

Knowledge acquisition tools based on personal construct psychology

BRIAN R. GAINES AND MILDRED L. G. SHAW

Knowledge Science Institute, University of Calgary, Calgary, Alberta, Canada T2N 1N4 gaines@cpsc.ucalgary.ca & mildred@cpsc.ucalgary.ca

Abstract

Knowledge acquisition research supports the generation of knowledge-based systems through the development of principles, techniques, methodologies and tools. What differentiates knowledge-based system development from conventional system development is the emphasis on in-depth understanding and formalization of the relations between the conceptual structures underlying expert performance and the computational structures capable of emulating that performance.

Personal construct psychology is a theory of individual and group psychological and social processes that has been used extensively in knowledge acquisition research to model the cognitive processes of human experts. The psychology takes a constructivist position appropriate to the modelling of human knowledge processes, but develops this through the characterization of human conceptual structures in axiomatic terms that translate directly to computational form. In particular, there is a close correspondence between the intensional logics of knowledge, belief and action developed in personal construct psychology, and the intensional logics for formal knowledge representation developed in artificial intelligence research as term subsumption, or KL-ONE-like, systems.

This paper gives an overview of personal construct psychology and its expression as an intensional logic describing the cognitive processes of anticipatory agents, and uses this to survey knowledge acquisition tools deriving from personal construct psychology.

1 Introduction

Knowledge-based systems development is targeted on the emulation of human high-level skilled performance in a program running on a digital computer. Knowledge acquisition research supports the generation of such programs through the development of principles, techniques, methodologies and tools. Knowledge acquisition is not a monolithic process but, like all software engineering, draws on many sources of information in diverse forms, such as specifications, experience, principles, laws, observation, and so on, recorded in a variety of media. A major part of the task of the knowledge-based system designer is to collect this diverse material, organize it effectively, and develop from it structures suitable for computer implementation.

What distinguishes software engineering for knowledge-based systems from that for conventional systems is the relative emphasis on two aspects of the development process. First, the knowledge-based system approach emphasizes the use of human experts as the primary sources of information. Second, it emphasizes the use of knowledge representation schema as the primary basis of implementation. From an integrative viewpoint, one can see these as particular foci of attention, rather than radical changes in approach: human experts have always been an important source of information in system development; and knowledge representation schema may be regarded as the culmination of a trend towards the use of very high level, declarative specification languages. What has differentiated knowledge-based system development from conventional system development is the emphasis on in-depth understanding and formalization of the relations

between the conceptual structures underlying expert performance and the computational structures capable of emulating that performance. That is, the relationship between the psychology of the expert performer and the ontology of the computational emulator has been viewed as a research challenge on the one hand, and as the basis of supporting practical system development on the other.

We have been very careful in the wording of the preceding paragraphs to avoid presuppositions about the nature and status of ‘conceptual structures underlying expert performance’ and about the intrinsic possibility and scope of ‘emulation of human high-level skilled performance’. In particular, nothing is presupposed about the reality and location of conceptual structures, and neither is anything presupposed about the suitability of such structures for modelling or emulating any or all aspects of human expertise. These are deep issues important to future directions in research and practice relating to knowledge-based systems, and we and others have addressed them in depth elsewhere (Clancey, 1991; Clancey & Roschelle, 1991; Gaines, 1992; Gaines, Shaw & Woodward, 1992). This paper adopts a pragmatic and retrospective stance, reporting on what has been achieved in the past decade in knowledge acquisition research based on personal construct psychology.

Figure 1 summarizes the issues addressed above, and gives an overview of the current state of the art described in this paper. At the top left is the person regarded as capable of expert performance of some task in some domain. At the top right is the computer regarded as capable of emulating the performance of that task in that domain. The arrow from person to computer indicates that an objective of knowledge acquisition research is to support the required expertise transfer from the person to the computer.

Research on knowledge-based systems has emphasized the importance of understanding the conceptual structures underlying expertise, and of having knowledge representation schema that represent them as directly as possible. On the far left of Figure 1 is shown the psychological model of the person that has been the basis of the work surveyed in this paper, Kelly’s personal construct psychology (Kelly, 1955). This psychological model focuses on people as anticipatory systems,




Figure 1 Psychological and computational foundations for theories, methodologies and tools supporting expertise transfer

themselves developing conceptual models so as to better understand, predict and control a world. On the far right of Figure 1 is shown the ontological model of the computational knowledge representation that has been the basis of the work surveyed in this paper, Brachman's term subsumption model for KL-ONE-like schema (Brachman & Schmolze, 1985). This ontological model focuses on computers as anticipatory systems, using conceptual models to understand, predict and control a world. The arrow at the bottom of Figure 1 linking the psychological and computational models indicates that one significant aspect of the work reported in this paper is the unification of these two frameworks. There is a very direct relationship between the psychological conceptual structures presupposed in personal construct psychology and those implemented in term subsumption logics, and this is highly significant in supporting expertise transfer as a process of modelling the basis of human skilled performance in operational terms.

The central section of Figure 1 gives an overview of the theories, methodologies and tools presented in this paper. The central loop from person to computer shows: knowledge elicitation tools used to develop a psychological model of skilled performance through representations of the skills in terms of conceptual structures; knowledge modelling tools used to develop this to be a computational model of the skill in terms of logical structures; and knowledge structuring tools used to further develop this to be an operational model that will run within a particular computer environment.

At the very centre of Figure 1 is the most important aspect of supporting expertise transfer effectively, that of providing feedback to the human expert for purposes of validating that the knowledge structures at each stage of the transfer properly represent the basis of the skilled performance, and for verifying them operationally as far as possible at every stage of development. It is an important side-effect of the emphasis on the human expert in knowledge-based system development that much of the functionality of the tools involved is targeted on supporting human understanding of the development process. This is again consistent with trends in the support of conventional software engineering such as tools for 'requirements tracking', and knowledge-based system development can be seen within this framework as focusing on the way in which the requirements themselves develop and are refined through the very process of system development. That is, in the terms of the discussion above, the tools supporting 'expertise transfer' are developing and refining a model of the basis of the expertise as a specification for a system to emulate it. Expertise is being transferred through the development of a model in terms of knowledge that provides a computational basis for emulating that expertise.

2 Personal construct psychology and its applications

George Kelly was a clinical psychologist who lived between 1905 and 1967, published a two volume work (Kelly, 1955) defining personal construct psychology in 1955, and went on to publish a large number of papers further developing the theory many of which have been issued in collected form (Maher, 1969). Kelly was working on a second book when he died, of which, unfortunately, only the preface has been published (Kelly, 1970), and he also became involved with the computer simulation of personality (Kelly, 1963). However, his work never became part of the mainstream cognitive science literature, although it has attracted widespread attention and application in management, education and clinical psychology (Shaw, 1981).

Kelly was a keen geometer with experience in navigation and an interest in multi-dimensional geometry. When he came to formalize his theory he took as his model Euclid's *Elements* and axiomatized personal construct psychology as a *fundamental postulate* together with eleven *corollaries*, terming the primitives involved *elements* and *constructs*. Kelly presented his theory as a *geometry of psychological space* (Kelly, 1969), and his conceptual framework is very clear if seen in these terms. It may seem strange to base cognitive science on geometry rather than logic until one remembers that the reasoning structure of the *Elements* was the basis for both Greek and modern logic (Lachterman, 1989), and that geometry and logic in a category-theoretic framework are equivalent (Mac Lane, 1971). What Kelly achieved through the use of geometry was an intensional

logic, one in which predicates are defined in terms of their properties rather than extensionally in terms of those entities that fall under them. Logics of knowledge and belief are essentially intensional (Hintikka, 1962), and in his time there were no adequate formal foundations for intensional logic. It was not until 1963 that Hintikka published the model sets formulation that gave intensional logic its *possible worlds* formal foundation (Hintikka, 1963), and hence formal foundations for cognitive science in logical terms only became possible in the late 1960s. The intensional nature of semantic networks in artificial intelligence was recognized in the late 1970s (Woods, 1975; Brachman, 1977; Shapiro, 1979), and their philosophical and logical structure as cognitive models has been detailed by Zalta (1988). In a computational context, Kelly's approach is attractive in leading directly to representations and algorithms, and its rephrasing in terms of modern intensional logic leads to a model of human psychological processes that readily transfers to their support and emulation by computers.

Kelly's 'repertory grid' methodology for eliciting conceptual structures had become a widely used and accepted technique for knowledge elicitation, and has been implemented as a major component of many computer-based knowledge acquisition systems. The grid itself dates back to Kelly's (1955) application of the *personal construct psychology* that he had developed on the basis of his clinical and teaching experience. A comprehensive computer-based elicitation and analysis system for repertory grids was developed by Shaw (1979) with applications mainly in educational, clinical and management studies. Gaines and Shaw (1980) suggested that repertory grids would provide a useful development technique for expert systems, and later published a validation study of the elicitation of expertise from accountants and accounting students using computer-based repertory grid elicitation (Shaw & Gaines, 1983). Boose, in an independent parallel study, reported success in a wide range of industrial expert system developments using computer elicitation of repertory grids (Boose, 1984), and since then many knowledge acquisition systems have incorporated repertory grids as a major elicitation technique (Boose & Bradshaw, 1987; Diederich, Ruhmann & May, 1987; Garg-Janardan & Salvendy, 1987; Shaw & Gaines, 1987; Ford et al., 1990).

The repertory grid methodology has evolved in the light of application experience, and now has major differences from that described by Kelly. Shaw took advantage of the processing power and interactivity of computers to introduce on-line analysis and feedback to the person from whom the grid was being elicited (Shaw, 1980). In expert systems terms, this feedback can be seen as highlighting correlations that might be spurious and lead to incorrect rules in later analysis. Shaw and Gaines introduced new forms of analysis of the repertory grid based on fuzzy sets theory (Shaw, 1979) which became the basis of rule extraction (Gaines & Shaw, 1986). Boose and Bradshaw (1987) made changes to the grid structure introducing hierarchical data structures to cope with more complex domains. Bradshaw, et al. (1988) showed how many problems that did not seem appropriate to repertory grids could be formulated in terms of them.

The repertory grid was an instrument designed by Kelly to bypass cognitive defences and give access to a person's underlying construction system by asking the person to compare and contrast relevant examples (significant people in the person's life in the original application). This use of relevant examples has also been important in the application of repertory grids to expertise modelling, it is often easier and more accurate for the expert to provide critical cases rather than a domain ontology. The power of a few critical cases described in terms of relevant attributes to build a domain ontology is remarkable but statistically well-founded (Gaines, 1989), and may be seen as the basis of the success of repertory-grid techniques in knowledge elicitation.

The original repertory grid methodology is based primarily on only one aspect of Kelly's personal construct psychology, his *dichotomy corollary* that "A person's construction system is composed of a finite number of dichotomous constructs." The standard grid is a flat structure of elements described in terms of dichotomous attributes or constructs that does not represent the hierarchical structure of Kelly's organization corollary that "Each person characteristically evolves, for his convenience of anticipating events, a construction system embracing ordinal relationships between constructs." Hinkle developed a technique of *laddering*, based on 'why' and

'how' questions, for investigating ordinal relations between constructs (Hinkle, 1965), and Boose incorporated a laddering tool in ETS (Boose, 1986a). However, ordinal relations between constructs were not the primary focus in initial applications of repertory grid tools.

This emphasis on dichotomous constructs alone changed as the second generation toolbench, AQUINAS (Boose & Bradshaw, 1987), was developed in the light of experience with ETS, and hierarchical structures of tasks, experts, elements and constructs were introduced into the data structures and interfaces. It also changed as conceptual induction techniques were used to derive hierarchical concept structures from the rules extracted from repertory grids (Gaines & Shaw, 1992). Recently, the intensional logic underlying the psychological primitives of personal construct psychology has been developed in detail (Gaines & Shaw, 1990), and this has been used to develop knowledge acquisition tools based on a visual language that corresponds to a formal semantics for semantic nets (Gaines, 1991(c)).

These later developments suggest that personal construct psychology can also provide new foundations for existing tools in which ordinal relations are a primary focus, such as those that use some form of semantic network to build an overall domain and task ontologies directly. These include a wide range of significant knowledge acquisition tools such as MOLE (Eshelman, et al., 1987), KNACK (Klinker et al., 1987), SALT (Marcus, 1987), Cognosys (Woodward, 1990), KEATS (Motta et al., 1988), and CAMEO (Jones, 1989). The tools vary in their interfaces, sources, emphasis on declarative or procedural knowledge, and domain dependence, but all result in domain models on a network of directed relations between concepts. These tools are different in approach from those based on repertory grids, but as the trend towards the integration of tools and techniques continues it is becoming important to understand the relationship between them. Similar considerations apply to other tools now being incorporated in integrated knowledge acquisition systems such as those for empirical induction and conceptual clustering.

3 Knowledge representation in personal construct psychology

Kelly's 'fundamental postulate' for personal construct psychology (Kelly, 1955) was that:

A person's processes are psychologically channelized by the way in which he anticipates events. (p. 46)

This was stated as a postulate to emphasize that it was presented as a convenient viewpoint from which to understand human behaviour, not imputed to an underlying physiological or psychological reality. Kelly saw all men as 'personal scientists' in anticipating the world, and attempted to develop techniques where this anticipatory modelling activity was reflexively applied to the self. His first corollary, the construction corollary, states:

A person anticipates events by construing their replications. (p. 50)

This emphasis on the role in behaviour of a view to the future is what distinguishes Kelly's approach to psychology. He saw man as driven by the need to cope with coming events in the world and all other aspects of behaviour as deriving from this:

A person's processes, psychologically speaking, slip into the grooves which are cut out by the mechanisms he adopts for realizing his objectives. (p. 49)

These grooves provide templates for construing events which he termed 'personal constructs':

Man looks at his world through transparent templates which he creates and then attempts to fit over the realities of which the world is composed (pp. 8-9)

Constructs are used for predictions of things to come, and the world keeps on rolling on and revealing these predictions to be either correct or misleading. This fact provides the basis for the revision of constructs and, eventually, of whole construct systems. (p. 14)

Kelly introduces the notion of a *psychological space* as a term for a region in which we may place and classify elements of our experience. It is important to note that he did not suppose this space to

pre-exist as a world of such elements, but rather to come into being through a process of construction by which we create a space in which to place elements as we come to construe them. He sees us as creating dimensions in personal psychological space as a way of providing a coordinate system for our experience, and emphasizes that the topology of the space comes into existence as it is divided:

Our psychological geometry is a geometry of dichotomies rather than the geometry of areas envisioned by the classical logic of concepts, or the geometry of lines envisioned by classical mathematical geometries. Each of our dichotomies has both a differentiating and an integrating function. That is to say it is the generalized form of the differentiating and integrating act by which man intervenes in his world. By such an act he interposes a difference between incidents—incidents that would otherwise be imperceptible to him because they are infinitely homogeneous. But also, by such an intervening act, he ascribes integrity to incidents that are otherwise imperceptible because they are infinitesimally fragmented. In this kind of geometrically structured world there are no distances. Each axis of reference represents not a line or continuum, as in analytic geometry, but one, and only one, distinction. However, there are angles. These are represented by contingencies or overlapping frequencies of incidents. Moreover, these angles of relationship between personal constructs change with the context of incidents to which the constructs are applied. Thus our psychological space is a space without distance, and, as in the case of non-Euclidean geometries, the relationships between directions change with the context. (Kelly, 1969)

It is this emphasis on the space itself being created by a process of making distinctions rather than being defined by the elements distinguished that gives personal construct psychology its intensional nature:

the construct is a basis of making a distinction . . . not a class of objects, or an abstraction of a class, but a dichotomous reference axis. (Kelly, 1970)

One obvious question about Kelly's use of the term construct was how it differs from the more conventional term 'concept'. He discusses this in the following terms:

We use the term *construct* in a manner which is somewhat parallel to the common usage of 'concept'. However, if one attempts to translate our *construct* into the more familiar term, 'concept', he may find some confusion. We have included, as indeed some recent users of the term 'concept' have done, the more concretistic concepts which nineteenth-century psychologists would have insisted on calling 'percepts'. The notion of 'percept' has always carried the idea of its being a personal act—in that sense, our *construct* is in the tradition of 'percepts.' But we also see our construct as involving abstraction—in that sense our *construct* bears a resemblance to the traditional usage of 'concept'. . . . Now when we assume that the construct is basically dichotomous, that it includes percepts, and that it is a better term for our purposes than the term 'concept,' we are not quarreling with those who would use it otherwise. Within some systems of logic the notion of contrast as something distinct from irrelevancy is not part of the assumptive structure. We, on the other hand are simply assuming that this is the way people do, in fact, think. (Kelly, 1955)

Since the term 'concept' is used in psychology and in knowledge representation in somewhat different ways, it is useful to define it clearly in these contexts for the purposes of this paper. A psychological concept is defined to be that mental entity *imputed* to a distinction making agent as enabling it to make a particular distinction. Not that concepts are separated both from the distinctions they support and the entities they distinguish, and are not reified but seen as imputed to the agent. They are themselves distinctions made by an observer—concepts are state variables we impute to a knowledgeable agent. This definition also corresponds to Anglin's:

a concept is all of the knowledge possessed by an individual about a category of objects or events—Concepts mediate categorization but concepts are not the resultant categories. (Anglin, 1977)

A computational knowledge representation concept is defined to be that data structure used by a distinction making algorithm to enable it to make a particular distinction.

The dichotomous aspect of constructs is the most significant aspect of the difference between Kelly's constructs and current usage of the term 'concept'. His *dichotomy corollary* states this (Kelly, 1955):

A person's construction system is composed of a finite number of dichotomous constructs. (p. 59)

and it is a consequence of the two-sided nature of a distinction represented in the geometry. That people tend to conceptualize the world in terms of restructured sorts that are then dichotomized is a phenomenon identified in antiquity (Lloyd, 1966), and is common across many cultures (Maybury-Lewis & Almagor, 1989).

The taxonomic, abstraction, or subsumption, hierarchy between concepts is Kelly's organization corollary (Kelly, 1955):

Each person characteristically evolves, for his convenience of anticipating events, a construction system embracing ordinal relationships between constructs. (p. 56)

He uses this ordinal relation in the development of the psychology to model the dynamics of change in conceptual systems. For example, that we have 'core constructs' that we are very reluctant to change because of the dependencies, such as subsumption, that exist within our construction systems.

4 The intensional logic of personal construct psychology

It is very easy and highly instructive to move from Kelly's topological geometry to a corresponding intensional logic of knowledge representation. We may take Kelly's notion of a distinction as primitive, and see how distinctions may relate to each other in psychological space. If an element is placed within the region carved out by a distinction, then we may say that the distinction is *asserted* to apply to the element. If one distinction carves out a region that contains that carved out by another, then the first distinction may be said to *subsume* the second. If one distinction carves out a region that does not overlap that carved out by another, then the first distinction may be said to be *disjoint* to the second. These relations are in themselves sufficient to define an intensional logic of distinctions in that more complex relations may be composed from them.

The subsumption and disjoint relations may be defined in an algebraic formalism by representing distinctions by bold lower case letters such that a distinction applied to another distinction is concatenated to the right of it. Then the definition above translates as one distinction will be said to *subsume* another if it can always be applied whenever the other can. It can be represented formally as:

$$\text{'b subsumes a'} \quad a \rightarrow b \Leftrightarrow \vdash xa \Rightarrow \vdash xb \quad (1)$$

That is, **b** subsumes **a** if and only if, whenever one asserts **xa**, one also asserts **xb**. The definition is to be read intensionally in terms of a *commitment* to the way in which distinctions will be made, such that if **a** is made then there is a commitment to **b** being made also. This is why the form $\forall x$ is avoided—the notion of all the distinctions to which **a** and **b** may be applied is not well-defined.

Subsumption corresponds to increasing generality, since the subsuming distinction can be applied to at least as many things as that subsumed. In (1) concept **a** is said to be *subordinate* to concept **b**, and **b** *superordinate* to **a**. Subsumption is an asymmetric, transitive relation, a partial pre-order, over distinctions, that supports the ordinal relations of Kelly's organization corollary. This notion of subsumption is adequate to capture Kelly's use of the term that one construct subsumes another, and also the use of the same term in knowledge representation, that one concept subsumes another. Subsumption between computational concepts corresponds to the 'is-a' relation in knowledge representation schema. The interpretation of subsumption in terms of commitment above corresponds to the definitional form of the 'is-a' relation. The computed form of 'is-a' requires some further structures, which are developed in the next section, when primitive and non-primitive concepts are differentiated.

The disjoint relation is definable in similar terms, that one distinction is disjoint with another in that one can never be applied whenever the other can. It can be represented formally as:

$$\text{'a disjoint b'} \quad \mathbf{a} \text{---} \mathbf{b} \Leftrightarrow \vdash \mathbf{x}\mathbf{a} \Rightarrow \neg \vdash \mathbf{x}\mathbf{b} \quad (2)$$

that is, **a** is disjoint with **b** if and only if, whenever one asserts **xa**, one does not assert **xb**. The definition is again to be read intensionally in terms of a commitment to the way in which distinctions will be made, such that if **a** is made then there is a commitment to **b** not being made. Disjoint is a symmetric, intransitive relation over distinctions, and support Kelly's *dichotomy corollary* and the definition of disjoint concepts in knowledge representation.

It is interesting to note that definition (2) is an asymmetric definition of what is clearly a symmetric relation. Logically, this is possible because the reverse implication can be derived from (2), i.e., if one asserts **xb** one cannot assert **xa** because that would imply $\neg \vdash \mathbf{x}\mathbf{b}$. This derivation of symmetry from asymmetry may be logically simple, but is not semantically trivial. In terms of knowledge representation it corresponds to the essential sequence of definitions: if we define **a** first we cannot define it to be disjoint with **b** because **b** is not yet defined. Psychologically, this asymmetry appears to be related to the empirical asymmetries Adams–Webber has observed in the use of the, apparently symmetric, poles of a construct (Adams–Webber, 1979).

The $\rightarrow$ and --- relations are complementary in establishing four possible binary relations between distinctions, that $\mathbf{a} \rightarrow \mathbf{b}$, $\mathbf{b} \rightarrow \mathbf{a}$, $\mathbf{a} \text{---} \mathbf{b}$, or none of these. The two subsumption relations can hold together giving an equivalence relation on distinctions. The disjoint relation is incompatible with the subsumptions relation, and is *inherited* through subsumption, i.e.:

$$\mathbf{a} \text{---} \mathbf{b} \quad \mathbf{c} \rightarrow \mathbf{a} \Rightarrow \mathbf{c} \text{---} \mathbf{b} \quad (3)$$

4.1 A visual language for the logic

The arrow and line notion adopted in definitions (1) and (2) translates nicely to a graphical notation defining a *visual language* for the logic (Gaines, 1991c), a tool for which has been used to draw the structures in Figure 2. As shown at the top of Figure 2, Kelly's 'construct' in psychological space can be represented by a pair of disjoint concepts corresponding to what he terms the construct 'poles', both subsumed by a third concept corresponding to what he terms the 'range of convenience'. It is this fundamental conceptual unit, or templet that we fit over the world, being a pair of disjoint concepts applied to a restricted domain that characterizes Kelly's use of the logic as a foundation for cognitive psychology. In logical terms, he emphasizes the importance of *opposition* as relative negation applied within a context, rather than absolute negation free of any context. The psychological unit is the triple of concepts in the relation shown, rather than the individual concept, or logical predicate, in isolation.

At the centre of Figure 2, the abstract components of a concept are given specific instances to exemplify their application. 'Evaluable' things may be classified into two disjoint classes, 'good' and 'bad'. The emphasis on dichotomous concepts may give the impression that constructs are essentially binary in nature. However, at the bottom of Figure 2 is shown how what Kelly terms 'shades of gray' arise naturally through the addition of related concepts compatible with the original dichotomy. The dichotomy has been split into two such that 'bad' is now disjoint both from 'good' and 'fairly good', and 'good' is now disjoint from both 'bad' and 'fairly bad'. 'Mediocre' has been added as an additional concept intermediate between 'good' and 'bad', defined as 'fairly good' and 'fairly bad'. In tools such as the repertory grid, these intermediate concepts are represented on a numeric scale, as shown under the bottom structure of Figure 2.

The structures in Figure 2 are simple semantic networks in the style of KL-ONE (Brachman & Schmolze, 1985) or KRS (Gaines, 1991a), but they have well-defined logical semantics as defined above, and also strong psychological foundations in personal construct psychology. There is an analogy between the visual language and the representation of chemical structures as atoms and bonds. Distinctions are the atomic primitives in personal construct psychology, and further constructions may be seen as complex 'molecules' formed by distinctions joined through subsumption and disjoint 'bonds'. For example, the complex structure at the bottom of Figure 2 may be seen as the composition of two of the basic construct structures shown at the top.

Multiple constructs in psychological space correspond to multiple axes of reference, and the planes representing their distinctions and ranges of convenience intersect to define regions of the space corresponding to conjunction, composition and multiple inheritance in the logic, as shown at the top of Figure 3.

There is an important distinction to be made between the concepts defined by basic distinctions and those defined by intersections. The former are said to be *primitive concepts* and the latter non-primitive, or computed, concepts. In the visual language, primitive concepts are distinguished by having a small internal horizontal line at their left and right edges. Primitive concepts are incompletely defined in that we have complete freedom of choice as to where to place an element relative to the planes defining their distinction and range of convenience. However, no such freedom exists for non-primitive concepts, since they are defined as the intersection of primitive concepts. Logically, we have to *assert* that a primitive concept applies to an element, whereas we can either assert that a non-primitive applies or *recognize* that it applies through the previous assertions of the primitives that define it. In knowledge representation this recognition is termed *classification* (Borgida et al., 1989).

The definition of subsumption in (1) applies to non-primitive concepts, but it is no longer a matter of direct commitment but rather of derivation from the composition of commitments for concepts defining the intersection. In knowledge representation terms, the ‘is-a’ relation for non-primitive concepts is computable rather than definable—the commitment to their definition in terms of their structure entails a commitment to a derived, rather than a defined, ‘is-a’ relation. Confusion about these two forms of concept, and associated ‘is-a’ relations, caused problems in early developments of semantic nets (Brachman, 1983).




Figure 2 Representation of abstract and specific constructs and scales in a visual language for specifying definitions and assertions in the intensional logic




Figure 3 Concepts defined in terms of others, and their application to representing anticipations as rules supporting inference

Kelly’s theory of anticipation is based on attaching significance to such recognizable intersections:

What one predicts is not a fully fleshed-out event, but simply the common intersect of a set of properties. (Kelly, 1955)

The logic remains intensional because there is no implication that elements have already been construed within the intersections. The attachment of an anticipation to the intersect correspond to a commitment to place an element that falls in this intersect in the region defined by the pole of some other construct also. In logic, this is a *material implication* rather than an entailment, in that it is not necessitated by the way in which the distinctions are defined but is instead an auxiliary commitment or *rule*. Rules allow a cognitive system to be anticipatory in containing structures which from one set of distinctions made about an event will imply that others should be made leading to prediction or action. Rules play a similar role in computational systems in generating recommendations for decision or action. Overtly modelling the conceptual system of an expert as such a structure is a basis for emulating the expert’s performance in a knowledge-based system.

As shown in Figure 3, Kelly’s model of anticipation is represented in the visual language by an additional primitive, a rectangle with vertical bars, representing material implication or a rule. The rule in the centre has the premise that ‘if a room is large and unoccupied’ then the conclusion is that it is ‘suitable for a manager’. At the bottom right of Figure 3, an individual ‘Room 119’, represented in the visual language as a rectangle, is asserted to be ‘large’ and ‘unoccupied’, represented by arrows from the individual to these concepts. When the entire knowledge structure of concept definitions, rules and assertions, is then compiled and run through the inference engine, the graph

output is that shown at the bottom right of Figure 3. Room 119 has been inferred to be suitable for a manager.

This anticipatory behaviour of rules is not so readily represented in the geometry because it is a dynamic principle, ‘if you place an element in this intersection then move it to this sub-intersection’. A rule corresponds to an *attractor* in dynamics, that an element placed in a region of space is unstable and falls into the basin of an attractor (Abraham & Shaw, 1984). This can be represented in the geometry, but the diagrams for activities of any reasonable complexity become very difficult to visualize and understand. It does, however, provide a nice link to the corresponding phenomena in the geometrical dynamics of neural networks underlying the phenomenon (Domany, Hemmen & Schulten, 1991).

The logic based on Kelly’s axiomatic presentation of personal construct psychology, and the visual language representing it, both extend to support the additional features normal in term subsumption knowledge representation systems, such as attributes and relations, or ‘roles’ as they have been termed generically (Brachman & Schmolze, 1985), rules with exceptions (Gaines, 1991b), and contexts (Sowa, 1984). Figures 2 and 3 have been presented in a graphing tool KDraw (described in section 11), that provides a fully operational semantics for the input and output of knowledge structures in the visual language, and further illustrations of its application are given in later sections.

5 The repertory grid

Kelly introduces the ‘role repertory grid’ (Kelly, 1955) as a means for investigating a person’s conceptual structure relevant to inter-personal relations by having that person classify a set of people significant to them in terms of elicited personal constructs. Figure 4 shows the general form of a repertory grid and its relation to the conceptual structures already discussed. If one takes a particular concept somewhere in the lattice, and a set of individuals asserted to fall under that concept, then each distinction that may be made about individuals falling under that concept forms the rows of a matrix, the individuals form the columns, and the constraints applying to a particular individual relative to a particular distinction form the values in the matrix.

In simple applications of the repertory grid, these constraints are taken to be the values of the individuals on the roles corresponding to the distinctions. However, it is apparent from Figure 4




Figure 4 The repertory grid as a matrix of concepts, individuals and constraints




Figure 5 A repertory grid about spatial mapping techniques

that concepts subordinate to those defining the scope of the grid may also be used as if they were individuals, and these may be expected to have more general constraints than single values. Hence in extended repertory grid elicitation, such as that of AQUINAS (Boose & Bradshaw, 1987) the 'values' in the matrix can themselves be complex constraints.

Figure 5 shows a basic repertory grid elicited from a geographer about spatial mapping techniques. The mapping techniques used as elements are listed as column names at the bottom. The poles of the constructs elicited are listed on the left and the right as row names. The ratings of the mapping techniques along the dimensions of the constructs form the body of the grid. A 1–9 scale has been used based on a construction similar to that of Figure 2. For example, 'probability mapping' is rated 8 on the dimension 'qualitative and quantitative—quantitative' which means that it is construed as primarily 'quantitative'. Note that the predicates defined by the pole names may through their linguistic labelling indicate that they themselves are compounds of more primitive predicates. The grid is the starting point for analysis and refinement, not necessarily in itself a conceptual structure, but rather a set of data that must be accounted for by any conceptual structure developed. The grid of Figure 5 is a typical one based on Kelly's original specification. Later developments allow structural relations between elements and between constructs to be specified, and the ratings to be extended to more complex constraints.

The psychological function of the repertory grid is to provide a technique for building the conceptual structure without direct elicitation of concepts and their structures and relationships. The assumption is that it may be easier for a person to provide exemplary individuals in the domain of interest, and then to state in fairly concrete terms how they would distinguish them in terms of properties relevant to the purpose of eliciting the grid. In terms of the intensional logic of the

concept structure, the extensional specification of how concepts apply to individuals is clearly inadequate to fully specify the concept structure. However, the structure must be consistent with its model, and hence it is possible, through suitable analysis techniques, to approximate the structure from the extensional data, as is discussed in the next section.

5.1 Using repertory grids

The use of the repertory grid to elicit concept structures thus involves a variety of psychological and analytical techniques, including:

1. Careful definition of the purpose of the elicitation and the appropriate sub-domain to be considered. Maintaining this context so that the purpose and domain do not tacitly change during elicitation is also very important.
2. Choice of exemplary individuals that characterize the relevant features of a domain. This choice is very important, and is a major focus of attention in both tool design and application. Fortunately, experts often find it natural to discuss a domain in terms of stereotypical cases, but much care is required to elicit a full range of stereotypes adequate to characterize a domain.
3. Various techniques may be used for initial element elicitation including interviews, protocol analysis, brainstorming with groups of experts, and keyword extraction from relevant textual material (Shaw & Gaines, 1987; Shaw & Woodward, 1990).
4. Online analysis of the interim conceptual structures may be used to detect closely related distinctions, and use these relations to request information on potential stereotypes that might specifically reduce the closeness of the distinctions (Shaw, 1980).
5. The elicitation of some initial distinctions may again derive from interviews, protocols, brainstorming and text analysis.
6. *When no prior information is available, triadic elicitation can be effective, in which a randomly selected set of three individuals is presented with a request to state in what way are two alike and differ from the third.*
7. Online analysis of the interim conceptual structures may be used to detect closely related individuals, and use these relations to request information on potential distinctions that might specifically reduce the closeness of the individuals (Shaw, 1980).
8. The conceptual structure can be developed through various forms of hierarchical and spatial cluster analysis such as FOCUS (Shaw, 1980) and principal components analysis (Slater, 1976; 1977).
9. Rule induction may be used both to derive potential implications between concepts and also, since the premise of a rule is itself a concept, to develop non-primitive concepts and their subsumption relations (Gaines & Shaw, 1992).
10. Direct elicitation of the concept structure may be mixed with indirect development of the grid (Boose & Bradshaw, 1987).

In clinical psychology, repertory grids have been used extensively without computer support, and techniques for manual elicitation are well-documented (Fransella & Bannister, 1977). Repertory grid methodologies are difficult to undertake manually as they require feedback and management from the elicitor while at the same time attempting to avoid inter-personal interactions that would distort the elicitée's conceptual structures. Hence the advent of the personal computer in the mid-1970s and its evolution into the graphic workstations of the 1980s has been very significant for the practical application of the approach (Shaw, 1980, 1981; Mancuso & Shaw, 1988).

6 Distance measures, conceptual clustering and induction

In analysing repertory grid data, distance measures play an important role in conceptual clustering and induction. This section shows how such measures may be derived from the logical structures already defined by considering the extensions of the concepts.




Figure 6 Calculation of distance measures between concepts and between constructs

6.1 Conceptual clustering

In terms of the logic and visual language, there is a natural construction of a distance between two concepts, *x* and *y*, as shown on the left of Figure 6. Let *u* be some minimal upper bound of *x* and *y* subsuming both of them, and *I* some maximal lower bound subsumed by both of them, and *U* be the extension of *u* and *L* the extension of *I* over some universe of individuals. If *x* and *y* are identical so will be *U* and *L*, whereas if they are disjoint *L* will be empty. Hence a natural distance measure is the number of individuals that are in *U* but not *L*:

'*x* distance *y*'

$$d(x, y) = CU - CL$$

(4)

where *CU* and *CL* are the cardinalities of *U* and *L*, respectively. This measure satisfies the triangle inequality, and can be normalized by dividing by its maximum possible value, *CU*. It is clearly dependent on the universe of individuals involved, but this is appropriate to measuring concept distance in an extensional context. Intensional concept 'distance' independent of context is reflected in the relational structures already developed.

The distance measure defined readily extends to dichotomous constructs through the comparison of poles as shown on the right of Figure 6:

'*b*—*c* distance *d*—*e*'

$$d(b-c, d-e) = CA - CF - CG$$

(5)

This measure is a count of the numbers of individuals that fall under the opposite pole of the other construct. Note that it is not invariant if one construct is reversed. This construction generalizes to scales with more than three points. If these scales are numbered linearly it computes a 'city block' distance measure—which is precisely that used in construct clustering algorithms such as FOCUS (Shaw, 1980). These distance measures, when applied to the constructs of one person, enable natural clusters to be seen that may be grouped as part of a coherent concept, for example, in that they are all contributors to an evaluatory dimension. When applied to the constructs of different people in the same domain, they enable consensus, conflict, correspondence and contrast in the use of constructs to be measured (Shaw & Gaines, 1989), as discussed in the next section.

Figure 7 shows a FOCUS analysis of the grid of Figure 5 in which the distance measure defined in (5) has been used to develop two matrices of inter-element and inter-construct distances. The sets of elements and constructs have then each been sorted to re-order the grid in such a way that similar elements and similar constructs are close together. The hierarchical clusters on the right portray this similarity by developing groups of elements and constructs, starting with the most similar and adding the next most similar, and so on, until all the elements and constructs are clustered.

Thus, at the bottom of the construct clusters, it can be seen that the dimension 'incorporates geologic model—doesn't incorporate geologic model' is used very similarly to 'qualitative and quantitative—quantitative', and that both of these relate closely to 'interpretive—representative'. This geographer construes mappings that incorporate a geologic model as qualitative and quantitative, as opposed to solely quantitative, and as interpretive as opposed to representative.

FOCUS: Expert2

Elements: 12, Constructs: 18, Range: 1 to 9, Context: To evaluate spatial interpolation techniques




Figure 7 FOCUS hierarchical clustering of spatial mapping grid

Similarly, in the element clusters, 'nonparametric kriging', 'punctual kriging' and 'universal kriging' are construed as closely related techniques with little distinction between them.

6.2 Rule induction

The measures used in the induction of a rule linking to concepts are also readily derived as shown in Figure 8. CX is the number of 'predictions' made by concept x as the left hand side of a rule, and CL is the number which are correct. Thus, the measures of the validity of inducing the rule $x \rightarrow y$ are:

$$\text{'prior probability of } y\text{' } p(y) = CY/CU \quad (6)$$

$$\text{'probability correct } x \rightarrow y\text{' } p(x \rightarrow y) = CL/CX \quad (7)$$

$$\text{'probability by chance } x \rightarrow y\text{' } c(x \rightarrow y) = Ip(y)(CX-CL, CL+1) \quad (8)$$

where I is the incomplete beta function summing a binomial distribution tail. These measures are precisely those used by Induct (Gaines, 1989) in inducing rules from datasets. In the application to repertory grids Induct searches for potential rules whereby a target predicate may be deduced from some of the others, and constrains the search to rules whereby the probability that they arise by chance is less than some prescribed threshold. The basic search techniques have been well documented by Cendrowska (1987), but for practical applications they need to be controlled by these probabilistic measures, and also to be extended to generate rules with exceptions as these are both more compact and more in accordance with human practice (Gaines, 1991b).




Figure 8 Induction of rules between concepts




Figure 9 Grid based on Cendrowska's contact lens data

RULE 1: -> lens = none
RULE 2: age = young & astigmatism = astigmatic & tear production = normal -> lens = hard (EXCEPTION TO: 1)
RULE 3: prescription = myope & astigmatism = astigmatic & tear production = normal -> lens = hard (EXCEPTION TO: 1)
RULE 4: astigmatism = not astigmatic & tear production = normal -> lens = soft (EXCEPTION TO: 1)
RULE 5: astigmatism = not astigmatic & age = presbyopic & prescription = myope -> lens = none (EXCEPTION TO: 4)

Figure 10 Induct analysis of contact lens data




Figure 11 Contact lens domain and rules represented in the visual language

To illustrate rule induction from repertory grids, Figure 9 shows a grid for Cendrowska's contact lens example (Cendrowska, 1987) involving five constructs and 14 stereotypical elements. The objective is to derive the prescription as 'hard', 'soft' or 'none' (this last having been specified as a rating value of 2 on the 'hard-soft' dimension). Figure 10 shows the five rules derived by Induct from the contact lens data. The first rule is a default to be applied when none of the others apply, recommending 'none'. The next two rules specify the conditions for 'hard'. The fourth rule specifies the conditions for 'soft', and the fifth specifies an exception to this rule.

Figure 11 shows the construct data in the grid used to define a class of patients together with the rules represented in the graphical notation described in this paper in the grapher implementing the visual language. The only extension to what has been already described is that an arrow from one rule to another means that the first is an exception to the second. The additional terminology has been introduced by having the user specify names for the constructs in the grid, for the central value of the 'young-presbyopic' dimension, and for the rules. The graph produced by automatic output from the analysis of grid and rules has been manually edited to give a pleasing appearance but otherwise is a direct analysis of the grid.

7 Eliciting conceptual structures from groups

In developing knowledge-based systems, one becomes aware very rapidly that the required conceptual structures can only rarely be elicited through access to one expert. The basic cognitive




Figure 12 Four-quadrant representation of consensus, correspondence, conflict and contrast between distinctions in conceptual systems

agent is not the person but the group. Kelly took group processes into account in personal construct psychology, and the psychology and logic described extend readily to the group situation. The primary additional consideration in modelling the knowledge structures of groups is a linguistic one. We have so far assumed that the terminology used for distinctions is well-defined and unambiguous. This is not reasonable in a group situation. People may use the same term for different distinctions, and different terms for the same distinction. Elicitation methodologies for the group situation are much the same as those for the individual situation, except for the need to track such terminological differences. Figure 12 shows the four situations that may arise through interactions between terminology and distinctions (Shaw & Gaines, 1989).

The recognition of *consensual* concepts is important because it establishes a basis for communication using shared concepts and terminologies. The recognition of *conflicting* concepts establishes a basis for avoiding confusion over the labelling of differing concepts with the same term. The recognition of *corresponding* concepts establishes a basis for mutual understanding of differing terms through the availability of common concepts. The recognition of *contrasting* concepts establishes that there are aspects of the differing knowledge about which communication and understanding may be very difficult, even though this should not lead to confusion. Such contrasts are more common than is generally realized. For example, it is possible to derive the same theorem in mathematics either by using an algebraic perspective, or a geometric one. There is nothing in common in these two approaches except the final result. It may still be possible to discuss the same domain using consensual and corresponding concepts that were not fundamental to the problem solving activities.

Figure 13 shows grid data elicited from two geographers on spatial mapping techniques analysed in terms of consensus, correspondence, conflict and contrast. Some raw data and its individual analysis has already been shown in Figures 5 and 7. There is no significant example of *contrast* in this data, possibly because the two experts work very closely together. The basis of the consensus and conflict analysis is shown in detail in the difference analysis of Figure 14, where rating values (in this case 1 to 5) for expert B's ratings of expert A's elements on his constructs are subtracted from expert A's similar rating values respectively. Figure 14 shows this with the elements and constructs about which they agree the most in the top right corner, shown by no difference or a difference of only 1; and those with most disagreement towards the bottom left, shown by the maximum difference of 4 or a large difference of 3. The darker grey areas indicate the greater difference. The graphs on the right show the declining matches, and may be used to decide where to place thresholds distinguishing consensus and conflict. From the difference grid, the consensus and conflicts can easily be identified and discussed by the experts.




Figure 13 Four-quadrant analysis of grid data from two experts in spatial mapping

The basis of the correspondence and contrast analysis is shown in detail in the construct comparison analysis of Figure 15, in which construct matches are sorted with highest first. The cumulative percentage is given of those with matches greater than the value shown, and the construct from one expert which best matches each from the other is shown beneath it. It can be seen that:

- Both experts are using the construct *local-global* in the same way. This is an example of *correspondence*. It is not an interesting one because we can see that the terms used are the same and it is effectively arising from a *consensus*.
- When one expert uses the term *low-level-data-high-level data*, the other is using the term *nominal data–interval or ratio data*. This shows a difference in terminology which can be




Figure 14 Difference analysis for two spatial mapping experts

G1<:G2 62.5% over 80.0 (ExpertA construct-construed-by ExpertB)		
1:	6.2% ≥ 88.5	G1C2: Local - Global G2C2: local - global
2:	12.5% ≥ 87.5	G1C3: Low level data - High level data G2C8: nominal data - interval or ratio data
3:	86.5	G1C1: Does not honour data points - Honours data points G2C4: doesn't honour data points - honours data points
4:	86.5	G1C7: Short distance autocorrelation - Long distance autocorrelation G2C2: local - global
5:	31.2% ≥ 86.5	G1C9: New geographical technique - Old geographical technique G2C18: not widely used - widely used
6:	37.5% ≥ 85.4	G1C16: Not widely used - Widely used G2C18: not widely used - widely used
7:	43.8% ≥ 83.3	G1C5: Discontinuous - Continuous G2C2: local - global
8:	50.0% ≥ 82.3	G1C4: Mathematically complex - Mathematically simple G2C11: heavy computing load - no computing load
9:	81.2	G1C10: Hard to adapt to multivariate - Easy to adapt to multivariate G2C5: usually one variable considered - multiple variables considered
10:	62.5% ≥ 81.2	G1C12: Does not require spatial search - Requires spatial search G2C13: estimates susceptible to clusters - not as susceptible to clusters

Figure 15 Construct comparison for two spatial mapping experts

interpreted as their *levels of abstraction* being different in the construing of this topic, and is an interesting *correspondence*.

- The construct *heavy computing load–no computing load* is being used by one expert to correspond to *mathematically complex–mathematically simple* used by the other. We can interpret this as a difference in terminology corresponding to a correlation in the real-world.

8 Tools based on personal construct psychology—KSS0

The preceding sections have given the historic, psychological and technical foundations of personal construct psychology, and illustrated some of the techniques that are used in its application to knowledge acquisition. In doing this, some aspects of repertory grid and semantic net elicitation and analysis tools based on personal construct psychology have also been illustrated. There is not space in this paper to present even the major tools in great detail—the tools have many features, the user manuals are large, and each application in itself requires a long paper to describe. The details are available in the knowledge acquisition literature and in many papers in journals and conference proceedings that have already been cited. This and the following sections give an overview of the objectives, activities and developmental directions of the two major research groups that have developed a diversity of knowledge acquisition tools: that associated with the authors at the Centre for Person-Computer Studies and the Knowledge Science Institute, and that associated with John Boose and Jeffrey Bradshaw at Boeing Computer Services. A brief history of developments in the first group is given in Gaines (1988) and of the second in Boose et al. (1989).

KSS0 (CPCS, 1991) is a repertory grid elicitation and analysis system running on Macintosh computers that aims to allow end-users, the experts rather than the knowledge engineers, to develop their conceptual structures through direct computer interaction. Hence it focuses on ease of use through attractive user interfaces to both elicitation and analysis. Figure 16 gives an overview of the tools in KSS0 and their interrelations.

The user interface to KSS0 is designed for use by those who are not computer specialists. This shows up particularly in the knowledge elicitation tools that are highly graphical and emphasize knowledge visualization at all stages of elicitation:

- *Elicit* accepts specifications of cases as elements within a sub-domain, and provides an interactive graphical elicitation environment within which the experts can distinguish elements to derive their constructs. The resultant class is continuously analyzed to provide feedback prompting the expert to enter further elements and constructs.

The visualization tools consist of an interactive interface to represent the abstractions derived from those elements in terms of hierarchical clusters using FOCUS, and relational diagrams such as non-hierarchical conceptual maps derived through principal components analysis. The objectives are to validate the raw domain knowledge and to suggest further structure:




Figure 16 KSS0 architecture

- *FOCUS* hierarchically clusters elements and constructs within a sub-domain prompting the experts to add higher-level constructs structuring the domain.
- *PrinCom* spatially clusters elements and constructs within a sub-domain prompting the experts to add higher-level constructs structuring the domain.

The group comparison tools consist of an interactive interface to represent the relations between the terminologies and conceptual systems of different experts, or experts and clients:

- *Socio* compares the structures for the same sub-domain generated by different experts, or the same expert at different times or from varying perspectives.

The inductive part consists in the derivation of constraints within the conceptual structures through logical entailment analyses. The objective is to suggest further structure at a higher level that translates into concept subsumptions or rules in the expert system shell:

- *Induct* induces logical entailments enabling the constructs applying to an element or the evaluations of a decision-making situation in a domain, to be derived from other constructs.

The generative part consists in the transformation of the knowledge analysis made by the previous tools into formalisms understandable by knowledge-based system shells such as *NEXPERT* (Rappaport, 1987), *Babylon* (Christaller et al., 1992) and *KRS* (Gaines, 1991a):

- *Export* formats the specifications of sub-domains as classes, of elements as objects, of constructs as properties, and of entailments as methods, and transfers them to the performance tool.

8.1 Examples of elicitation and matching

The main aspects of KSS0 that have not been illustrated so far in this paper are the *Elicit* module for the interactive graphical elicitation of repertory grids, and the *Export* module for transfer of the




Figure 17 Construct elicitation from a triad

results of grid elicitation and analysis to an expert system shell. Figure 17 shows triadic elicitation occurring during the elicitation of a grid on spatial mapping techniques, and Figure 18 shows the elements being rated on the construct elicited in this way. Triadic elicitation is used to elicit the first few constructs, but as more data is obtained other methods become possible based on a continuous analysis of the knowledge structure. Figure 19 shows a construct match being used to suggest the entry of another element which might break the match. KSS0 generates a wide variety of such




Figure 18 Click and drag rating of elements on constructs




Figure 19 Eliciting an element through a construct match

```
(@CLASS= Rules (@SUBCLASSES= contact_lens_prescription)
  (@PROPERTIES=
    hypothesis @TYPE=Boolean;
  )
)

(@CLASS= contact_lens_prescription
  (@SUBCLASSES=
    lens_recommendation
  )
)

(@CLASS= lens_recommendation)
```

Figure 20 Rule class and subclass definitions from contact lens dataset

screens during elicitation to maintain the expert's interest and to explore his or her psychological space. Its operation is non-modal, so that cluster and induction analyses may be requested during elicitation also.

8.2 Examples of induction and export to an expert system shell

Figures 20 through 23 are based on the contact lens data of Figures 9 through 11, and show some of the data structures exported from KSS0 to NEXPERT. These are given in detail because experience has shown that the way in which repertory grid data translates into expert system knowledge structures requires considerable attention when implementing knowledge acquisition tools. The structures for another shell such as BABYLON are similar in form but very different in detail.

Figure 20 shows rule class and subclass definitions used to establish an object-oriented framework for the rules.

```

(@OBJECT=    Global_Hypothesis)
(@RULE=      Global
  (@LHS=
    (Yes    (<Rules>.hypothesis))
  )
  (@HYPO=    Global_Hypothesis)
)
(@GLOBALS=   @SUGLIST=    Global_Hypothesis;)

```

Figure 21 The top level control rule

```

(@CLASS=     people
  (@PROPERTIES=
    age       @TYPE=String;
    prescription @TYPE=String;
    astigmatism @TYPE=String;
    tear_production @TYPE=String;
    lens_recommendation @TYPE=String;
  )
)

(@SLOT=      people.age    @PROMPT="Double click on the age from the list
below that applies to @SELF";)

(@OBJECT=    person (@CLASSES=  people))

```

Figure 22 Case class definition, prompting and instantiation

```

(@OBJECT=    Ind_1  (@CLASSES=  lens_recommendation))
(@RULE=      Ind_1
  (@LHS=
    (IsNot (<people>.lens_recommendation) (KNOWN))
    (Is    (<people>.prescription) ("hypermetrope"))
    (Is    (<people>.astigmatism) ("not_astigmatic"))
    (Is    (<people>.tear_production) ("normal"))
  )
  (@HYPO=    Ind_1.hypothesis)
  (@RHS=
    (Let    (<people>.lens_recommendation) ("soft"))
  )
)
(@SLOT=      Ind_1.hypothesis    @INFCAT=    10098;)

```

Figure 23 Induced rule transferred to *NEXPERT*

Figure 21 shows the top level control rule generated. It uses the class of Rules defined in Figure 20 in a pattern-matching clause that tests whether the hypothesis of any rule is true. The global hypothesis of this rule is also put on the list of suggestions in the shell so that this rule may be triggered very simply.

Figure 22 shows the class definition that transfers the primary knowledge structure about the class of cases defined by the dataset. It is followed by meta-slot definitions, of which one is shown, giving the prompts that the shell should use in requesting the values of attributes. These are followed by the instantiation of one object in the class that may be used as a test case.

Figure 23 shows a rule generated by *Induct* translated into *NEXPERT* knowledge base format. The first line instantiates a rule object for the rule. The left hand side tests the premise of the rule

with an added test to determine whether the value of the attribute to be set on the right hand side of the rule is already known.

9 Tools based on personal construct psychology—AQUINAS

AQUINAS and KSS0 have very similar origins and some related functionality, since the two groups of developers have collaborated through very open information interchange. However, there are major differences in the directions of development and objectives that make for very significant variations in the operations and functionalities of the two systems. The major difference in performance objectives is that while ETS, the precursor to AQUINAS, was aimed at knowledge acquisition for expert system development, AQUINAS has been designed as a complete decision support environment. It uses case-based reasoning to provide its own inference engine independent of any expert system shell. The major difference in technical objectives is that AQUINAS, as already noted, greatly extended the basic repertory grid to support hierarchical structures of elements, constructs, multiple experts and multiple tasks, and to allow complex constraints to be specified rather than just rating values.

Figure 24 shows the way in which a rating grid in AQUINAS is essentially a four-dimensional structure indexed by hierarchies of solutions (elements), traits (constructs), experts (elicitees) and cases (application contexts). The particular grid under consideration at any time is determined by the sub-hierarchies selected as indicated in the diagram.

Figure 25 is a concrete instance of these hierarchies developed for a programming language advisor, and Figure 26 shows how a simple grid might be defined through selection in the hierarchies. AQUINAS provides a wide range of user facilities to navigate the multi-dimensional grid structure. In its inference procedures it also provides techniques to use grid data from multiple sources and levels to provide advice in decision support problems. For example, if one expert has not rated a solution on a particular trait, then this information may be obtained from the data of another expert. More complexly, ratings available only at higher levels in the hierarchy may be propagated to lower levels and, conversely, ratings available at only lower levels may be used to estimate those at higher levels. This need to estimate expert judgments from a variety of sources, and to use them directly to support decision has let to the use of a variety of existing decision-theoretic methodologies in AQUINAS, and also to the development of many novel approaches.

As Boose and Bradshaw (1987) note, in AQUINAS various trait (attribute) scale types can be elicited, analyzed, and used by the reasoning engine. Traits are currently described according to the level of measurement of their rating scales, which is determined by the expert. The level of measurement depends on the presence or absence of four characteristics: distinctiveness, ordering




Figure 24 AQUINAS extended repertory grid structure (Boose & Bradshaw, 1987)




Figure 25 Programming language advisor hierarchies in AQUINAS (Boose & Bradshaw, 1987)

in magnitude, equal intervals, and absolute zero (Coombs, Dawes & Tversky, 1970). These four characteristics describe the four major levels of measurement, or types of traits: nominal (unordered symbols), ordinal, interval, and ratio. The additional information about trait types gives increased power to analytical tools within AQUINAS and allows experts to represent information at the level of precision they desire.




Figure 26 Defining a grid from a hierarchy in AQUINAS (Boose & Bradshaw, 1987)

- Ratings may be generated through several methods:
- 1. Direct** An expert directly assigns a rating value for a trait and an element. If an exact value is unknown, AQUINAS helps the expert derive an estimate (Beyth-Marom & Dekel, 1985). If fine judgments are needed, AQUINAS can derive a set of ratio scaled ratings from a series of pairwise comparisons (Saaty, 1985). AQUINAS also contains tools for encoding of probability distributions on specific values. The value with the highest probability is displayed in the grid, but all appropriate values are used in reasoning and may be edited with graphic distribution aids.
 - 2. Derived** Incomplete grids can be automatically filled through propagation of rating values from another grid through the hierarchies (e.g., from lower to higher level grids, different experts, or different cases).

One particularly attractive feature of AQUINAS is the way in which the expert hierarchy may be used to provide advice based on the opinion of a particular expert or weighted combinations of experts, and also used to provide an analysis of the extent to which other experts might dissent from this advice (Boose, 1986b). The dissenting opinion is found by computing a correlation score between each member and the consensus; the member with the lowest correlation score is listed as the dissenting opinion. Dissenting opinions show the users the range of opinion about a decision, not just the top rated list, and give decision makers confidence that the top rated alternatives are sound choices or point out areas of disagreement for further exploration.

The Boeing group has developed many other decision support tools that have been combined in a variety of ways with AQUINAS. For example, Canard is an engineering design aid that helps synthesize alternatives from potentially large search spaces (Bradshaw et al., 1990). Canard helps generate and structure complex alternatives in a possibility table. The possibility table representation was adopted from manually developed strategy tables (McNamee & Celona, 1987) and morphological charts (Zwicky, 1969). Canard automates this representation, and extends its logic and structure to allow knowledge-based inference and the representation of more complex problems (e.g., hierarchical tables, explicit representation of constraints). Figure 27 shows a Canard possibility table for a portion of a network design problem. The thick horizontal line is a partial solution path linking component alternatives. Shaded cells show hard and soft constraints associated with the path. The left-most column shows classes of solutions associated with paths.




Figure 27 Engineering design in Canard (Shema, 1990)

Repertory grids 'plug in' to each column, recording criteria relationships and enabling analysis. Repertory grid reasoning methods help rank the components in a column. Automated methods combine ranking information with constraints to produce a set of best possible solutions.

DART (Design Alternatives Rationale Tradeoffs) is an application of repertory-grid-based knowledge acquisition techniques to *trade studies* in engineering design. It was originally developed for NASA as part of an effort to capture design knowledge for the Space Station Freedom program (Boose et al., 1990). In recent years, a major focus of activity for the group has become the development of *group decision support systems* based on extended repertory grids and decision analysis (Boose et al., 1992).

10 Tools based on personal construct psychology—KSSn

KSS0 and AQUINAS were both developed commencing with the standard, rating-scale based repertory grid and extending it in various ways. In the light of several years experience with the systems, and with developments in the logical expression of personal construct psychology, it became timely in the 1990s to develop a new range of systems based on the cognitive and logical foundations described in this paper. In particular, the weaknesses of the basic repertory grid in representing more 'objective' data such as categories, numbers, dates, and so on, have restricted those methodologies without major extensions to use only in the initial stages of knowledge acquisition. Similar considerations apply to the lack of representation of hierarchical relations between elements and constructs in the basic repertory grid. As knowledge acquisition tools mature it is reasonable to expect them to extend in scope from initial elicitation, through detailed knowledge modelling, to validation and knowledge base maintenance. This has been achieved experimentally through the various extensions already described, and also by integrations with other systems, as described in the next section. However, it is important that the somewhat *ad hoc* systems developed in this way be used as a basis for a new generation of more principled designs.

KSSn is an ongoing experiment in the development of knowledge acquisition tools based on personal construct psychology. It extends KSS0 by basing all the elicitation and analysis tools on a common knowledge representation server supporting a wide range of data structures, including the hierarchical conceptual relations of KL-ONE-like systems. KSSn is implemented in C++ (Gaines, 1991a), and the representation component provides services similar to those of CLASSIC (Borgida et al., 1989) 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 (Gaines, 1991b). The server supports the operations of the intensional logic underlying personal construct psychology as described in this paper, and one of the modules attached to it is a graphic knowledge editor, KDraw, supporting the associated visual language (Gaines, 1991c). KDraw was used to represent the knowledge structures shown in Figures 2, 3 and 11.

The user interfaces to the server in KSSn are designed to extend those of KSS0 to a wide range of data types and to hierarchical knowledge structures, without losing the essential simplicity of the original interfaces which have proved very effective with end users. Visualization interfaces such as those of Figures 17 through 19 are still used initially to capture basic repertory grid data, but the scales can be converted to categories, numbers or dates as the knowledge structures are refined. The elements and constructs are also visually represented not as lists but rather as a knowledge structure in the grapher. Figure 28 shows the grid of Figure 5 during development in KDraw. This is a simple flat structure of individuals instantiating roles, but the grapher allows this basic grid structure to be extended to hierarchical relations. Figure 29 shows the AQUINAS knowledge structure of Figure 25 in the grapher.

While KSSn still 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, Figure 30 shows KSSn in use on a room allocation problem (Gaines, 1991d) derived from an ESPRIT project (Voß, 1990) that has recently been placed in the public domain as part of Project Sisypus. Sisypus is research program to encourage international




Figure 28 Spatial mapping grid elements and constructs in the grapher




Figure 29 AQUINAS language advisor solution and trait hierarchies in the grapher




Figure 30 Room allocation task solution in the visual language

collaboration in knowledge-based system development initiated by the European Knowledge Acquisition Workshop (EKAW) in 1989. A number of problem datasets have been made available through Sisyphus, and a major part of the EKAW'91 program was devoted to reports on the solution of these problems using different approaches and techniques (Linster, 1991).

At the top, Figure 30 shows the top level conceptual structures of an organizational domain, and at the bottom, the assertion of facts about a particular organization. Human-computer interaction in the editor is modelled on Apple's MacDraw with additional features appropriate to the language, such as arcs remaining attached to nodes when they are dragged. A popup menu as shown at the bottom of Figure 30 allows connecting lines to be entered easily. 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 popup 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. However, partitioning data structures over several screens is encouraged, and has proved practical in managing large knowledge structures.

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 analysed through the KSS0-style modules, may be exported to the grapher for visual analysis and editing. Thus KSSn is a step towards a new family of tools that instantiate more aspects of personal construct psychology than those based on repertory grids alone.

11 Integration with related systems

The development of Canard as an extension to AQUINAS is an indication of the possibilities of using personal construct psychology-based knowledge acquisition tools as embedded components in larger systems. A particularly attractive combination is to integrate hypermedia tools for informal knowledge acquisition with repertory grid knowledge acquisition tools and an expert system shell. For example, KSS0 has been extended to support the informal representations of knowledge that are prior to those within the tool, such as text, pictures, diagrams, semi-structured interviews, protocols, and so on. This has been done by providing an interapplication protocol allowing KSS0 to interact with Apple's HyperCard to provide the appearance of a seamless single application to users. Each element and construct has a card associated with it on which can appear a detailed description, diagrams, or whatever other form of annotation the expert wants to include. KSS0-specific functionality in HyperCard is supported by scripts that allow conceptual structures on the KSS0 side to be linked to informal sources and annotation on the HyperCard side. The same inter-application protocol has also been used to provide run-time integration with the knowledge-based system shells to which KSS0 exports, such as NEXPERT (Gaines, Rappaport & Shaw, 1989) and BABYLON (Gaines & Linster, 1990).

Figure 31 shows the distributed knowledge base and inter-application protocols linking the hypermedia, knowledge acquisition and knowledge-based system shell, BABYLON, in such an integrated system. Also shown are the types of interaction which are supported with experts, clients and knowledge engineers. It is important to note that all three types of user see a single,




Figure 31 Integration of hypermedia, knowledge acquisition and expert system shell




Figure 32 Sequence of activities in using integrated system for knowledge base development

uniform interface, so it is easy for people to change roles. The ‘expert’ can become a ‘client’ for testing, or a ‘knowledge engineer’ for deeper technical development.

The operation of each component of this system is similar to its stand alone use, but the inter-application protocol provides a new level of functionality. For example, in the KSS0 analysis and elicitation screens already shown, popup menus become available as one moves the mouse cursor over element and construct pole names. These give access to the HyperCard cards that can be used to annotate these elements and constructs. These same cards may also be used as a decision-support interface to interrogate the shell, and as annotation accessible from the shell.

One important result of on-line integration between the knowledge acquisition and performance tools is that new modes of knowledge acquisition become possible. Figure 32 shows a typical development cycle with the integrated system. The most interesting new mode is that, when the knowledge base exported from KSS0 is run in BABYLON, it can be tested on new cases and, if these lead to erroneous decisions, the cases may be corrected and posted to KSS0 to extend the grid and hence affect the rule induction and knowledge base generated. This capability allows knowledge acquisition to be integrated with the validation of the knowledge base in the applications environment.

Ideally, the loop between BABYLON, KSS0 and HyperCard could be left in place indefinitely for system upgrade and maintenance. However, since the knowledge acquisition tools do not yet support the full functionality of the knowledge base, particularly its capabilities for procedural attachment and linkage to data-processing subsystems, it is usual for the final system development to take place in BABYLON as shown at the lower right of Figure 32. Note that the links to the HyperCard annotation are retained in the final performance system.

The integration of knowledge acquisition tools based on personal construct psychology with other related tools such as hypermedia and performance systems offers many important possibilities for future research, development and application. Ford, Stahl, Adams–Webber, Novak and Jones have suggested integration with neural nets as a grid generating technique, and with conceptual maps as a conceptual structure representation (Ford et al., 1990). RepGrid-Net combines repertory grids and electronic mail to allow geographically dispersed groups to work together and exchange and mutually develop knowledge structures (Shaw & Gaines, 1991a,b), AQUINAS incorporates its own performance tool, and has been used as part of other systems such

as Canard through heterogeneous integration. KSSn has an open architecture supporting various forms of integration at different levels, through the C++ class library, through direct calls to the server from associated modules, and through indirect calls, possibly across a network, from heterogeneously integrated programs.

The following section outlines the types of knowledge acquisition methodologies appropriate to such integrated systems.

12 Integration of knowledge acquisition methodologies

The foundations for personal construct psychology and knowledge representation developed in this paper also provide a framework for the analysis of knowledge acquisition techniques and the design of integrated tools. To use these tools effectively requires the development of knowledge acquisition methodologies that themselves integrate a variety of approaches. Figure 33 illustrates the way in which the techniques described in this paper may be seen as supporting an overall development methodology:

- Stage 0 consists of the acquisition of informal knowledge from interviews, protocols and media.
- Stage 1 consists of using the informal knowledge to elicit the major coherent sub-domains that together encompass the significant phenomena in the overall domain.
- Stage 2 consists of using the repertory grid methodology in each sub-domain to elicit relevant attributes and critical cases.
- Stage 3 consists of using the conceptual induction methodology to derive concepts and rules from the grids.




Figure 33 Integration of knowledge acquisition techniques

Stage 4 consists of linking the sub-domains together, generally by specifying as a constraint in one domain that some role in that domain has as value an individual in another.

Stage 5 consists of testing the overall knowledge base and iterating back through any of the earlier stages in order to develop and refine the knowledge base further.

The architecture of Figure 33 is not intended to be a rigid framework for integrated tools, but rather to illustrate ways in which the various knowledge acquisition techniques already described may be naturally combined. For example, some sub-domain structures may be developed by direct editing, others by induction from databases, others by text analysis, and others by tools not yet defined or developed. The techniques and tools we have described come together naturally in this framework, and it also provides an open architecture for incorporating further techniques and tools. This paper gives strong theoretical foundations for what is shown. In the longer term, the architecture shown may need major extension, but currently it seems to encompass what is available to us in knowledge acquisition.

13 Conclusions

Personal construct psychology is a theory of individual and group psychological and social processes that has been used extensively in knowledge acquisition research to model the cognitive processes of human experts. The psychology has the advantage of taking a constructivist position appropriate to the modelling of specialist human knowledge, but basing this on a positivist scientific position that characterizes human conceptual structures in axiomatic terms that translate directly to computational form.

The repertory grid knowledge elicitation methodology is directly derived from personal construct psychology. In its original form, this methodology was based primarily on the notion of dichotomous constructs and did not encompass the ordinal relations between them captured in semantic net elicitation. However, it has been extended in successive tools developed for applied knowledge acquisition and tested in a wide variety of applications.

This paper has given an overview of personal construct psychology and its expression as an intensional logic describing the cognitive processes of anticipatory agents. A theoretical framework has been developed and shown to provide logical foundations for personal construct psychology and computational knowledge representation schema. The framework is generated from the single primitive of 'making a distinction'. It has been used to provide cognitive and logical foundations for existing knowledge acquisition tools and techniques, and for the design of integrated knowledge acquisition systems.

In conclusion, we suggest that personal construct psychology provides a very attractive foundation for knowledge acquisition methodologies and tools that have to bridge between human cognitive processes and computational knowledge representation. The many tools so far developed based on personal construct psychology encourage this conclusion. However, what has so far been developed taps only a small part of the potential, and we may expect to see major new developments in the future. In particular, personal construct psychology offers the opportunity to integrate many different knowledge acquisition and knowledge representation approaches within a single principled framework.

Acknowledgements

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada. We are grateful to many colleagues for discussions over the years that have influenced the research described in this paper, in particular John Boose, Jeff Bradshaw, Marc Linster, Alain Rappaport, Brian Woodward and other colleagues at the Knowledge Acquisition Workshops worldwide. We are grateful to Mark Musen for encouraging us to write this review, and to the anonymous referees for many helpful comments. We are grateful to John Boose and Jeff Bradshaw for giving us access to the illustrations of AQUINAS used in this paper.

References

- Abraham, RH and Shaw, CD, 1984. *Dynamics: The Geometry of Behavior* Aerial Press.
- Adams-Webber, JR, 1979. *Personal Construct Theory: Concepts and Applications* Wiley.
- Anglin, JM, 1977. *Word, Object, and Conceptual Development* Norton.
- Beyth-Marom, R and Dekel, S, 1985. *An Elementary Approach to Thinking Under Uncertainty* Lawrence Erlbaum.
- Boose, JH, 1984. "Personal construct theory and the transfer of human expertise". In: *Proceedings AAAI-84*, pp 27–33.
- Boose, JH, 1986a. *Expertise Transfer for Expert Systems*, Elsevier.
- Boose, JH, 1986c. "Rapid acquisition and combination of knowledge from multiple experts in the same domain" *Future Computing Systems* 1 (2) 191–216.
- Boose, JH and Bradshaw, JM, 1987. "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.
- Boose, JH Bradshaw, JM, Kitto, CM and Shema, DB, 1989. "From ETS to Aquinas: Six years of knowledge acquisition tool development". In: *Proceedings of EKAW-89: 3rd European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Paris, pp 502–516. University of Paris.
- Boose, JH, Bradshaw, JM, Koszarek, JL and Shema, DB, 1992. "Better group decisions: using knowledge acquisition techniques to build richer decision models". In: *Proceedings of the 25th Annual Hawaii International Conference on System Sciences*.
- Boose, JH, Shema, DB and Bradshaw, JM, 1990. "Capturing design knowledge for engineering trade studies" *Current Trends in Knowledge Acquisition*, pp 78–89, IOS Press,.
- Borgida, A, Brachman, RJ, McGuinness, DL and Resnick, LA, 1989. "CLASSIC: a structural data model for objects". In: *Proceedings SIGMOD Conference on the Management of Data*, pp 58–67.
- Brachman, RJ, 1977. "What's in a concept: structural foundations for semantic nets" *International Journal of Man-Machine Studies* 9 127–152.
- Brachman, RJ, 1983. "What IS-A is and isn't" *IEEE Computer* 16 (10) 30–36.
- Brachman, RJ and Schmolze, J, 1985. "An overview of the KL-ONE knowledge representation system" *Cognitive Science* 9 (2) 171–216.
- Bradshaw, JM, Boose, JH, Covington, SP and Russo, PJ, 1988. "How to do with grids what people say you can't". In: *Proceedings 3rd AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, Banff, University of Calgary, pp 3-1-3-15.
- Bradshaw, JM, Covington, S, Russo, P and Boose, JH, 1990. "Knowledge acquisition techniques for intelligent decision systems: Integrating Axotl and Aquinas in DDUCKS" *Uncertainty in Artificial Intelligence* 5 Elsevier.
- Cendrowska, J, 1987. "An algorithm for inducing modular rules" *International Journal of Man-Machine Studies* 27 (4) 349–370.
- Christaller, T di Primio, F Schnepf, U and Voss, A, (ed.), 1992. *The AI Workbench Babylon: An Open and Portable Development Environment for Expert Systems*, Academic Press.
- Clancey, WJ, 1991. "The knowledge level reinterpreted: modeling how systems interact" *Machine Learning* 4 (3/4) 285–291.
- Clancey, WJ and Roschelle, J, 1991. "Situated cognition: how representations are created and given meaning" to appear *Educational Psychologist*.
- Coombs, CH, Dawes, RM and Tversky A, 1970. *Mathematical Psychology*, Prentice-Hall.
- CPCS, 1991. *KSS0 Manual*. Centre for Person-Computer Studies, 2019 Underhill Drive, Calgary, Alberta, Canada T2N 4E4.
- Diederich, J, Ruhmann, I and May, M, 1987. "KRITON: A knowledge acquisition tool for expert systems" *International Journal of Man-Machine Studies* 26 (1) 29–40.
- Domany, E, Hemmen, JL and Schulten, K, (ed.), 1991. *Models of Neural Networks* Springer Verlag.
- Eshelman, L, Ehret, D, McDermott, J and Tan, M, 1987. "MOLE: A tenacious knowledge acquisition tool" *International Journal of Man-Machine Studies* 26 (1) 41–54.
- Ford, KM, Cañas, A, Jones, J, Stahl, H, Novak, J and Adams-Webber, J, 1990. "ICONKAT: an integrated constructivist knowledge acquisition tool" *Knowledge Acquisition* 3 (2) 215–236.
- Fransella, F and Bannister, D, (ed.), 1977. *A Manual for Repertory Grid Technique*, Academic Press.
- Gaines, BR, 1988. "Knowledge acquisition: developments and advances". In: *Expert Systems and Intelligent Manufacturing: Proceedings of the Second International Conference and the Leading Edge in Production Planning and Control*, pp 410–434.
- Gaines, BR, 1989. "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:. In: *Proceedings 6th International Workshop on Machine Learning* pp 156–159.

- Gaines, BR, 1991a "Empirical investigations of knowledge representation servers: Design issues and applications experience with KRS" *ACM SIGART Bulletin* 2 (3) 45–56.
- Gaines, BR, 1991b. "Integrating rules in term subsumption knowledge representation servers" *AAAI'91: Proceedings 9th National Conference on Artificial Intelligence* pp 458–463.
- Gaines, BR, 1991c. "An interactive visual language for term subsumption visual languages" *IJCAI'91: Proceedings 12th International Joint Conference on Artificial Intelligence* pp 817–823.
- Gaines, BR, 1991d. "Organizational modeling and problem solving using an object-oriented knowledge representation server and visual language". In: *COCOS'91: Proceedings of Conference on Organizational Computing Systems* pp 80–94
- Gaines, BR, 1992. "Modeling practical reasoning" (to appear in *International Journal of Intelligent Systems*).
- Gaines, BR and Linster, M, 1990. *Development of second generation knowledge acquisition systems*, IOS Press. pp 143–160.
- Gaines, BR, Rappaport, A and Shaw, MLG, 1989. "A heterogeneous knowledge support system". In: *Proceedings of 4th AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, Calgary, Canada, pp 13-1–13-20.
- Gaines, BR and Shaw, MLG, 1980. "New directions in the analysis and interactive elicitation of personal construct systems" *International Journal Man-Machine Studies* 13 81–116.
- Gaines, BR and Shaw, MLG, 1986. "Induction of inference rules for expert systems" *Fuzzy Sets and Systems* 18 (3) 315–328.
- Gaines, BR and Shaw, MLG, 1990. "Cognitive and logical foundations of knowledge acquisition". In: *Proceedings of 5th AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop Banff, Canada*, pp 9-1–9-24.
- Gaines, BR and Shaw, MLG, 1992. "Integrated knowledge acquisition architectures" *Journal for Intelligent Information Systems* 1 (1).
- Gaines, BR Shaw, MLG and Woodward, JB, 1992. "Modeling as a framework for knowledge acquisition methodologies and tools" (to appear in *International Journal of Intelligent Systems*).
- Garg-Janardan, C and Salvendy, G, 1987. "A conceptual framework for knowledge elicitation" *International Journal of Man-Machine Studies* 26 (4) 521–531.
- Hinkle, DN, 1965. The change of personal constructs from the viewpoint of a theory of implications. PhD thesis, Ohio State University.
- Hintikka, J, 1962. *Knowledge and Belief* Cornell University Press.
- Hintikka, J, 1963. "The modes of modality" *Acta Philosophica Fennica* 16 65–81.
- Jones, WP, 1989. "Bringing corporate knowledge into focus with CAMEO" *Knowledge Acquisition* 2 (3) 207–239.
- Kelly, GA, 1955. *The Psychology of Personal Constructs*, Norton.
- Kelly, GA, 1963. "Aldous: the personable computer" *Computer Simulation of Personality*, pp 221–229, Wiley.
- Kelly, GA, 1969. "A mathematical approach to psychology". In: *Clinical Psychology and Personality: The Selected Papers of George Kelly*, pp 94–113, Wiley.
- Kelly, GA, 1970. "A brief introduction to personal construct theory". In: *Perspectives in Personal Construct Theory*, pp 1–29, Academic Press.
- Klinker, G, Bentolila, J, Genetet, S, Grimes, M and Mc Dermott, J, 1987. "KNACK—report-driven knowledge acquisition" *International Journal of Man-Machine Studies* 26 (1) 65–79.
- Lachterman, DR, 1989. *The Ethics of Geometry*, Routledge.
- Linster, M (ed.), 1991. *Sisyphus Working Papers Part 2: Models of Problem Solving* GMD.
- Lloyd, GER, 1966. *Polarity and Analogy*, Cambridge University Press.
- Mac Lane, S, 1971. *Categories for the Working Mathematician*, Springer-Verlag.
- Maher, B, (ed.), 1969. *Clinical Psychology and Personality: The Selected Papers of George Kelly*, Wiley.
- Mancuso, JC and Shaw, MLG, (eds.), 1988. *Cognition and Personal structure: Computer Access and Analysis*, Praeger.
- Marcus, S, 1987. "Taking backtracking with a grain of SALT" *International Journal of Man-Machine Studies* 26 (4) 393–398.
- Maybury-Lewis, D and Almagor, U, (eds.), 1989. *The Attraction of Opposites*, University of Michigan Press.
- McNamee, P and Celona, J, 1987. *Decision Analysis for the Professional – With Supertree*, Scientific Press.
- Motta, E, Eisenstadt, M, Pitman, K and West, M, 1988. "Support for knowledge acquisition in the knowledge engineer's assistant (KEATS)" *Expert Systems* 5 (1) 6–28.
- Rappaport, A, 1987. "Multiple-problem subspaces in the knowledge design process" *International Journal of Man-Machine Studies* 26 (4) 435–452.
- Saaty, TL, 1985. *The Analytic Hierarchy Process*, McGraw-Hill.
- Shapiro, SC, 1979. "The SNePS semantic network processing system". In: *Associative Networks: Representation and Use of Knowledge by Computers*, pp 179–203, Academic Press.

- Shaw, MLG, 1979. "Conversational heuristics for enhancing personal understanding of the world". In: *General Systems Research: A Science, A Methodology, A Technology*, pp 270–277. Society for General Systems Research.
- Shaw, MLG, 1980. *On Becoming A Personal Scientist: Interactive Computer Elicitation of Personal Models Of The World*. Academic Press.
- Shaw, MLG, (ed.), 1981. *Recent Advances in Personal Construct Technology*, Academic Press.
- Shaw, MLG and Gaines, BR, 1983. "A computer aid to knowledge engineering". In: *Proceedings of British Computer Society Conference on Expert Systems* pp 263–271.
- Shaw, MLG and Gaines, BR, 1987. "KITTEN: Knowledge initiation and transfer tools for experts and novices" *International Journal of Man-Machine Studies* 27 (3) 251–280.
- Shaw, MLG and Gaines, BR, 1989. "A methodology for recognizing conflict, correspondence, consensus and contrast in a knowledge acquisition system" *Knowledge Acquisition* 1 (4) 341–363.
- Shaw, MLG and Gaines, BR, 1991a. "Extending electronic mail with conceptual modeling to provide group decision support". In: *COCOS'91: Proceedings of Conference on Organization Computing Systems*, pp 153–158.
- Shaw, MLG and Gaines, BR, 1991b. "Supporting personal networking through computer networking". In: *Proceedings of CHI'91: Human Factors in Computing Systems*, pp 437–438.
- Shaw, MLG and Woodward, B, 1990. "Mental models in the knowledge acquisition process" *Knowledge Acquisition* 2 (3) 179–206.
- Slater, P, (ed.), 1976. *Dimensions of Intrapersonal Space: Volume 1* John Wiley.
- Slater, P, (ed.), 1977. *Dimensions of Intrapersonal Space: Volume 2* John Wiley.
- Sowa, JF, 1984. *Conceptual Structures: Information Processing in Mind and Machine*, Addison-Wesley.
- Voß, A, Karbach, W, Drouven, U, Lorek, D and Schuckey, R, 1990. Operationalization of a synthetic problem. ESPRIT Basic Research Project P3178 REFLECT Task I.2.1 Report. GMD, Bonn.
- Woods, WA, 1975. "What's in a link: Foundations for semantic networks". In: *Representation and Understanding: Studies in Cognitive Science*, pp 35–82, Academic Press.
- Woodward, B, 1990. "Knowledge engineering at the front-end: defining the domain" *Knowledge Acquisition* 2 (1) 73–94.
- Zalta, EN, 1988. *Intensional Logic and the Metaphysics of Intentionality*, MIT Press.
- Zwicky, F, 1969. *Discovery, Invention, Research through the Morphological Approach*, Macmillan.