool for exploratory policy discussion," *ACM Transactions on Office Information Systems*, **6**, 303-331 (1988).
154. M. Linster, "Closing the gap between *modeling to make sense* and *modeling to implement systems*," in special issue on knowledge acquisition of *International Journal of Intelligent Systems,* K.M. Ford and J.M. Bradshaw (Eds.), 1993. Also to appear in K.M. Ford and J.M. Bradshaw (Eds.), *Knowledge Acquisition as Modeling,* Wiley, New York, 1993.
155. H. Akkermans, F. van Harmelen, G. Schreiber, and B. Wielinga, "A formalisation of knowledge-level models for knowledge acquisition," in special issue on knowledge acquisition of *International Journal of Intelligent Systems,* K.M. Ford and J.M. Bradshaw (Eds.), 1993. Also to

appear in K.M. Ford and J.M. Bradshaw (Eds.), *Knowledge Acquisition as Modeling,* Wiley, New York, 1993.

156. B.R. Gaines, M.L.G. Shaw, and J.B. Woodward, "Modeling as a framework for knowledge acquisition methodologies and tools," in special issue on knowledge acquisition of *International Journal of Intelligent Systems,* K.M. Ford and J.M. Bradshaw (Eds.), 1993. Also to appear in K.M. Ford and J.M. Bradshaw (Eds.), *Knowledge Acquisition as Modeling,* Wiley, New York, 1993.
157. B.R. Gaines, "Modeling practical reasoning," in special issue on knowledge acquisition of *International Journal of Intelligent Systems,* K.M. Ford and J.M. Bradshaw (Eds.), 1993. Also to appear in K.M. Ford and J.M. Bradshaw (Eds.), *Knowledge Acquisition as Modeling,* Wiley, New York, 1993.
158. M.L. Brodie, "Future intelligent information systems: AI and database technologies working together," in *Readings in Artificial Intelligence and Databases,* J. Mylopoulos and M.L. Brodie (Eds.), Morgan Kaufmann, San Mateo, CA, 1989, pp. 623-642.
159. J.F. Sowa and J.A. Zachman, "Extending and formalizing the framework for information systems architecture," *IBM Systems Journal*, **31**(3), 590-616 (1992).
160. S. Perez, *Model Unification for Data Repositories*, American National Standards Institute (ANSI), Technical Report No. X3H4.6/92-001, Version 0.2, 1992.
161. J.A. Fulton, J. Zimmerman, P. Eirich, R. Burkhart, G.F. Lake, M.H. Law, B. Speyer, and J. Tyler, *The Semantic Unification Meta-Model: Technical Approach*, Dictionary/Methodology Committee of the IGES/PDES Organization, ISO TC184/SC4, 1991.
162. M.R. Genesereth and R. Fikes, *Knowledge Interchange Format Version 3.0 Reference Manual*, Stanford University, Department of Computer Science, Logic Group Report, Logic-92-1, 1992.
163. T. Finin, D. McKay, and R. Fritzson, *An Overview of KQML: A Knowledge Query and Manipulation Language*, University of Maryland, Department of Computer Science, 1992.
164. T.R. Gruber, *Ontolingua: A Mechanism to Support Portable Ontologies, Version 3.0*, Stanford University, Department of Computer Science, Stanford Knowledge Systems Laboratory Technical Report KSL 91-66, 1992.
165. T.R. Gruber, "A translation approach to portable ontology specifications," *Proceedings of the Seventh Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Department of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1992.
166. M. Weiser, "The computer for the 21st century," *Scientific American* (September), pp. 94-104 (1991).

167. L.G. Tesler, "Networked computers in the 1990s," *Scientific American* (September), pp. 86-93 (1991).
168. B. Ryan, "Dynabook revisited with Alan Kay," *BYTE* (February), pp. 203-208 (1991).
169. T.R. Gruber, "Model formulation as a problem solving task: Computer-assisted engineering modeling," in special issue on knowledge acquisition of *International Journal of Intelligent Systems,* K.M. Ford and J.M. Bradshaw (Eds.), 1993. Also to appear in K.M. Ford and J.M. Bradshaw (Eds.), *Knowledge Acquisition as Modeling,* Wiley, New York, 1993.
170. A. Rappaport, "A theory of local and interstitial knowledge," *Proceedings of the Sixth Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Dept. of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1991.
171. Y. Arens, S. Feiner, J. Foley, E. Hovy, B. John, R. Neches, R. Pausch, H. Schorr, and W. Swartout, *Intelligent User Interfaces*, University of Southern California/Information Sciences Institute, Report No. ISI/RR-91-288, 1991.
172. X/Open and O.M. Group, *The Common Object Request Broker: Architecture and Specification*, Reading, England: X/Open; Framingham, MA: OMG, OMB Document No. 91.12.1, Revision 1.1, 1991.
173. D. Marques, G. Klinker, G. Dallemagne, P. Gautier, J. McDermott, and D. Tung, "More data on usable and reusable programming constructs," *Proceedings of the Sixth Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, SRDG Publications, University of Calgary, Department of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1991.
174. D. Marques, G. Dallemagne, G. Klinker, J. McDermott, and D. Tung, "Easy programming: Empowering people to build their own applications," *IEEE Expert*, **7**(3), 16-29 (June, 1992).
175. R. Neches, R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, and W.R. Swartout, "Enabling technology for knowledge sharing," *AI Magazine*, **12**(3), 36-55 (Fall, 1991).
176. R. Neches, J. Foley, P. Szekely, P. Sukaviriya, P. Luo, S. Kovacevic, and S. Hudson, *Knowledgeable Development Environments Using Shared Design Models*, USC/Information Sciences Institute and Georgia Institute of Technology, 1992.
177. D.A. Norman, "Cognitive artifacts," in *Designing Interaction: Psychology at the Human-Computer Interface,* J.M. Carroll (Ed.), Cambridge University Press, Cambridge, 1992, pp. 17-38.
178. B. Chandrasekaran, N.H. Narayanan, and Y. Iwasaki, "Reasoning with diagrammatic representations: A report on the AAAI Spring Symposium, March 25-27, 1992," *AI Magazine* (in press).

179. Y. Shoham, *Agent-Oriented Programming*, Stanford University, Department of Computer Science Technical Report STAN-CS-1335-90 (Revised), 1990.
180. A. Kay, "Computer software," *Scientific American,* **251**(3), 53-59 (1984).
181. B. Laurel (Ed.), "Interface agents: Metaphors with character," in *The Art of Human Computer Interface Design,* Addison-Wesley, Reading, MA, 1990, pp. 355-365.
182. T. Erickson, *Design Issues for Enhanced Interfaces: Models and Metaphors*, Apple Computer Human Interface Group, 1991.
183. G.A. Boy and N. Mathé, "Operator assistant systems: An experimental approach using a telerobotics application," in special issue on knowledge acquisition of *International Journal of Intelligent Systems,* K.M. Ford and J.M. Bradshaw (Eds.), 1993. Also to appear in K.M. Ford and J.M. Bradshaw (Eds.), *Knowledge Acquisition as Modeling,* Wiley, New York, 1993.
184. G.A. Boy, "Acquiring and refining procedures according to context," *Proceedings of the Fifth Banff Knowledge Acquisition for Knowledge-based Systems Workshop*, SRDG Publications, University of Calgary, Department of Computer Science, Calgary, Alberta, Canada T2N 1N4, 1990, pp. 2.1-2.14.
185. G.A. Boy, "Indexing hypertext documents in context," *Proceedings of the Third ACM Conference on Hypertext*, San Antonio, TX, 1991, pp. 1-11.
186. J. Glicksman, J.C. Weber, and T.R. Gruber, "The NoteMail project for computer supported cooperative mechanical design," *Working Notes of the AAAI-92 Workshop on Design Rationale Capture and Use*, San Jose, CA, 1992.
187. T.R. Gruber, J.M. Tenenbaum, and J.C. Weber, "Toward a knowledge medium for collaborative product development," *Workshop Notes of the AAAI-92 AI in Enterprise Integration Workshop,* San Jose, CA, 1992.
188. J.C. Spohrer, D. Vronay, and R. Kleiman, "Authoring intelligent multimedia applications: Finding familiar representations for expressing knowledge," in *Proceedings of the 1991 IEEE International Conference on Systems, Man, and Cybernetics*, Charlottesville, VA, 1991, pp. 1751-1756.
189. J.M. Bradshaw and G.A. Boy, "Adaptive documents," submitted to JAVA-93 (Journees sur l'Acquisition, la Validation et l'Apprentissage), 29 Mars - 2 Avril 1993, Saint-Raphael, France.