Knowledge Acquisition (1991) **3**, 215–236

ICONKAT: an integrated constructivist knowledge acquisition tool†

KENNETH FORD, ALBERTO CAÑAS AND JEREMY JONES
Institute for Human and Machine Cognition, University of West Florida, Pensacola, FL 32514, USA

HOWARD STAHL AND JOSEPH NOVAK
Department of Education, Roberts Hall, Cornell University, Ithaca, NY 14853, USA

JACK ADAMS-WEBBER
Department of Psychology, Brock University, St. Catharines, Ontario, Canada L2S 3A1

(Received 31 January 1991 and accepted in revised form 3 March 1991)

In this paper, we report on a continuing research effort aimed at the development of an integrated knowledge acquisition system, ICONKAT. We describe the components of the tool and discuss how they may be used to facilitate the design, construction, testing, maintenance and explanation of knowledge bases. ICONKAT's knowledge elicitation subsystem, based on both personal construct theory and assimilation theory, interactively assists the domain expert in the task of building a model of his or her expertise. ICONKAT employs a collection of modeling primitives (i.e. the glue) as the material basis for the construction of a conceptual domain model. The maintenance subsystem provides support tools for use by the knowledge engineering team, as well as the domain expert, when testing the system's performance, refining the knowledge base, and maintaining the overall system. The components of the maintenance subsystem employ a variety of mediating representations (e.g. concept maps, repertory grids) to furnish various perspectives of the evolving domain model as embodied in the modeling primitives. Moreover, 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 is currently employed in the design and construction of an expert system for the diagnosis of first pass functional cardiac images.

1. Introduction

The first and most fundamental step in the knowledge acquisition process is the elicitation of human expertise and its subsequent constructive modeling in a computer system. The mining analogy notwithstanding, we assume that expertise is not like a natural resource that may be harvested, transferred or captured, but rather is constructed by the expert and reconstructed by the knowledge engineer. A pivotal aspect of the knowledge acquisition process is the knowledge engineer's act of construction in which the utterances/tool-interactions of the domain expert are transformed into an implementation formalism. This *modeling* view of knowledge

† Based on a paper presented at the 5th Knowledge Acquisition Workshop Banff, October 1990.

1042-8183/91/020215 + 22$03.00/0

acquisition implies a need for tools that directly support both the domain expert and knowledge engineer.

The knowledge acquisition process does not culminate at some arbitrary point in the development life-cycle, but rather extends for the life of the system. It encompasses the following facets (among others): elicitation and modeling of human expertise; testing the system's efficacy; refinement of the knowledge base; maintenance of the resulting system; and elaboration of an explanation capability. This *process* view implies that, in principle, knowledge acquisition tools should support all of these activities.

1.1. AUTOMATED KNOWLEDGE ACQUISITION TOOLS

Theories without tools are empty, tools without *theories* are blind.

(with apologies to Immanuel Kant).

A commonly proposed partial solution to the difficulties associated with the knowledge acquisition process is the design and implementation of automated knowledge acquisition tools. To what extent are we justified in optimistic expectations about the usefulness of such tools? Upon what foundation can we base these sanguine anticipations? It seems clear that we should expect a proposed knowledge acquisition tool (as with any other tool) to be successful to the extent that we have a serviceable theory explaining the envisioned tool's basis of operation and delineating its approximate range of application. For example, we are able to satisfactorily specify the circumstances under which a thermometer is a useful tool only because we have a valid theory about its use.

However, as noted by Bradshaw and Boose (1990), a significant difficulty has been that much of the current work in knowledge acquisition lacks a plausible theoretical foundation.

> As a consequence of incomplete theory and a limited repertoire of practical approaches to the dynamics of the modelling process, knowledge engineers have had to rely on intuition and experience as the primary means of developing and testing effective procedures (p. 6).

Rarely do working knowledge engineers explicitly consider the epistemological underpinnings of the methods and tools that they employ in their task.

Many workers engaged in knowledge acquisition (as researchers or practitioners) may be classified as tool-makers and/or tool-users. Tool-makers should exploit theory as a means of developing their tools on a sound footing, and also as a framework in which to make explicit their epistemological assumptions (Ford & Adams-Webber, 1991). Furthermore, theory may offer tool-makers a useful infrastructure upon which to build highly integrated collections of tools and techniques (i.e. hybrid knowledge acquisition workbenches). Tool-users may employ theory as the basis for a principled application of the myriad tools now becoming available. That is, a robust theory can endow a knowledge acquisition tool with the kind of conceptual supports (guides) that an operators manual alone cannot provide.

Integration is the battle cry as tool-makers rush to produce hybrid knowledge acquisition tools. Nonetheless, it is becoming widely realized† that *ad hoc*

† General concensus expressed in the group meeting of the knowledge acquisition tools special interest group at the Fourth Knowledge Acquisition for Knowledge-Based Sysems Workshop, Banff (cf. Bradshaw & Woodward, 1989, p. 12).

combinations of techniques and tools—a sort of Occam's hell—may not contribute much to ameliorating the knowledge acquisition bottleneck. As noted earlier, theory can offer an infrastructure upon which to construct highly integrated hybrid knowledge acquisition tools in a principled way. ICONKAT is one effort in this direction.

The main sections of this paper are concerned with the following related topics: Section 2, Kelly's personal construct theory and its major methodological tool—repertory grids; Section 3, Ausubel's assimilation theory and its major methodological tool—concept maps; Section 4, ICONKAT, a constructivist hybrid knowledge acquisition tool based upon an integration of repertory grids and concept maps. The theoretical rational of the system is explicated from the perspective of our evolving notions about the knowledge acquisition process, as well as in terms of principles of both personal construct theory and assimilation theory. Section 5 provides a general summary.

2. Personal construct theory

Personal construct theory, as formulated by Kelly (1955, 1969*a*, 1970) and summarized by Adams-Webber (1987, 1989), is essentially a constructivist model of human representational processes. The basic units of analysis in this theory are bipolar dimensions termed "constructs". Kelly posited that the underlying distinction which lends each construct meaning is dichotomous in form, for example, happy/sad, young/old, odd/even. He argued further that the dichotomous structure of personal constructs is an essential feature of the way in which people organize information (cf. Adams-Webber, 1990). For example, if the same event could be perceived simultaneously as equally pleasant and unpleasant in the same respect, then this distinction would have no definite meaning.

We use our personal constructs to represent perceived similarities and differences among events, and then we organize these representations into coherent patterns or "schemata" within the framework of which we are able to detect certain recurrent patterns in our experience over time. We then feed these representations forward in the form of expectations about future events (cf. Ford, 1989). Each construct has a specific range of convenience, which comprises "all those things to which the user would find its application useful" (Kelly, 1955). Accordingly, the range of convenience of a construct defines its extension in terms of a single aspect of a limited domain of events (Mancuso & Eimer, 1982). A particular construct seldom, if ever, stands alone in our experience, because it is usually deployed together with one or more other hierarchically related constructs in interpreting and predicting events. Indeed, a necessary condition for organized thought is some degree of overlap between constructs in terms of their respective ranges of convenience (Adams-Webber, 1970). It is this overlap, or intersection, between the constructs' extensions that enables us to form hypotheses. That is, in interpreting an event we essentially categorize it in terms of one or more constructs, and then by considering related constructs, we can derive predictive inferences from that initial categorization.

2.1. REPERTORY GRIDS

Kelly's (1955) *Role Construct Repertory Grid* ("repertory grid") was developed to assess an individual's personal construct system, that is, "one's own finite system of cross references between the personal observations he has made and the personal constructs he has erected" (Kelly, 1969*b*, p. 291). A repertory grid is essentially a complex sorting task in which a list of elements are judged successively on the basis of a set of bipolar constructs (Adams-Webber, 1987). These elements can be either concrete or abstract entities, and may be said to operationally define the grid's universe of discourse. They should be chosen carefully to represent completely the topic of interest, and also be roughly of the same type and level of complexity.

Constructs represent the ways in which elements are judged to be similar or different from each other: that is, they permit the subject to make relevant distinctions among the elements. Thus, a repertory grid may be considered a mapping of elements onto constructs. The data generated by each respondent are entered into a separate two-dimensional table, or "grid", in which there is a column for every element and a row for every construct. Each row-column intersect in this table contains a number indicating how a given construct was applied to a partiuclar element.

The repertory grid in Figure 1 was produced, using ICONKAT, by a domain expert (radiologist) participating in the development of a large-scale expert system for the diagnosis of heart wall motion abnormalities. During the last 30 years, there has been a considerable proliferation of new forms of repertory grid in which people, objects, situations or other kinds of elements, are either categorized, rated or rank-ordered on constructs. The act of a respondent's assigning a rating to an element on a given construct has been interpreted in a variety of ways, which has led to several different approaches to grid analysis, including information theoretic measures, non-parametric factor analysis, conventional factor analysis, principal component analysis, multidimensional scaling, and hierarchical cluster analysis, among others (Adams-Webber, 1987). These procedures are widely used in clinical psychology, as well as in the study of education and management (cf. Adams-Webber & Mancuso, 1983). In addition, several recent approaches to knowledge acquisition have been based on personal construct psychology. For example, repertory grid analysis is becoming an increasingly useful tool in knowledge acquisition for expert systems (Gaines & Shaw, 1980, 1986; Boose, 1986; Ford, Petry, Adams-Webber & Chang, 1991).

Element Index:

1 - Normal	5 - Abnormal Fx/Ischemia
2 - Nonspecific Wall Abnormalities	6 - Abnormal Fx/Cardiomyopathy
3 - Nonspecific WMA/MVP	7 - Abnormal Fx/Valve Disease
4 - Suspicion for Ischemia	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/cxer	worsen/exer	1	2	2	3	2	2	2	3
C6	symmetric	assymetric	1	1	2	3	3	1	2	3

FIGURE 1. A repertory grid from the domain of nuclear cardiology.

Nonetheless, the repertory grid is not a "magic matrix". The root of the method's relative success lies in the fact that it is the principal methodological component of personal construct theory—an epistemologically sophisticated psychology (Ford & Adams-Webber, 1991). Within the framework of personal construct theory, the use of repertory grids may assist the knowledge engineer in avoiding the domain expert's cognitive defenses, and thus elicit deep knowledge that the expert would not have been able to express otherwise (see Section 4).

3. Assimilation theory

Ausubel's assimilation theory (Ausubel, 1963; Ausubel, Novak & Hanesian, 1978), as does Kelly's personal construct theory (Kelly, 1955), belongs to the family of theories contributing to a constructivist model of human representational processes. Assimilation theory is essentially a cognitive learning theory that has been applied to education. Ausubel argues that learning is synonymous with a change in the meaning of experience. His fundamental premise seems deceptively simple:

> 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. (Ausubel *et al.*, 1978, p. 159)

Ausubel posits that meaningful learning involves the assimilation of new concepts and propositions into existing cognitive structures. In this context, concepts are defined as perceived regularities in events or objects (or the records of events or objects) which have been designated by a label. This assimilation of new meaning leads to progressive differentiation and reintegration of cognitive structures.

Ausubel's theory deals with the processes involved in linking new information to existing cognitive structure in a propositional manner. He explicates various forms of meaningful, as opposed to rote learning that involve the assimilation of new information. Ausubel assumes that meaningful learning requires that the learner's cognitive framework contain relevant anchoring ideas to which new material can be related. Indeed, he argues that "the most important single factor influencing learning is what the learner already knows. Ascertain this and teach him accordingly".

Specifically, "concept maps", developed by Novak (1977), have been designed to tap into a person's cognitive structure and externalize concepts and propositions. In the next section, we will describe concept mapping, which is assimilation theory's major methodological tool for ascertaining what is already known.

3.1. CONCEPT MAPS

A concept map is a two-dimensional representation of a set of concepts that is constructed so that the interrelationships among them are evident. The vertical axis expresses a hierarchical framework for the concepts. More general, inclusive concepts are found at the highest levels, with progressively more specific, less inclusive concepts arranged below them. These maps emphasize the most general concepts by linking them to supporting ideas with propositions.

Along the vertical axis, concept maps display Ausubel's notion of subsumption, namely that new information is often relatable to and subsumable under more inclusive concepts. Along the horizontal axis, maps express the idea of progressive differentiation by showing how concepts acquire richer meaning as new propositions are acquired. This axis shows that meaning making is a continuous process. Concepts are always being modified and made more explicit and inclusive.

The overall structure of a concept map constitutes a hierarchical framework for the concepts included in it. All concepts at any given level in the hierarchy will tend to have a similar degree of generality. Consider the concept map about concept mapping in Figure 2 (by convention links run top-down unless annotated with an arrowhead).

Concept maps represent meaningful relationships between concepts in the form of propositions. Propositions are two or more concepts linked by words to form a semantic unit. In the simplest form, a concept map would contain just two concepts connected by a linking word to form a single proposition. For example, "grass is green" would represent a simple map forming a valid proposition about the concepts "grass" and "green". A concept acquires additional meaning as more propositions include the concept. Thus, "grass is green", "grass is a plant", "grass grows" and so on, all expand the meaning of the concept "grass". In this sense, concept maps represent meaning in a framework of embedded propositions.

In educational settings, concept mapping techniques have aided people of every age to examine many fields of knowledge (Novak & Gowin, 1984). Their rich expressive power derives from each map's ability to allow its creator the use of a virtually 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 observing nuances of meaning. Mapping techniques have also been employed to help students "learning how to learn" by bringing to the surface cognitive structures and self-constructed knowledge. Problems of this sort currently

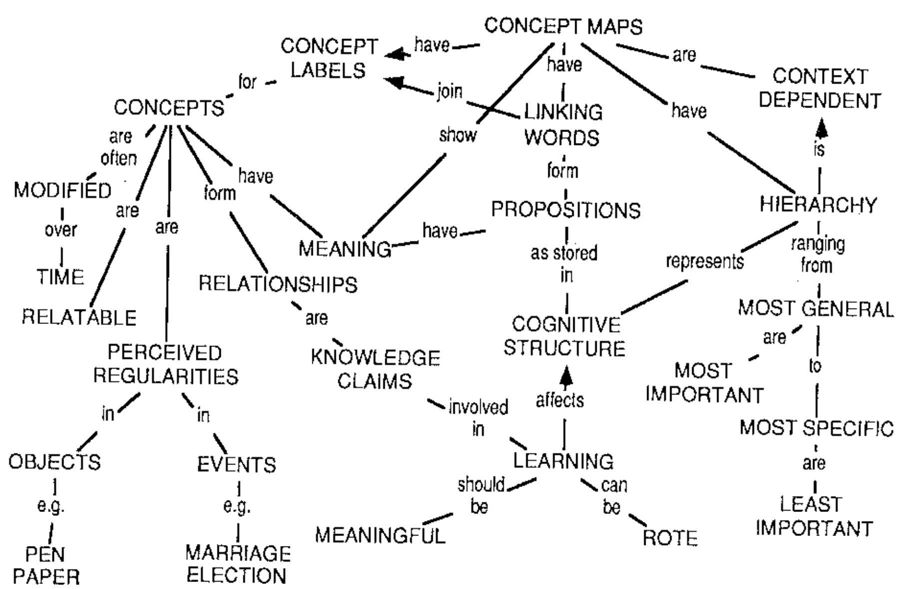

FIGURE 2. A concept map about concept mapping.

face the knowledge acquisition community. As a knowledge elicitation technique, concept maps provide a framework for eliciting the most significant part of a specialist's expertise, that is, his or her personally constructed knowledge (Ford & Adams-Webber, 1991).

The structure of a concept map is dependent on its context. Consequently, maps having similar concepts can vary from one context to another and are highly idiosyncratic. The strength of concept maps lies in their ability to measure a particular person's knowledge about a given topic in a specific context (Novak & Gowin, 1984). There is evidence that when experts and novices in the same domain construct concept maps, the former produce maps which are both more differentiated and integrated (Novak, 1989).

4. Knowledge acquisition in the ICONKAT environment

The constructivist tool described here (i.e. ICONKAT) interactively assists the domain expert in the task of building a model of his or her expertise. In keeping with the view that the knowledge acquisition process continues for the life of the system, ICONKAT provides explicit support for the elicitation and modeling of human expertise, documentation, testing, maintenance, and explanation of the knowledge base under development. Figure 3 depicts a concept map overview of the entire system.The three major components of ICONKAT—the elicitation, maintenance, and explanation subsystems are described below in Sections 4.2, 4.3 and 4.4 respectively. In addition, Section 4.1 offers a discussion of ICONKAT's modeling environment—the framework that integrates these constituent subsystems.

4.1. THE MODELING ENVIRONMENT

ICONKAT's modeling environment consists of both mediating and intermediate representations. Although the terms *mediating* and *intermediate* representations have various interpretations in the literature (Ford, Bradshaw & Adams-Webber, in press), we take the term mediating representation to "convey the sense of synthesis and coming to understand through the representation", while the expression intermediate representation is defined as one "which only exists between flanking

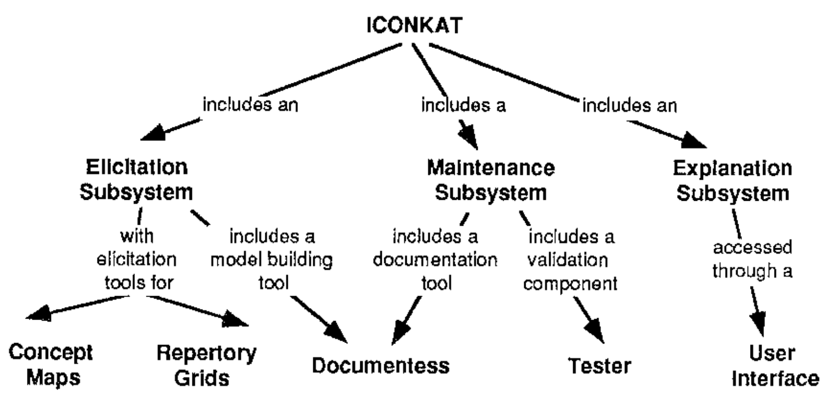

FIGURE 3. ICONKAT system architecture.

representations and is bound to them by clearly defined projection rules which map one representation to the next" (Johnson, 1989, p. 184).

ICONKAT's modeling environment (including both mediating and intermediate representations) contains all the various informational components of the knowledge acquisition process, thus providing direct support for model creation, documentation, maintenance, knowledge base generation, and the subsequent expert system's explanation facility.

4.1.1. Mediating representations

The essential function of a mediating representation is that of promoting communication and understanding. This may entail enhancing communication amongst the participants in the knowledge acquisition process, and/or providing for improved understanding of the evolving conceptual domain model. ICONKAT's principal mediating representations are the concept map and repertory grid.

Mediating representations help bridge the gap between the knowledge engineer's view of the domain and the domain expert's perspective. An adequate mediating representation fosters the constructive modeling processes (e.g. meaning making and meaning sharing) by assisting the knowledge engineer in bringing the experts' self-constructed knowledge to the surface, making explicit the valuable heuristic knowledge that experts possess but are frequently unable to articulate. Thus, mediating representations empower domain experts and knowledge engineers to build models of expert knowledge cooperatively.

In addition to their role in ameliorating the communication difficulties among the human participants, mediating representations help bridge the gap between the domain expert (and/or knowledge engineer) and the implementation formalism required by the performance environment. Typically, knowledge bases are not organized from the perspective of humans, but rather for the convenience of the representation and reasoning mechanisms of the performance environment. Mediating representations can provide a mapping between human and machine-oriented representations.

Furthermore, mediating representations may facilitate explanation (see section 4.4) 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).

4.1.2. Intermediate representations

Intermediate representations are intended to provide the integrating glue between flanking representations, mediating and otherwise. ICONKAT's intermediate representations (sometimes referred to as the "glue") consist of a collection of modeling primitives implemented as abstract data types in C++. These modeling primitives are largely the product of the expert's interaction with the mediating representations (e.g. concept maps and repertory grids) found in ICONKAT's elicitation subsystem (see Section 4.2). In addition, the knowledge engineer can add supporting objects and primitives (e.g. images, audio, video, documents) by means of the Documentess (see Section 4.3.1) interface. Thus, the conceptual domain model is constructed within the framework afforded by the modeling primitives in conjunction with the mediating representations.

4.2. THE ELICITATION SUBSYSTEM

ICONKAT interacts with the domain expert largely through the mediating representations furnished by its concept mapping and repertory grid interfaces (see Figure 4). The system uses concept maps to elicit the "object space" from the domain expert, incrementally generating a hierarchically-structured collection of objects embodied in ICONKAT's intermediate representation (see Section 4.1.2). Concepts in the map are specified by the user as either classes, objects, sub-objects or properties. The links are used to propagate inheritance appropriately.

ICONKAT elicits much of the reasoning dimension (e.g. rules) by means of repertory grids. The hierarchically-grouped concepts found in the elicited concept map may provide the elements for a corresponding hierarchical collection of

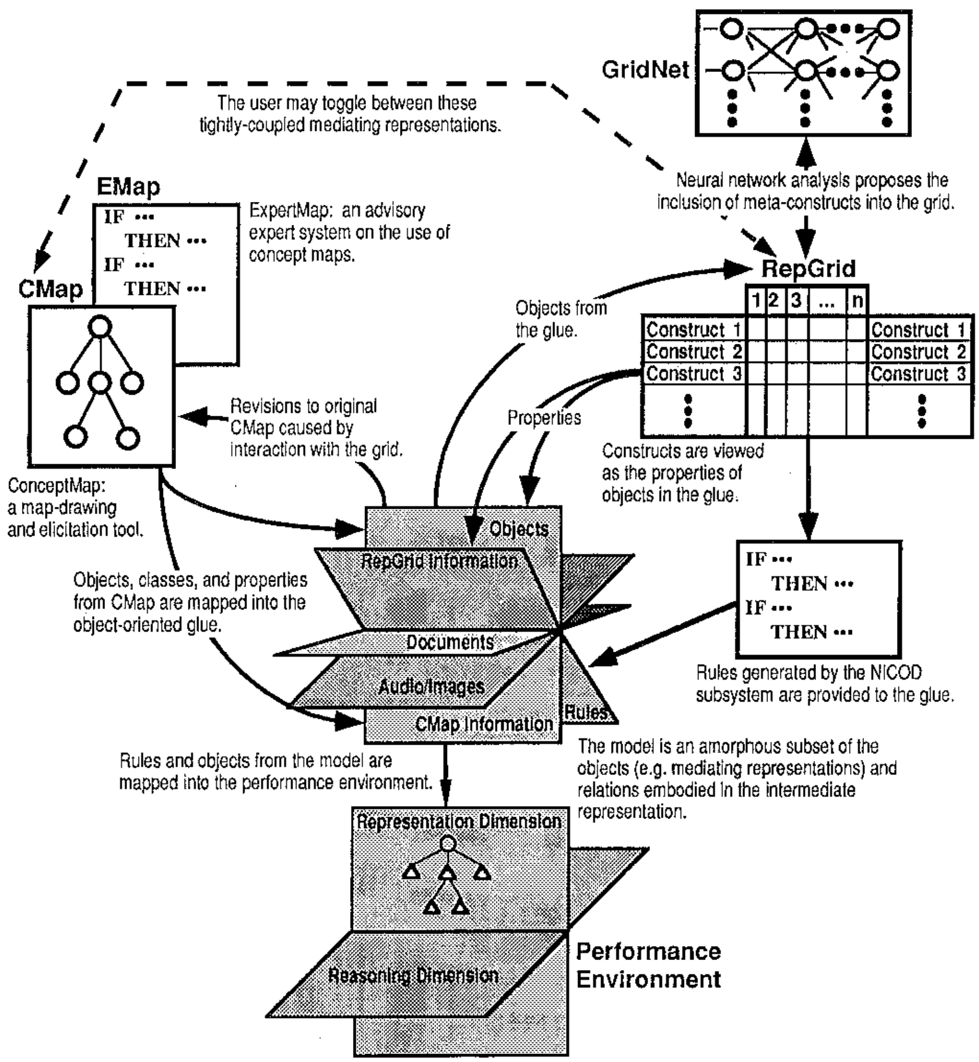

FIGURE 4. Elicitation subsystem overview.

repertory grids, with each grid at a different level of abstraction. Using GridNet (the neural net construct elicitation method described in Section 4.2.2) as well as triads and other well-known repertory grid techniques (cf. Adams-Webber, 1979), constructs are elicited on the basis of the elements already defined as objects in the intermediate representation. The constructs can then serve as the properties of these objects. Further analysis of the repertory grids will derive the construct entailments necessary to generate production rules (Ford, Petry, Adams-Webber & Chang, 1991). Thus the system creates a reasoning space structured on the basis of the object space already represented.

ICONKAT employs the complementary mediating representations of concept maps and repertory grids in a synergetic fashion. Concept maps afford the knowledge engineer a conceptual overview of the expert's domain model, thus providing an essential initial structuring of the domain. This initial structuring informs the knowledge engineer about the appropriate use of repertory grids and other knowledge elicitation techniques within this domain.

4.2.1. The concept mapping components

The concept mapping elements of ICONKAT are primarily based on CMap (Hunter, Stahl & Novak, 1990), an elicitation tool for assisting in the design and construction of concept maps. CMap provides a graphical concept mapping environment which users can employ to codify their knowledge of a specific domain. Initially, the user supplies concepts to the drawing program one at a time; and later, as the user develops the map, propositions can be included by linking concepts together. Propositions can, if desired, be left unlabelled to signify a relationship between two concepts which, for whatever reason, cannot be explicitly named at the time the link is established.

To illustrate the powerful visual component of CMap, consider the screen-dump of the system in operation shown in Figure 5. The sample screen displays a concept map (reproduced from Novak & Gowin, 1984) as it appears in the CMap

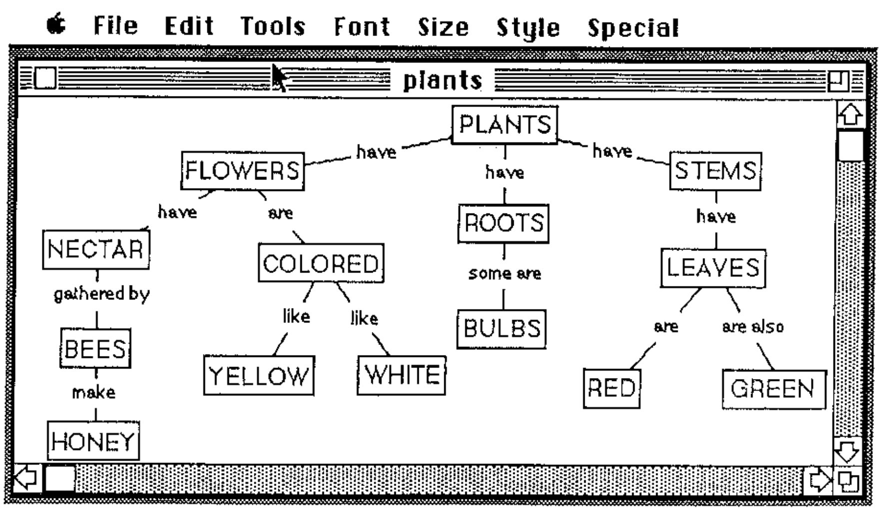

FIGURE 5. A sample screen from the CMap subsystem.

environment. During an interactive session with CMap, users are free to edit, remove or relocate the concepts and links in their map. Throughout this entire development process, maps are continuously displayed in their two-dimensional graphical representation. This constant visual feedback provides essential information to users when evaluating the effectiveness of a specific modification or alteration, and empowers them to consider alternative conceptual hierarchies, thereby constructing an effective map, one which perspicuously conveys the desired meaning and relationships they are attempting to express.

The concept mapping portion of ICONKAT also includes a real-time advisory expert system (known as EMap) which assists the domain expert in creating well-structured maps. EMap responds to user-initiated requests for advice about the map under construction with information about actions the user might reasonably take to improve his or her map. By monitoring the structure of the concept map under development, EMap can focus attention on regions within a concept map in need of restructuring or further elaboration. EMap aids the process of knowledge base construction by exploiting both the tool's underlying theory (assimilation theory) and the heuristic knowledge of its inventor (Joe Novak served as the domain expert for EMap).

In application to the nuclear cardiology domain, we have found this concept space elicitation method quite successful. The concept map in Figure 6 was produced, using ICONKAT, by a domain expert after assimilating some of the map-structuring advice suggested by EMap. The concept map provides a high-level conceptual overview of the domain as constructed by the expert. In particular, this map 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. Moreover, the relevant disease states about which this expert renders a diagnosis can be found at the lowest levels of the map, and have been incorporated as elements in the repertory grid components of ICONKAT. 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", "Blue Bull's-eyes" and "Ice Cream Cones"). These markers are the basis on which the expert differentiates the various disease states, and as such, they may be included as constructs in the developing repertory grid.

In summary, the concept mapping components of ICONKAT assist the user (i.e. domain expert) to concentrate on introspection, and facilitate the elicitation of personally constructed knowledge, which in our view, often represents the most significant and difficult to formulate portion of expertise (Ford & Adams-Webber, 1991).

4.2.2. The repertory grid components

The repertory grid aspects of ICONKAT have evolved from NICOD (Ford *et al.*, 1991) and OMNIGRID (Mitterer & Adams-Webber, 1988). In addition, ICONKAT includes a new neural net method (described below) for construct elicitation (Jones & Ford, 1991). The repertory grid component of ICONKAT are illustrated in Figure 7.

NICOD is a semi-automated knowledge acquisition system supporting the

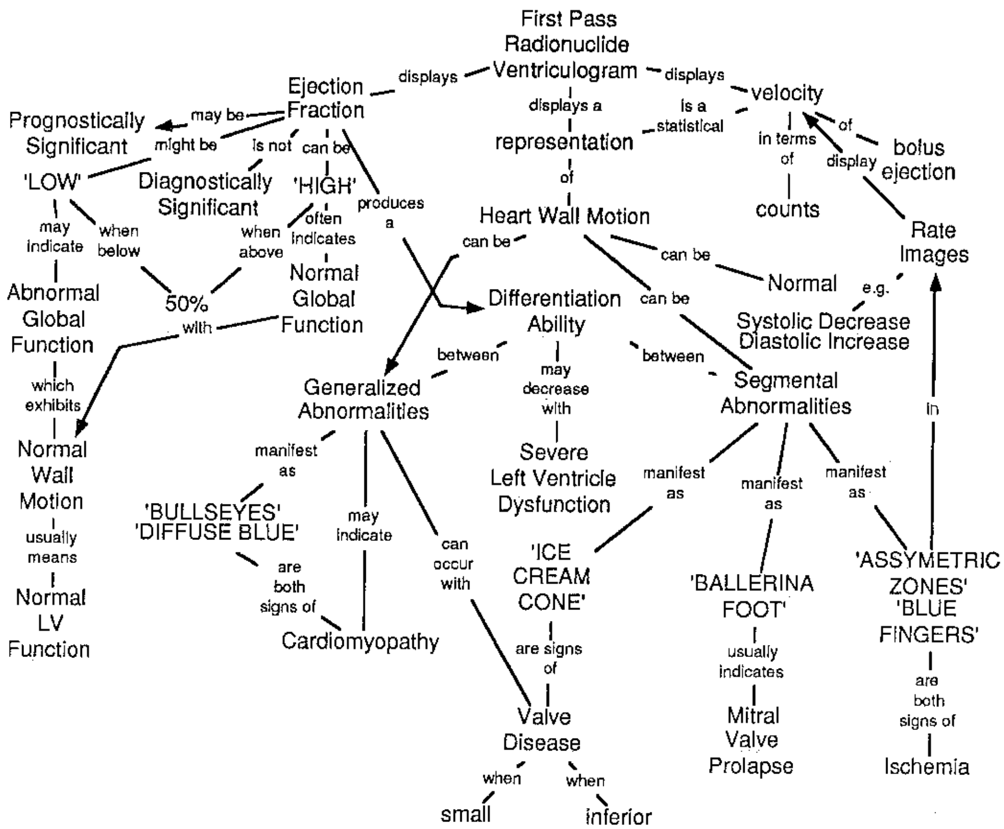

FIGURE 6. A segment of a concept map from the domain of nuclear cardiology.

elicitation, representation and analysis of repertory grid data. It has been used by radiologists to produce quickly, useful rule-sets in several medical domains, including mammography and nuclear cardiology. NICOD is grounded in a theory of confirmation that incorporates tenets of personal construct theory directly into the logic as a basis for the determination of relevance, thus strengthening the logic, and extending personal construct theory. More specifically, Kelly's notion of "range of convenience" is used to define operationally the range of relevance of hypotheses which are derived from relations of entailment between constructs. The concept of α-planes (Ford, 1987; Ford *et al.*, 1991) is used for the binary decomposition of repertory grid data elicited from a domain expert. This procedure furnishes the uniquely determined strings of incidences required for the application of Bundy's (1985) truth functional incidence calculus.

OMNIGRID is a general tool for designing, administering and analysing repertory grids that affords a variety of different formats, including rating scales, nominal categorization and rank-order procedures, with either elicited or supplied constructs and elements. In the analysis of repertory grid data elicited by this program, we are able to select from a wide range of composite indices and summary statistics tailored to the specific format options, including construct intercorrelations, element distances, information statistics and rating extremities.

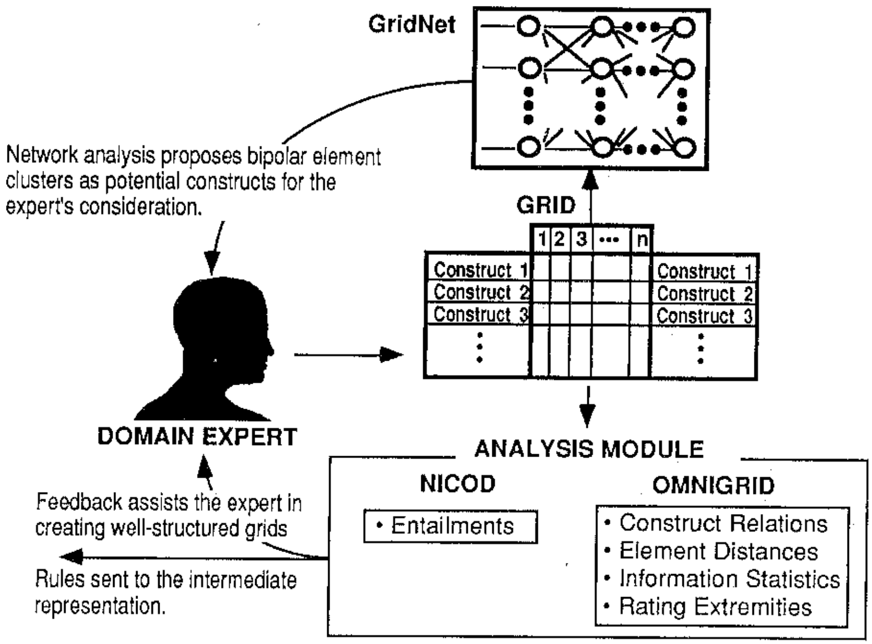

FIGURE 7. The repertory grid components of ICONKAT.

4.2.2.1. GridNet: a neural net approach to construct elicitation. Although the repertory grid method of knowledge acquisition has proved quite successful, there remains room for methodological improvement. One difficulty for repertory grid oriented knowledge acquisition tools is the domain expert's inability to provide a reasonably complete sample of their constructs. This problem may stem partly from the limitations of currently available methods of construct elicitation. It is somewhat ironic that a knowledge acquisition tool should suffer from this problem.

Typically, constructs are elicited by asking the expert to compare the grid's elements (which are usually listed first) with each other in groups of two or three. When groups of two are used, the expert is asked to name a construct that distinguishes between them. When Kelly's (1955) triad method is used, the expert is asked to name a characteristic which two elements have in common that distinguishes them from the third. As noted by Butler and Corter (1986), there are logical problems for construct elicitation methods that rely exclusively on questions about grouped elements. Contrasting two elements with a third can fail to discover an important construct (characteristic) that distinguishes the pair. Conversely, if two objects are contrasted with one another, an important common characteristic may be overlooked. These objections can be met if the knowledge engineer asks the expert to (1) give the common characteristics of a pair of elements A and B, and (2) give the features that A has that B does not (distinctive characteristics of A) and that B has that A does not (distinctive characteristics of B) (cf. Tversky, 1977). However, this technique suffers from the problem that the number of pairs and triples increases exponentially with the number of elements.

GridNet, the neural net element of ICONKAT, assists the expert to uncover non-linear patterns, as they occur during the construction of a repertory grid, suggesting higher level "meta-constructs" (see Figure 7). Kelly (1955, 1970) implies that relatively superordinate constructs might function as pivotal axes of reference in terms of which alternative construct subsystems can be logically integrated at higher levels of abstraction.

The expert's preliminary repertory grid data will serve as input to a self-organizing, multilevel, artificial neural net which uses back propagation, in which abstractions taken from this non-linear hierarchy, called "hidden features", can be identified. These artificial neural net abstractions are then fed back as suggestive prompts to the domain expert for the elicitation of new constructs. Artificial neural net learning algorithms seem especially well suited to the extraction of meaningful patterns from repertory grid data that can be used as useful prompts for the elicitation of new constructs.

Latent features (i.e. bipolar element groupings) are derived from repertory grid data in the following manner. The construct vector patterns (i.e. rows, see Figure 1) from a given repertory grid are repeatedly and randomly input to a basic feedforward network during the back propagation net training session. In each case, the data are compressed through only four initial hidden units, and then are expanded to produce an output as close to the original input vector as possible.

To the extent that this identity mapping suceeds for all construct vectors (six in the case of Figure 1), the first layer of hidden units can be said to have captured or encoded the structure of the entire repertory grid in the following respect: the output from each hidden unit is the threshholded inner product of the input pattern vector with a weight (feature) vector describing the hidden unit. This scalar product determines the component or amount of the feature vector contained in an input pattern in the ordinary Euclidean Hilbert Space sense.

A feature embodies one of the most recurrent patterns in the expert's point of view as represented by the multilayer, non-linear structure, extracted from the repertory grid by the network learning algorithm. Note that since multilayer connectionist systems are generally highly non-linear and hierarchical, these recurrent features may be latent within a number of hierarchical views, where each feature is the input cluster to a deeper hierarchy of clusters. It is this capacity for uncovering many-layered, non-linear hierarchies that make possible the qualitative feature sets that can assist the domain expert in identifying new constructs for subsequent use in the expert's expanded repertory grid.

GridNet and the knowledge engineer assist the domain expert in completion of a repertory grid as follows:

- first, the domain expert provides the elements (perhaps with the assistance of the concept mapping components of ICONKAT), and then the knowledge engineer elicits an initial collection of constructs on the basis of these elements using standard techniques (e.g. triads);
- the expert rates the elements on the constructs available thus far—at this point the knowledge engineer may ask GridNet (running in the background) to provide a feature set that replicates the expert's current grid;
- each of the qualitative features (in a given feature set) are presented to the expert as potential bipolar groupings in which the elements are clustered at the opposing poles, and the expert is asked to provide a construct for each grouping;

- the domain expert now returns to the repertory grid and rates the elements on the basis of the newly derived construct(s);
- the knowledge engineer may at any time invoke ICONKAT's analysis module (i.e. NICOD and OMNIGRID) to produce a rule-set or other analysis of the newly expanded repertory grid.

The feature set of bipolar element clusters illustrated in Figure 8 are derived from the repertory grid in Figure 1. The domain expert is asked to provide a bipolar construct to differentiate the element clusters from each other. Inasmuch as this clustering maximally separates the elements in some decision space, the proposed features may tend to be more useful in teasing out missing constructs (as reflected in an overall structure in the repertory grid) than features arising from less distinct sets of elements.

For example, consider the first feature vector as presented to the domain expert in Figure 8(a). GridNet asks the expert to furnish a bipolar construct which disinguishes the disease state, Abnormal Function with suspicion of Valve Disease, from the disease state, Nonspecific Wall Motion Abnormalities. In other words, "how is the first group of diseases (in this case only one) most significantly different from the second group?" If the expert cannot provide a construct for this feature, the next item in the feature set will be presented to the expert in a similar fashion. This process continues until the domain expert has provided a construct corresponding to one or more of the features presented.

In application to the nuclear cardiology domain, we have found this construct elicitation method successful in that the domain expert was able to provide a construct for both feature (b) and feature (c) from Figure 8. The construct that the domain expert associated with feature (b), *segmental/generalized,* is a high-level, rather pivotal, construct. Note the important position the domain expert accorded this new construct when he incorporated it into the concept map in Figure 6. The construct represented by feature (c), no ballerina foot/ballerina foot, is the fundamental distinction used by the domain expert to discriminate Mitral Valve Prolapse from other heart wall motion abnormalities. In summary, the GridNet approach to construct elicitation can assist the expert to expand the repertory grid in a meaningful way, thus producing an improved expert system.

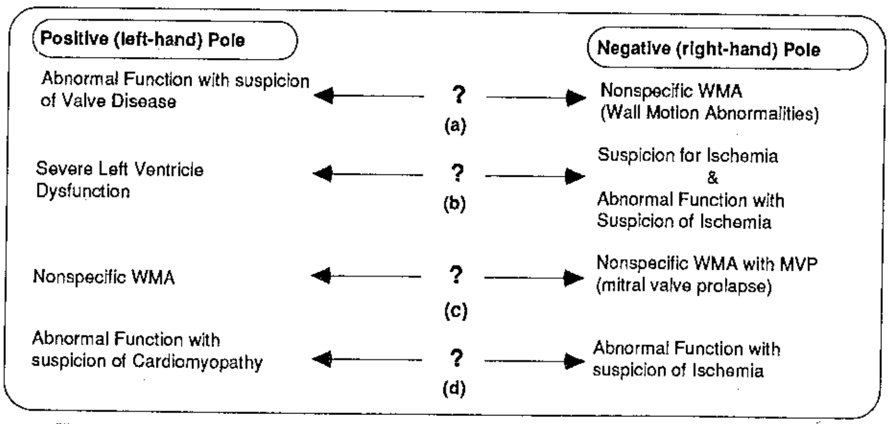

FIGURE 8. Extracted grid features (potential constructs) from the grid in Figure 1.

4.2.3. Integration of the elicitation subsystem: integrated how and why

The components of ICONKAT's elicitation subsystem are integrated on three levels, which we will briefly examine in a top-down fashion. In general, our efforts are directed toward the development of a unified elicitation subsystem informed by a constructivist approach to knowledge acquisition.

First, and most importantly, the principal components of ICONKAT's elicitation subsystem are integrated at the theoretical level. We (as tool-makers) have exploited the considerable epistemological overlap between Kelly's (1955) personal construct psychology and Ausubel's assimilation theory (Ausubel *et al.*, 1978), which enables repertory grids and concept maps to be integrated in a principled manner. Furthermore, this common theoretical framework, underlying a constructivist model of human representational processes, lends ICONKAT the kind of conceptual supports (guides) that a user's manual alone could never provide.

Second, in the ICONKAT system, repertory grids and concept maps have been integrated at the level of practice (domain level). The similarities of personal construct psychology and assimilation theory, as well as an appreciation of the differences in their intended ranges of application (i.e. psychology and education, respectively), empower the knowledge engineer to employ these otherwise orthogonal tools in a symbiotic fashion. Specifically, concept maps provide a hierarchically ordered, conceptual overview of the expert's model of the domain. Therefore, concept maps can provide knowledge landscapes (essentially topographical maps) of the domain that inform the knowledge engineer about the potential and appropriate use of grids within this domain. For example, the hierarchically-grouped concepts found in the elicited concept map may provide the elements for a single repertory grid, or a collection of grids, based on the hierarchical structure found in the map.

Third, the mediating representations (i.e. CMap/EMap and RepGrid) of ICONKAT's elicitation subsystem have been integrated at the representation level. CMap and RepGrid utilize the representation facilities (i.e. the modeling primitives or glue) of ICONKAT's modeling environment as the material basis for the construction of conceptual domain models.

4.3. THE MAINTENANCE SUBSYSTEM

A knowledge-based system typically evolves over its lifetime in a largely *ad hoc* fashion, thus potentially weakening the desired homomorphism with the domain expert's (and knowledge engineer's) original conception of the domain. This approach often leads to a "runaway" expert system, in which the actual knowledge base slides ever further out of conformance with the documentation and explanation modules, thus making the performance of its explanation facility unsatisfactory to users (see Section 4.4), and effectively rendering the knowledge base unmaintainable.

On the assumption that knowledge acquisition is an ongoing process (rather than a discrete phase), ICONKAT includes a maintenance subsystem. This subsystem provides support tools for use by the knowledge engineering team when testing the system's performance, refining the knowledge base, and maintaining the overall system. The maintenance subsystem includes two hypermedia components, Documentess and Tester, which fully exploit ICONKAT's mediating representations (e.g. concept maps, repertory grids) to furnish various perspectives of the evolving domain model as embodied in the modeling primitives.

4.3.1. Documentess (Documentation tool for expert systems)

Documentess is an interactive "intelligent" documentation tool for constructing and maintaining a domain model from which expert system knowledge bases can be derived. By exploiting the flexibility offered by hypermedia in structuring information, Documentess provides a single integrated interface to the conceptual domain model.

Essentially, Documentess furnishes a mechanism for hierarchically structuring a relatively unstructured collection of abstract data types (i.e. the modeling primitives or glue). It is through the Documentess interface that these comparatively labile and amorphous data structures comprising the intermediate representation are molded into a coherent domain model. Thus, it enables the underlying knowledge base to be viewed through the prism of a model largely constructed by the expert, and later reconstructed by the knowledge engineer.

For example, Documentess enables the user to shift levels of abstraction (up or down) from the elicited repertory grids, to the resulting rules expressed in a pseudo-natural language, to the actual rules formulated in the required syntactic formalism of the performance environment, and back. Documentess also associates the concepts found in the previously elicited concept maps with their related objects and classes as represented in the intermediate representation.

In summary, an appropriate hypermedia system integrated with an intermediate representation (the glue) enables the knowledge engineer to move smoothly between the methodological tools (i.e. mediating representations) employed in the knowledge acquisition process and the resulting rules and objects which comprise the knowledge base of the expert system under development.

4.3.2. Tester

Tester is a hypermedia-oriented system for evaluating the performance of knowledge-based systems. The tool volunteers data to the performance environment by extracting values as needed from either a database or a collection of spreadsheet files containing test cases. Through monitoring and recording the resulting inference process, the output is compared with the judgement made by the expert on the same data, thus providing a measure of the system's effectiveness.

Tester also includes an interactive mode by means of which the knowledge engineer can volunteer data either interactively or from the database of test cases. The system allows the knowledge engineer to monitor the session, interrupting the execution to examine the progress as necessary. Thus, this tool helps to automate partially the validation phase, a crucial step in the development of an expert system.

The interactive nature of Tester and the collection of test cases previously rated by the expert also serve as a foundation for the development of an instructional system for training users in the application domain. For example, the user is presented with all the relevant information (images, data, etc.) relating to a particular case that is typically available to practitioners in that domain. At this time, he or she is asked to render a diagnosis (or other significant judgement), which

Tester then compares with the conclusion arrived at by the domain expert about the same case. Alternatively, the user may request the expert system to diagnose the case in question. Upon learning of the result, the user may invoke the explanation subsystem (see Section 4.4) to illustrate how and why it reached the particular conclusion.

4.4. EXPLANATION SUBSYSTEM

The capacity of most current expert systems to explain their findings (i.e. conclusions), tends toward inadequate causal descriptions of the behavior of the performance environment's reasoning mechanism. A key to the design of explanation subsystems capable of more than shallow and/or mechanistic accounts is to recognize that the development of an explanation facility is a fundamental part of the knowledge acquisition process. Instead of arduously constructing a model of human expertise and then throwing it away (upon translation to the syntax of the performance environment), an explanation facility should exploit the mediating and intermediate representations formed during the knowledge acquisition process. Without these representations, the implicit connections that establish the "logical" structure of the domain model are lost, and as a result, much effort will be required to "put Humpty Dumpty back together again". This is precisely the task (i.e. reassembling Humpty) facing knowledge engineers who do not regard explanation as part of the knowledge acquisition process and, lacking an adequate model, attempt to build their explanation systems *post hoc.* As illustrated in Figure 9, the model resulting from the knowledge elicitation process is subsequently exported from the development environment to the delivery environment where it serves as the foundation for the explanation capability for the deployed system.

The explanation subsystem exploits a concept mapping interface, which is directly inherited from the domain expert's previous interactions with the elicitation subsystem. Specifically, the previously elicited concept maps provide conceptual guideposts as navigation aids to the user when browsing the knowledge landscape. Thus, changes to the overview (i.e. top level) concept map found in the model will be reflected in the explanation subsystem's user interface to that model.

Hence, in addition to its role in documentation and maintenance, ICONKAT's mediating and intermediate representations offer a flexible embodiment of the domain model for subsequent exploitation by the explanation subsystem. Because the explanation subsystem affords a domain model view of the knowledge base, explanations can be provided at a variety of conceptual levels. In response to a user-initiated request for explanation, the performance environment is interrupted and the user is switched into the context-sensitive explanation subsystem. At this point, users are free to explore the conceptual domain model until they are satisfied that they have discovered or constructed an adequate explanation.

5. Summary

The main purpose of this paper has been to describe the results of our research aimed at the development of an integrated constructivist knowledge acquisition system, ICONKAT. As noted at the outset, the most fundamental step in the knowledge acquisition phase of the construction of an expert system is the elicitation

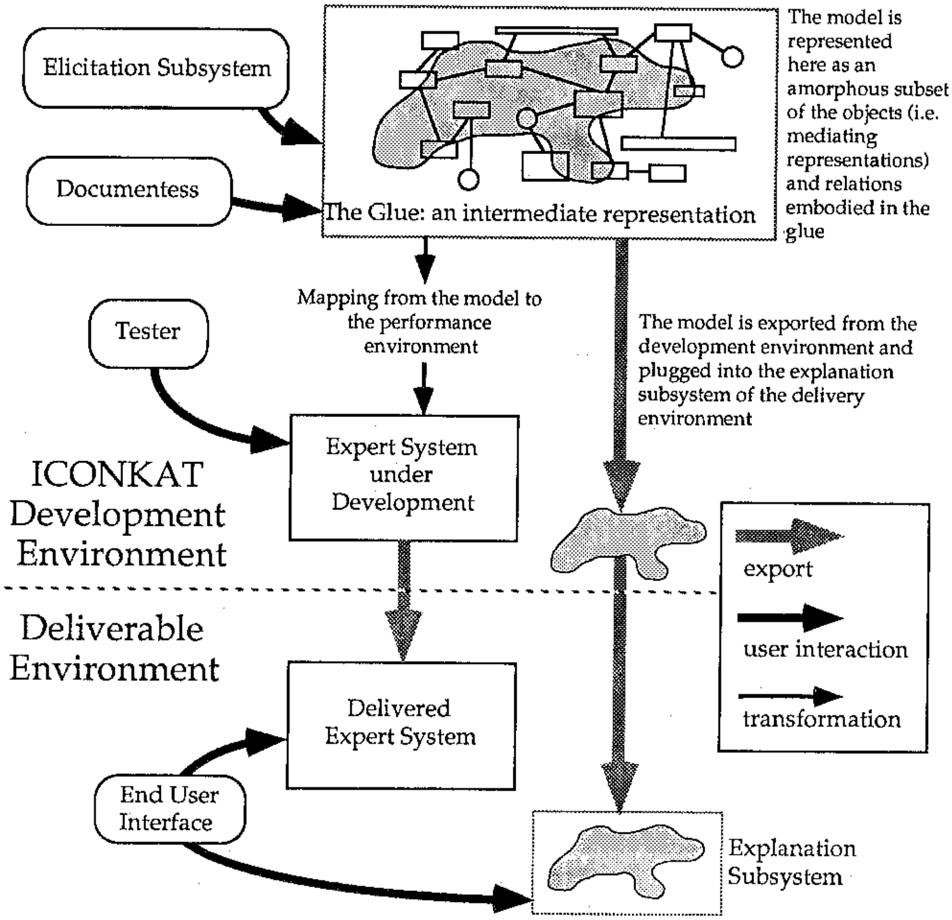

FIGURE 9. The transition from development to delivery.

of knowledge from a skilled individual. We have suggested elsewhere (Ford & Adams-Webber, 1991), that human "expertise" not only consists of explicit knowledge of the sort found in textbooks (i.e. widely shared consensual beliefs), but also includes a fund of personal experience comprising functional, but fallible, anticipations held with high confidence and uncertain validity (cf. Agnew & Brown, 1989). We have argued also that *ad hoc* combinations of techniques and tools cannot contribute significantly to resolving most of the difficulties associated with the knowledge acquisition process.

Specifically, we examined some of the epistemological assumptions of two related constructivist models of cognition. First, we explicated certain of the central principles of Kelly's personal construct theory and described its major methodological tool, repertory grids. Secondly, we summarized Ausubel's assimilation theory and described its major methodological tool, concept mapping.

Finally, we introduced a new constructivist knowledge acquisition tool (ICONKAT) supporting the entire knowledge acquisition process, which, in our view, encompasses elicitation, model construction, implementation, documentation, testing, maintenance and explanation of the domain model. The ICONKAT system architecture consists of three major subsystems; elicitation, maintenance and explanation.

The elicitation subsystem integrates repertory grids and concept maps as mediating representations. Its theoretical rationale was explicated in terms of some basic tenets of personal construct theory and assimilation theory. In addition, we discussed a novel method of construct elicitation exploiting neural nets to elicit new constructs on the basis of previously elicited repertory grid data. Integrated in a principled manner, the resulting knowledge elicitation system is designed to bring the expert's self-constructed knowledge to the surface.

The maintenance subsystem of ICONKAT provides support tools for use by the knowledge engineering team when testing the system's performance, refining the knowledge base and maintaining the overall system. Documentess provides a single integrated environment for interacting with the various informational components of the knowledge acquisition process (as embodied in the conceptual domain model). Tester evaluates the performance of knowledge-based systems, thus partially automating the validation phase.

In addition, we described an approach to explanation in which the conceptual domain model emerging from the knowledge acquisition process is exported to the delivery environment where it serves as the foundation for the explanation capability of the deployed system. Users are encouraged to construct their own explanations by exploring the conceptual domain model until they are satisfied that they have found an adequate explanation.

ICONKAT is currently employed in the design and construction of an expert system for the diagnosis of first pass cardiac functional images, a non-invasive radionuclide technique used to evaluate heart wall motion abnormalities.

We anticipate that personal construct theory, assimilation theory and other constructivist models of cognition will continue to serve as useful conceptual supports (guides) for the development of new practical tools for knowledge acquisition.

Many people have influenced the development of this work. We owe particular thanks to Jim Andrews, Jeff Bradshaw, John Brennan, John Coffey, Jack Fleming, Scott Hunter, Johnny Liseth, Steve Venticique, and Jeff Yerkes. This work was supported in part by Florida High Technology Enhancement Funds.

References

ADAMS-WEBBER, J. R. (1970). Actual structure and potential chaos. In D. BANNISTER, Ed. *Perspectives in Personal Construct Theory*, pp. 30–45. London: Academic Press.

ADAMS-WEBBER, J. R. (1979). *Personal Construct Theory: Concepts and Applications*. New York: Wiley.

ADAMS-WEBBER, J. R. (1987). Repertory grid technique. In R. CORSINI, Ed. *Concise Encyclopedia of Psychology*, pp. 973–974. New York: Wiley Interscience.

ADAMS-WEBBER, J. R. (1989). Kelly's pragmatic constructivism. *Canadian Psychology*, **30**, 190–193.

ADAMS-WEBBER, J. R. (1990). Some fundamental asymmetries in the structure of personal constructs. In G. NEIMEYER & R. A. NEIMEYER, Eds. *Advances in Personal Construct Psychology*, pp. 49–85. Greenwich, CT: JAI Press.

Adams-Webber, J. R. & Mancuso, J. C. (1983). The pragmatic logic of personal construct psychology. In J. R. Adams-Webber & J. C. Mancuso, Eds. *Applications of Personal Construct Psychology*, pp. 1–10. New York: Academic Press.

Agnew, N. M. & Brown, J. L. (1989). Foundations for a theory of knowing: II. Fallible but functional knowledge. *Canadian Psychology*, **30,** 168–183.

Ausubel, D. P. (1963). *The Psychology of Meaningful Verbal Learning.* New York: Grune and Stratton.

Ausubel, D. P., Novak, J. D. & Hanesian, H. (1978). *Educational Psychology: A Cognitive View* (2nd Edit). New York: Holt, Reinhart & Winston. Reprinted (1986), New York: Warbel & Peck.

Boose, J. H. (1986). Rapid acquisition and combination of knowledge from multiple domain experts in the same domain. *Future Computing Systems*, **1,** 191–216.

Bradshaw, J. M. & Boose, J. H. (1990). Decision analysis techniques for knowledge acquisition: Combining information and preferences using AQUINAS and AXOTL. *International Journal of Man–Machine Studies*, **32,** 121–186.

Bradshaw, J. M. & Woodward, B. (1989). Report of the working group on knowledge acquisition tools. In J. Boose & B. Gaines, Eds. *Knowledge Acquisition for Knowledge-Based Systems Newsletter*, **1,** 12–13.

Bundy, A. (1985). Incidence calculus: A mechanism for probabilistic reasoning. *Journal of Automated Reasoning*, **1,** 263–283.

Butler, K. A. & Corter, J. E. (1986). Use of psychometric tools for knowledge acquisition: A case study. In W. G. Gale, Ed. *Artificial Intelligence and Statistics*, pp. 295–319. Reading, MA: Addison-Wesley.

Ford, K. M. (1987). *An Approach to the Automated Acquisiton of Production Rules from Repertory Grid Data.* PhD dissertation, Tulane University.

Ford, K. M. (1989). A constructivist view of the frame problem in artificial intelligence. *Canadian Psychology*, **30,** 188–190.

Ford, K. M. & Adams-Webber, J. R. (1991). Knowledge acquisition and constructivist epistemology. In R. Hoffman, Ed. *The Psychology of experts: Cognitive Research and Empirical AI.* New York: Springer-Verlag.

Ford, K. M., Bradshaw, J. M. & Adams-Webber, J. R. (In press). Knowledge acquisition as a modeling activity. *International Journal of Intelligent Systems.*

Ford, K. M., Petry, F., Adams-Webber, J. R. & Chang, P. J. (1991). An approach to knowledge acquisition based on the structure of personal construct systems. *IEEE Transactions on Knowledge and Data Engineering*, **3,** 78–87.

Gaines, B. R. & Shaw, M. L. G. (1980). New directions in the analysis and interactive elicitation of personal construct systems. *International Journal of Man–Machine Studies*, **13,** 81–116.

Gaines, B. R. & Shaw, M. L. G. (1986). Induction of inference rules for expert systems. *Fuzzy Sets and Systems*, **18,** 315–328. Amsterdam: Elsevier Science Publishers.

Hunter, S., Stahl, H. A. & Novak, J. D. (1990). *CMap: A Concept Mapping Elicitation Tool*, Technical Report #90–006. Cornell University: Department of Education.

Johnson, N. E. (1989). Mediating representations in knowledge elicitation. In D. Diaper, Ed. *Knowledge Elicitation: Principles, Techniques and Applications.* New York: Wiley.

Jones, J. C. & Ford, K. M. (1991). A neural net approach to knowledge elicitation from repertory grids. In M. Fishman, Ed. *Advances in Artificial Intelligence Research*, vol. 2. Greenwich, CT: JAI Press.

Kelly, G. A. (1955). *The Psychology of Personal Constructs*, vols. 1 & 2. New York: W. W. Norton.

Kelly, G. A. (1969*a*). A mathematical approach to psychology. In B. Maher, Ed. *Clinical Psychology and Personality*, pp. 7–45. New York: Wiley.

Kelly, G. A. (1969*b*). The role of classification in personality theory. In B. Maher, Ed. *Clinical Psychology and Personality*, pp. 289–300. New York: Wiley.

Kelly, G. A. (1970). A brief introduction to personal contruct theory. In D. Bannister, Ed. *Perspectives in Personal Construct Theory*, pp. 1–29. London: Academic Press.

Mancuso, J. C. & Eimer, B. N. (1982). Fitting things into sorts: The range corollary. In J.

C. Mancuso & J. R. Adams-Webber, Eds. *The Construing Person,* pp. 130–151. New York: Praeger.

Mitterer, J. & Adams-Webber, J. R. (1988). OMNIGRID: A program for the construction, administration, and analysis of repertory grids. In J. C. Mancuso & M. L. G. Shaw, Eds. *Cognition and Personal Structure,* pp. 91–106. New York: Praeger.

Novak, J. D. (1977). *A Theory of Education.* Ithaca, NY: Cornell University Press.

Novak, J. D. (1989). Helping students learn how to learn: a view from a teacher-researcher. Paper presented to the *Third Congress on Research and Teaching of Science and Mathematics,* Santiago de Compostela, Spain.

Novak, J. D. & Gowin, D. B. (1984). *Learning How to Learn.* New York: Cambridge University Press.

Tversky, A. (1977). Features of similarity. *Psychological Review,* **84,** 372–352.