5th AAAI-Sponsored Knowledge Acquisition for Knowledge-Based Systems Workshop
Banff, Canada, November, 1990.

COGNITIVE AND LOGICAL FOUNDATIONS OF KNOWLEDGE ACQUISITION

Brian R Gaines & Mildred L G Shaw

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

ABSTRACT

The widely used repertory grid knowledge elicitation methodology is based on *personal construct psychology* but utilizes primarily the notion of dichotomous constructs and does not encompass the ordinal relations between them captured in semantic net elicitation. This paper reviews Kelly's formal development of personal construct psychology and gives a systemic basis for it in terms of a calculus of distinctions. This basis is shown to be general enough to encompass category-theory and specific enough to define the computational structures of knowledge representation schema: concepts, constructs, individuals, roles, subsumption, disjoint concepts, primitive concepts and rules. This provides unified cognitive and logical foundations for knowledge acquisition that unifies current techniques and provides design principles for a new family of integrated tools.

INTRODUCTION

The repertory grid has 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 *personal construct psychology* that he had developed on the basis of his clinical and teaching experience. The first 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 the BIAIT methodology from accountants and accounting students using computer-based repertory grid elicitation and analysis (Shaw & Gaines, 1983). Boose (1984) in an independent parallel study reported success in a wide range of industrial knowledge-based system developments using computer elicitation of repertory grids, and since then many knowledge acquisition systems have incorporated repertory grids as a major elicitation technique (Boose & Bradshaw, 1987; Diederich, Ruhmann & May, 1987; Gaines & Shaw, 1987; Garg-Janardan & Salvendy, 1987; Shaw & Gaines, 1987; Ford, Stahl, Adams-Webber, Novak & Jones, 1990).

The 'repertory grid' methodology has itself evolved in the light of application experience and now has major differences from that described by Kelly. Shaw (1980) 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. In knowledge-based systems terms, this can be seen as highlighting correlations that might be spurious and lead to incorrect rules in later analysis. Shaw and Gaines (1979) introduced new forms of analysis of the repertory grid based on fuzzy sets theory 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, Boose, Covington and Russo (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 them to compare and contrast relevant examples (significant people in their life in the original application). This has also been its significance in expertise transfer, that 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 is the basis of the success of repertory-grid techniques.

However, the repertory grid itself 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 (1965) developed a technique of *laddering*, based on "why" and "how" questions, for investigating ordinal relations between constructs, and Boose (1986) incorporated a laddering tool in ETS, but ordinal relations between constructs have not been the primary focus in applications of repertory grid tools.

Ordinal relations are a primary focus in tools that use some form of semantic network to build an overall domain ontology directly, such as MOLE (Eshelman, Ehret., McDermott & Tan, 1987), KNACK (Klinker, Bentolila, Genetet, Grimes, & McDermott, 1987), SALT (Marcus, 1987), ONTOS (Monarch & Nirenburg, 1987), Cognosys (Woodward, 1990), KEATS (Motta, Eisenstadt, Pitman & West, 1988), CODE (Skuce, Shenkang & Beauville, 1989), CAMEO (Jones, 1989), and in techniques based on domain ontologies such as KADS (Breuker & Wielinga, 1987). The tools vary in their interfaces, sources, emphasis on declarative or procedural knowledge, and domain dependence, but all result in domain models based 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.

This paper develops cognitive foundations for current knowledge acquisition tools and techniques in terms of Kelly's personal construct psychology. It uses a systemic analysis of the psychology in terms of a calculus of distinctions (Gaines & Shaw, 1981) to provide a framework for knowledge acquisition that is both psychologically and logically well-founded, and which has direct, operational interpretation in terms of knowledge acquisition tools and techniques.

PERSONAL CONSTRUCT PSYCHOLOGY

Kelly (1955) put forward the idea of an individual using "personal constructs" as filters through which he perceives events: "Man looks at his world through transparent templets which he creates and then attempts to fit over the realities of which the world is composed." (pp.8-9) He continually emphasizes the epistemological status of these constructs in predicting and controlling the world and their ontological status as personal conjectures rather than reality-derived absolutes: "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 was acutely aware of the psychologist as an observer modeling a person and objected to the introduction of theoretical constructs such as "learning" and "motivation" as if they were causal variables. Hence, he was extremely careful in the exposition of his own system to make its underlying presuppositions minimal and overt by developing personal construct psychology as a semi-formal system based on a postulate and corollaries. He made minor but significant changes to the corollaries after the publication of his book, and those below are updated from Bannister and Mair (1968).

"**Fundamental Postulate:** A person's processes are psychologically channelized by the way in which he anticipates events."

This was stated as a postulate to emphasize that it was presented as a convenient viewpoint from which to understand human behavior, not imputed to an underlying physiological or psychological reality. This is consistent with modern views of the role of theory in science, but was a fairly radical position in Kelly's time. The basis of Kelly's approach to psychotherapy was that this was also a convenient viewpoint from which someone could understand, and modify, their own behavior. He saw all people as "personal scientists" in anticipating the world, and attempted to develop techniques where this anticipatory modeling activity was reflexively applied to the self.

This fundamental emphasis on the role in behavior 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 behavior as deriving from this. He notes in expanding upon this point that "a person's processes, psychologically speaking, slip into the grooves which are cut out by the mechanisms he adopts for realizing his objectives." (Kelly, 1955, p.49)

"**Construction Corollary:** A person anticipates events by construing their replications."

Construction involves two basic operations, slicing out an "event" and interpreting it by "construing." Kelly notes "we look at the undifferentiated stream of circumstance flowing past us, and we try to find something about it that repeats itself. Once we have abstracted that property, we have a basis for slicing off chunks of time and reality and holding them up for inspection one at a time...this is done by construing" (p.120) "By construing we mean placing an interpretation: a person places an interpretation on what is construed. He erects a structure, within the framework of which the substance takes shape or assumes meaning." (p.50). "When one abstracts replicated properties in the events he has already experienced, it becomes possible for him to chart events to come in terms of these same properties. ... What one predicts is not a fully fleshed-out event, but simply the common intersect of a set of properties. If an event comes along in which all the properties intersect in the prescribed way, one identifies it as the event he expected." (p.121)

It is these "slices of circumstance" that we elicit from an expert as "critical cases", and these "interpretations" that we elicit as "relevant attributes."

"**Individuality Corollary:** Persons differ from each other in their constructions of events."

Kelly stresses the personal nature of construing but he emphasizes also that "persons can find common ground through construing the experiences of their neighbors along with their own", although "individuals can be found living out their existence next door to each other but in altogether different subjective worlds." (p.56) In knowledge-based system terms it is the personal, non-overt, nature of expertise that makes knowledge elicitation from experts a necessary developmental technique.

"**Organization Corollary:** Each person characteristically evolves, for his convenience of anticipating events, a construction system embracing ordinal relationships between constructs."

Kelly notes that "man systematizes his constructs by concretely arranging them in hierarchies and by abstracting them further" (p.58) He saw the hierarchy as one in which the stability of the upper levels allowed for change in the lower ones and that the arrangement of constructs in a particular hierarchy was more a determinant of personality than the constructs themselves. In knowledge-based system terms, hierarchies and abstractions are the material of semantic nets and software engineering. Kelly's approach emphasizes that they are constructs of the modeling system and characterize its particular "expertise."

"**Dichotomy Corollary**: A person's construction system is composed of a finite number of dichotomous constructs."

Kelly's emphasis on the dichotomous, bipolar nature of constructs is what most distinguishes them from what are conventionally called "concepts." "The construct, of itself, is the kind of contrast one perceives... while constructs do not represent or symbolize events, they do enable us to cope with events, which is a statement of a quite different order...They also enable us to put events into arrays or scales, if we wish." (pp.13-14) Kelly saw the linked concepts of *similarity* and *contrast* between elements forming the two "poles" of a construct. This does not necessarily involve a binary evaluation of elements since they may lie somewhere between the poles.

Kelly emphasizes that the opposition between the poles of a construct is not that of logical negation—"Now conventional logic would say that black and white should be treated as separate concepts. Moreover, it would say that the opposite of black can only be stated as not black, and the opposite of white can only be stated as not white. Thus the person whose field we mentioned would have shoes which would be just as much not white as the time of day, and he would write on paper which would be just as not black as the distance to his office." (p.106)

The notion that dichotomous pairs of concepts are fundamental to human mental process is an ancient one, strongly present in Greek philosophy (Lloyd, 1966), and found in many anthropological studies of the concepts defining a culture (Maybury-Lewis & Almagor, 1989). Eliciting dichotomous constructs by asking in what way are two entities in a domain alike and different from a third is the basic technique used in repertory grids. Semantic net elicitation focuses on unipolar concepts, although the meta-concept that two concepts are *disjoint* from one another can be significant.

The relation between constructs and concepts is both psychologically and theoretically important, and the next section analyzes it in detail.

"**Choice Corollary**: A person chooses for himself that alternative in a dichotomized construct through which he anticipates the greater possibility for the elaboration of his system."

This is Kelly's basis for human action which he saw as an alternative to concepts of motivation—"there is a continuing movement toward the anticipation of events, rather than a series of barters for temporal satisfactions" (p.68).

"**Range Corollary**: A construct is convenient for the anticipation of a finite range of events only."

This may be thought of as introducing a concept of relevance into the definition of construct systems.

"**Experience Corollary**: A person's construction system varies as he successively construes the replications of events."

This expresses the essential fallibility and conjectural nature of constructs. They are hypotheses to be judged by their utility and updated as they fail in anticipation. This updating is what the observer construes as "learning."

"**Modulation Corollary**: The variation in a person's construction system is limited by the permeability of the constructs within whose ranges of convenience the variants lie."

Kelly sees a "permeable" construct as one "so constituted that new experience and new events can be discriminatively added to those which it already embraces." (p.81)

"**Fragmentation Corollary**: A person may successively employ a variety of construction subsystems which are inferentially incompatible with each other."

Through this corollary Kelly allows for the operations of different psychological systems and hence different methods of inference and forms of rationality within a single individual. Pask (1975) sees "roles" within a person in similar terms, as forming separate psychological individuals, "P-Individuals." In these terms, an "expert" is not a person but a specialized and valued role within a person.

"**Commonality Corollary**: To the extent that one person embraces a construction of experience which is similar to that employed by another, his processes are psychologically similar to those of the other person."

This is complementary to the individuality corollary and allows for the "fact that certain groups of people behave similarly in certain respects." (p.93)

"**Sociality Corollary**: To the extent that one person construes the construction processes of another, he may play a role in a social process involving the other person."

Kelly is here allowing the construal processes of one person to become elements for construal by another and hence to model the process by which we anticipate the actions of others.

MAKING DISTINCTIONS

A theoretical framework that encompasses personal construct psychology, knowledge representation and acquisition for knowledge-based systems, and the relations between these topics, can be based on a single fundamental primitive, what Brown in his *Laws of Form* terms "making a distinction."

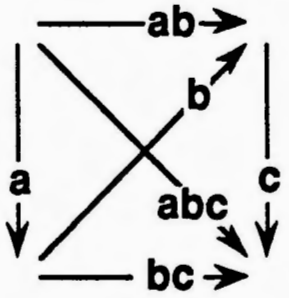
> "The theme of this book is that a universe comes into being when a space is severed or taken apart... By tracing the way we represent such a severance, we can begin to reconstruct, with an accuracy and coverage that appear almost uncanny, the basic forms underlying linguistic, mathematical, physical and biological science, and can begin to see how the familiar laws of our own experience follow inexorably from the original act of severance." (Brown, 1969 p.v)

We have shown that a wide range of psychological and philosophical frameworks are characterized by the systems of distinctions underlying them, and that there are a number of basic distinctions common to such systems (Gaines & Shaw, 1984). The frameworks included that of Kelly's personal construct psychology, and this section defines those parts of a calculus of distinctions adequate to do this.

Formally, a *distinction* is taken to be a primitive notion, but its connotations of cutting out something for some purpose, or because it satisfies a criterion, are informally intended. The connotations that there may be an agent making the distinction, and that something is distinguished are also informally useful. However, it should be noted that closed-world assumptions are specifically excluded, it is not assumed that what is *not* distinguished is meaningful or has the same categorial status as that which is distinguished. Informally, only that which is distinguished is known, new distinctions can be made indefinitely, and the world is assumed open to novelty. Formally, this affects the way relations between distinctions are defined. Extensional forms, such as "for all", are avoided because it is not reasonable to assume that statements can be made about an indefinite extension whose members are not yet identified and may never be known.

Distinction are the only primitives. One might ask "what is distinguished?", but the answer will be "distinctions." This corresponds to deriving the two phenomena in Kelly's *construction corollary* from a single primitive, that we distinguish events in the undifferentiated stream of circumstance, and then we further distinguish among the distinguished events by construing.

Translating this into an algebraic formalism, if distinctions are represented by bold lower case letters, and a distinction applied to another distinction is concatenated to the right of it, what has been defined is a partial groupoid (Bruck, 1971) over an alphabet of distinctions, that is a set of expressions of the form: **a, b, c, ad, bd, be, ce, cf,** and so on. Allowing further distinctions to be applied to those already distinguished, so that expressions of the form **adg** have meaning, extends the groupoid. If, in addition, sequences of distinctions can be named as a single distinction, with the naming convention that these are to be interpreted only through context-insensitive expansion, the groupoid becomes a semigroupoid. That is, it is associative: **a(bc) = (ab)c** where the equality may be interpreted as "means the same as." In graphical terms, the following diagram commutes:



This semigroupoid structure turns out to be adequate for exposition of personal construct psychology, and for current knowledge representation schema. This rich capability from such slight resources is, perhaps, less surprising if it is noted that adding the requirement for left and right identity distinctions to the above structure gives the formal structure of a *category* (MacLane, 1971), an algebraic structure sufficient to support all current systems of mathematics and logic (MacLane, 1986). Thus, the interesting questions are about the extent and naturality of the infrastructure that one has to erect in terms of the basic calculus of distinctions in order to derive an adequate semantics for personal construct psychology and knowledge representation.

DISTINCTIONS AMONG DISTINCTIONS

One distinction will be said to *subsume* another if it can always be applied whenever the other can. This notion of subsumption turns out to be 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. It can be represented formally as:

"**b** subsumes **a**" $\quad \mathbf{a} \rightarrow \mathbf{b} \Leftrightarrow \vdash \mathbf{xa} \Rightarrow \vdash \mathbf{xb}$ (1)

That is, **b** subsumes **a**, if and only if whenever one asserts **xa** one also asserts **xb**. Note that the form $\forall \mathbf{x}$ is avoided for the reasons given above, that the notion of all the distinctions to which **a** and **b** may be applied is not well-defined. However, if this extension is well-defined then the intensional formulation of (1) implies the extensional formulation in terms of $\forall \mathbf{x}$. Marcus' (1962) non-standard, *substitutional* reading of the universal quantifier would be appropriate to the definition above, but the notation is avoided because the extensional reading is prevalent and significant in its own right.

The definition in (1) is best 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. Subsumption corresponds to increasing generality since the subsuming distinction can be applied to at least as many things as that subsumed. It is an asymmetric, transitive relation, a partial pre-order, over distinctions, that supports the ordinal relations of Kelly's *organization corollary.*

There is a second relation definable in similar terms which is equally important, that one distinction is *disjoint* with another in that one can never be applied whenever the other can. This definition of a disjoint relation turns out to be adequate to capture Kelly's notion of a dichotomy, and also the notion of disjoint concepts in knowledge representation. It can be represented formally as:

"**a** disjoint **b**" $\quad \mathbf{a—b} \Leftrightarrow \vdash \mathbf{xa} \Rightarrow \neg \vdash \mathbf{xb}$ (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 best 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 supports Kelly's *dichotomy corollary.*

The → and — relations are complementary in establishing four possible binary relations between distinctions, that **a→b**, **b→a**, **a—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 subsumption relations, and is *inherited* through subsumption, that is:

$$\mathbf{a—b} \text{ and } \mathbf{a→c} \Rightarrow \mathbf{c—b} \qquad (3)$$

These notions are sufficient for the definition of concepts and constructs. A *concept* is defined to be *that mental entity imputed to a distinction making agent as enabling it to make a particular distinction.*

Note that concepts are separated both from the distinctions they support and the entities they distinguish. This definition corresponds to Anglin's (1977, p.3), "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." The notion of an agent and its mentality are deliberately introduced at this stage because they are part of the connotations of a concept. However, consistent with Kelly's position, concepts are defined to be *imputed* to an agent rather than as a mental structure. They are themselves distinctions made by an observer, possibly a reflective observer. This is equivalent to Zadeh's (1964) introduction of the notion of the *state* of a system as an imputed parametrization accounting for a system's behavior. *Concepts are state variables we impute to a knowledgeable agent.*

A construct is defined formally to be a triple of two disjoint distinctions mutually subsumed by a third as shown in Figure 1, and psychologically as the triple of concepts assumed to underlie the distinctions.

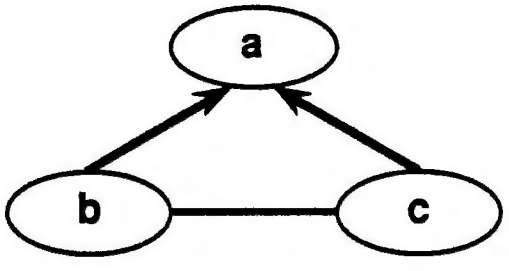


Fig.1 The triple of distinctions generating a construct

The triple captures the notions of *similarity* through the shared subsuming concept, **a**, and *contrast* through the disjoint subsumed concepts, **b** and **c**. The pair of disjoint concepts, **b** and **c**, captures Kelly's *dichotomy corollary*, that a person's construction system is composed of a finite number of dichotomous constructs. The concepts **b** and **c** characterize what Kelly terms the *poles* of the construct. A minimal subsuming concept captures Kelly's *range corollary*, that a construct is convenient for the anticipation of a finite range of events only. Such a minimal concept **a** characterizes what Kelly terms the *range of convenience* of the construct. The importance of such triples derives from the restricted closed world semantics generated, that the world carved out by concept **a** is divided by the disjoint concepts **b** and **c** into two parts that are now at the same categorial level.

The construction of Figure 1 can be simplified by noting that a subsuming concept **a** can be constructed for a pair of disjoint concepts **b** and **c**, and it is the disjunction which primarily defines the construct. In the graph of the disjoint relation between concepts, concepts are nodes and constructs are arcs. Hence *concept and construct are graph-theoretically dual relations.*

CONCEPTUAL STRUCTURES

The subsumes and disjoint relations, → and —, can between them characterize a wide variety of concept and construct structures. The left hand structure in Figure 2 shows that a given range of convenience can be construed through a number of unrelated constructs. The pairs **b**—**c** and **d**—**e** are two different dimensions of the domain characterized by **a**. If **a** characterizes a domain of people then **b**—**c** might be **young**—**old** and **d**—**e** might be **short**—**tall**.

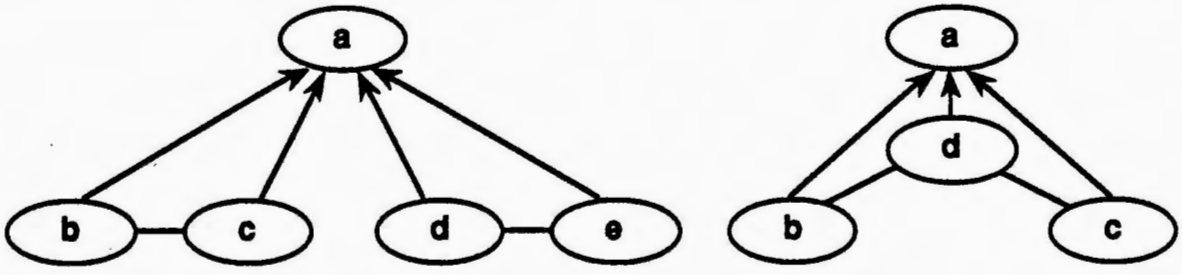


Fig.2 Alternative constructs and alternative poles

The right hand structure in Figure 2 shows that a pole of a construct may be paired with alternative poles that are not mutually disjoint. The pairs **b**—**d** and **c**—**d** are two different dimensions of the domain characterized by **a** that have some commonality through the shared pole, **d**. If **a** characterizes a domain of people then **b**—**d** might be **young**—**old** and **c**—**d** might be **fun**—**old**. It was because different contrasts were possible for the same concept that Kelly saw the concept pairs that form constructs as being a more stable chararacterization of psychological structure, that the way in which we react to something falling under a concept is dependent on what we see as the contrasting concept.

If the concepts **b** and **c** on the right of Figure 2 are also disjoint as shown on the left of Figure 3 then the structure may either be regarded as a trichotomy or a triple of dichotomous constructs. In knowledge representation terms, **b**, **c** and **d** might be seen as alternative values along some dimension such as young, old and middle-aged respectively. In personal construct terms one can see that the contrasts **young**—**middle-aged**, **middle-aged**—**old**, and **young**—**old**, might all be used in different situations.

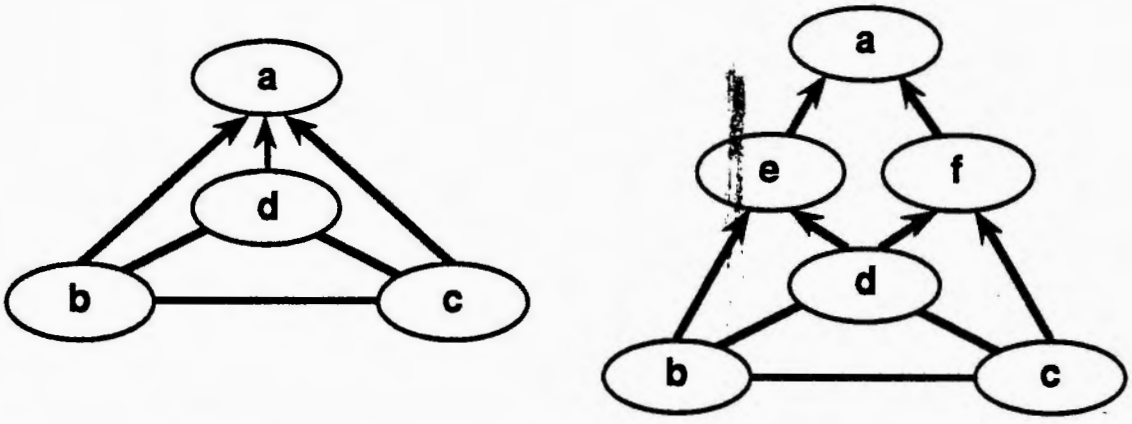


Fig.3 Trichotomous construct and a scalar trichotomy

What the structure on the left of Figure 3 does not capture is the possibility of putting events along scales which Kelly mentions in his dichotomy corollary. Middle-aged is not just an alternative to young and old, but an intermediary possibility between them. The structure on the right of Figure 3 extends that on the left to show how scales may be developed. The concepts **e** and **f** are not disjoint and **d** inherits their commonality while **b**, since it is disjoint with **d**, inherits a different part of **e**, and **c** similarly inherits a different part of **f**. For example, suppose that **b** is young interpreted as an age range [0, 35], **c** is old as an age range [50, 100], **d** middle-aged [35, 50], then **e** might be the age range [0, 50] and **f** [35, 100].

It is possible for the pair **e f** to be used as if it were a dichotomous construct even though the concepts are not disjoint. This generates the underlying infrastructure shown in Figure 4 where **d** captures the borderline cases (Sanford, 1975) that come under both **e** and **f**, **b**—**c** is a dichotomous construct underlying **e**—**f**, and **g**—**h** is a dichotomous construct characterizing that part of the borderline under **b** and that under **c**. The structure **b**, **c**, **d** is the basis for a fuzzy sets (Zadeh, 1965) characterization of the pseudo-construct **e f**, with, for example: **b** characterizing a membership of 1 to **e** and 0 to **f**; **c** a membership of 0 to **e** and 1 to **f**; and **d** 0.5 to **e** and 0.5 to **f**. Distinguishing more concepts subsumed by **d** allows for finer grades of membership (Gaines, 1976).

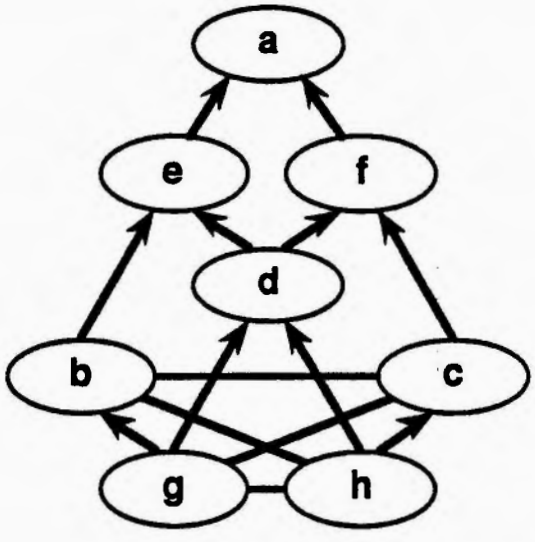


Fig.4 A non-dichotomous pseudo-construct, borderline and fuzzy infrastructure

The relation of **g—h** to **b—c** in Figure 4 is an example of an ordinal relation between constructs as described by Kelly in his *organization corollary*. Figure 5 shows two such ordinal relations discussed by Kelly. On the left, the concept characterized by the triple **d, e, f,** is subordinate to that characterized by **a, b, c.** For example, **b—c** might be **good—bad,** and **e—f** might be **fun—boring.** When the upper construct is evaluatory, such as **good—bad,** its subordinate constructs inherit the evaluatory connotations and typically have a *preferred pole*.

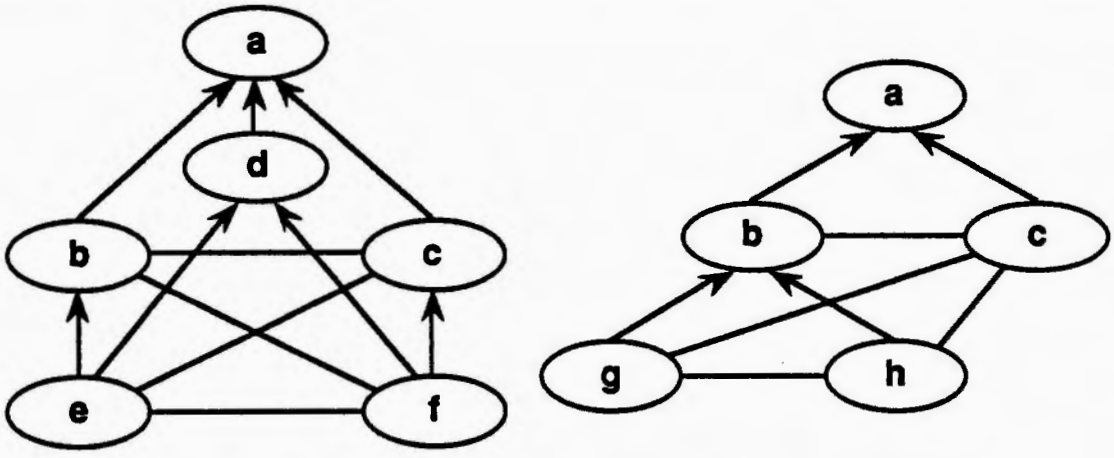


Fig.5 Some ordinal relations between constructs

On the right of Figure 5, the concept characterized by the triple **b, g, h,** is subordinate to that characterized by **a, b, c.** For example, **b—c** might be **good—bad,** and **g—h** might be **dancing—swimming.** The ordinal relations between constructs derive from the ordinal relation of subsumption between the underlying distinctions and an indefinite number of subordination structures are possible.

KNOWLEDGE REPRESENTATION

The calculus of distinctions developed above and discussed in terms of its cognitive connotations in personal construct psychology also provides computational foundations for knowledge representation schema such as those of KL-ONE (Brachman & Schmolze, 1985) and CLASSIC (Borgida, Brachman, McGuinness & Resnick, 1989). These schema involve concepts, primitive concepts, disjoint concepts, individuals, roles, constraints upon roles, individuals filling roles, rules, subsumption between concepts, and recognition of individuals by concepts. They encompass semantic nets, frames, and rule-based systems, as well as relational and post-relational databases. This section shows how such systems may be derived from primitives underlying systems of distinctions.

A *concept* in a computational context is defined to be *that data structure used by a distinction making algorithm to enable it to make a particular distinction.* No imputation is necessary because we are not dealing with identification of a black box from its behavior, but rather with the internal structure of an accessible white box. The *disjoint* relation between concepts already defined corresponds precisely to the disjoint relation in knowledge representation schema. Similarly, the *subsumption* relation between concepts already defined corresponds precisely to the "is-a" subsumption 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 below, and also allow primitive and non-primitive concepts to be defined.

Individuals in a computational context are defined to be *those distinctions which are initial in that they cannot be preceded by other distinctions*. This corresponds to the notion of an *initial object* in category theory and also to the analysis of individuals in metaphysics. Gracia (1983) notes that traditionally individuals are characterized by five properties; they are entities which:

(a) lose their fundamental character if they are divided into parts;
(b) are distinct from all other entities, even those that are of their same type or kind;
(c) are part of a group type or class which has or can have several members;
(d) remain fundamentally the same through time and various changes;
(e) are the subject of predication but not predicable of other things.

Gracia argues against each of these in turn as being a necessary or sufficient feature to characterize individuals and suggests instead that it is *non-instantiability* that characterizes an individual; that is we cannot instantiate an individual. This corresponds to an initial distinction which cannot have another, an instance, to its left.

Thus, while a distinction is the sole primitive, the difference between initial distinctions and non-initial ones is an important one. Individuals are concepts that cannot be applied to other concepts and hence have no extensions or instances and cannot subsume one another. However, since individuals are concepts it is possible to have knowledge representation schemas which do not make the distinction explicitly. It is also possible to treat non-initial concepts as *abstract individuals* that represent all the individuals that might be distinguished by them. When such an abstract individual has no extension, that is, there are not, or cannot be, any individuals that it distinguishes, then the paradoxes of sense and reference studied by Frege (1892), Meinong (1910) and Russell (1905), can arise. The "morning star, evening star" and "Scott, author of Waverley" equivalences are not problematic as they corresponding to two concepts subsuming the same individual. Non-existent objects such as "the present King of France" are also not problematic because there is no dependence on their empty extensions. They are interesting because, like fictional objects (Woods, 1974), they mimic individuals in that they do not (rather than cannot) have another distinction to their left. This may well underlie Meinong's intuition that they are, in some sense, equivalent to real objects.

A *role*, *property*, *attribute* or *relation* arises naturally by an important construction as shown in Figure 6. The triple of concepts defining a construct on the upper left may be used to generate that on the right by *defining* the concept **d** to be precisely the union of concepts **b** and **c**, in the sense that **d** subsumes a distinction if, and only if, it is subsumed by **b** or subsumed by **c**:

"**d** internalizes **b**—**c**" $\quad$ **d**←**b**—**c** $\Leftrightarrow \vdash \mathbf{xd} \Leftrightarrow \vdash \mathbf{xb} \vee \vdash \mathbf{xc}$ (4)

The definition in (4) reads as a commitment to making the distinction **d** if and only if there is a commitment to either **b** or **c** being made also. As shown in the lower left of Figure 6, the concept **d** is uniquely minimal in the subsumption ordering as a range of convenience for the construct defined by **a** , **b**, and **c**. Thus, the construct defined by **d**, **b**, and **c** precisely and uniquely captures the essence of the original construct.

The notation shown in the definition of **d** indicates that it is defined as subsuming the construct **b**—**c** and subsumes what is distinguished by **b** or by **c**. The construct **b**—**c** becomes a *role* and the alternation of the concepts **b** or **c** becomes a *constraint* on the role. This internalization process can be applied to other structures also such as the trichotomies of Figure 3 in which case the role can take more than two values, as a categorical role (left of Figure 3) or scale (right of Figure 3). Non-dichotomous pseudo-constructs, such as that of Figure 4, may also be internalized giving multiple-valued roles or relations. It is natural also, particularly as more complex structures are internalized, to name the role in its own right rather than through its poles.

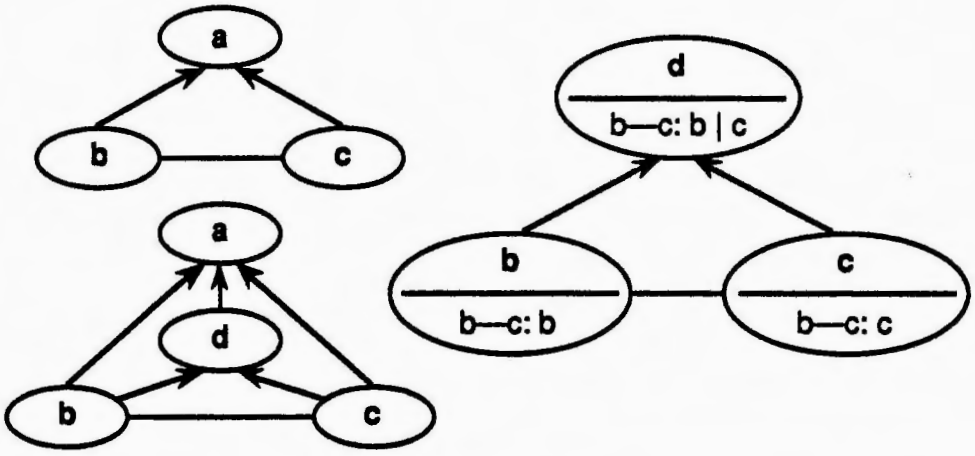


Fig.6 A role as the internalization of a construct

Multiple constructs with the same range of convenience may be internalized to define a concept with multiple roles as shown in Figure 7. The concept **i** may not be precisely superordinate to the constructs **b—c** and **d—e** because they may not cover the same range, but it does arise as the subordination of **g** and **h** which are the precise forms of **b—c** and **d—e.**

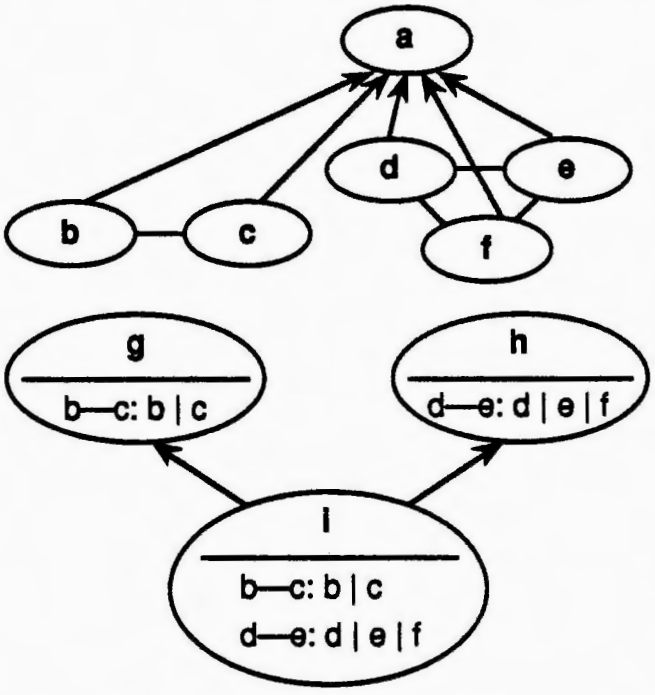


Fig.7 Multiple roles defining a concept

The internalization of constructs to form roles and constraints provides a natural coordinate system specifying individuals in terms of the concepts that distinguish them. The assertion of **xb** may be represented as the individual **x** having the value **b** on the role **b—c.** Note that, as individuals cannot subsume concepts, it is not possible to use the internalization construction. The coordinate system appears similar but is generated through a different construction. This corresponds to Zalta's (1983) distinction in his intensional logic of abstract objects that *abstract objects **encode** properties whereas existent objects cannot encode properties but do **exhibit** them.*

The internalization construction of Figures 6 and 7 gives some concepts a defined internal structure in which constructs define coordinate axes named as roles and constrained by the concepts forming the poles of the construct. Individuals have values on these axes and the values are themselves concepts or individuals. There is an important distinction between concepts that results from this construction: the subsumption relation between concepts defined completely in terms of their internal structure is determined by that structure rather than by a separate commitment. Such structurally defined concepts are said to be *non-primitive* and the others *primitive*.

Thus 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 computed, 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). A similar confusion was also a problem in differentiating the definition of concepts from the assertion of those applicable to individuals (Woods, 1975), the encoding—exhibiting distinction above. The generation of concepts, roles, individuals, constraints and values from the single primitive of making a distinction clarifies all these issues.

An important consequence of the structural definition of non-primitive concepts is that individuals may be *recognized* as being distinguished by one of these concepts even if this distinction has not been asserted. That is, other assertions made about an individual may be sufficient for them to satisfy the constraints of a non-primitive concept. This enables the anticipatory origins of distinction-making specified in Kelly's fundamental postulate to be captured. A *rule* is defined to be a link between concepts such that the recognition of an individual as distinguished by the first concept leads to the assertion that it is distinguished by the second:

"a rule b" $\quad \mathbf{a} \rightarrow\rightarrow \mathbf{b} \Leftrightarrow \mathbf{xa} \Leftrightarrow \vdash \mathbf{xb}$ (5)

Notice that this differs from subsumption as defined in (1) in that **xa** need not be asserted, the commitment is to assert **xb** if **xa** is noted through assertion or recognition. The double arrow notation is chosen to portray both the similarities and differences between a rule and a subsumption relation.

Rules allow a cognitive system to be anticipatory in that it contains structures which from one set of distinctions made about an event will infer that another should be made. This can lead to prediction or action. Rules play a similar role in computational systems in allowing recommendations for decision or action to be made. Overtly modeling the conceptual system of an expert as a structure of concepts, roles, individuals and rules is a basis for emulating the expert's performance in a knowledge-based system.

DISTANCE MEASURES, CONCEPTUAL CLUSTERING AND INDUCTION

Distance measures play an important role in conceptual clustering and induction. There is a natural construction, a distance between two concepts, **x** and **y**, as shown on the left of Figure 8. Let **u** be some minimal upper bound of **x** and **y** subsuming both of them, and **l** some maximal lower bound subsumed by both of them, and **U** be the extension of **u** and **L** the extension of **l** 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" $\quad d(\mathbf{a}, \mathbf{b}) = \mathcal{C}\mathbf{U} - \mathcal{C}\mathbf{L}$ (6)

where $\mathcal{C}\mathbf{U}$ and $\mathcal{C}\mathbf{L}$ 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, $\mathcal{C}\mathbf{U}$. 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.

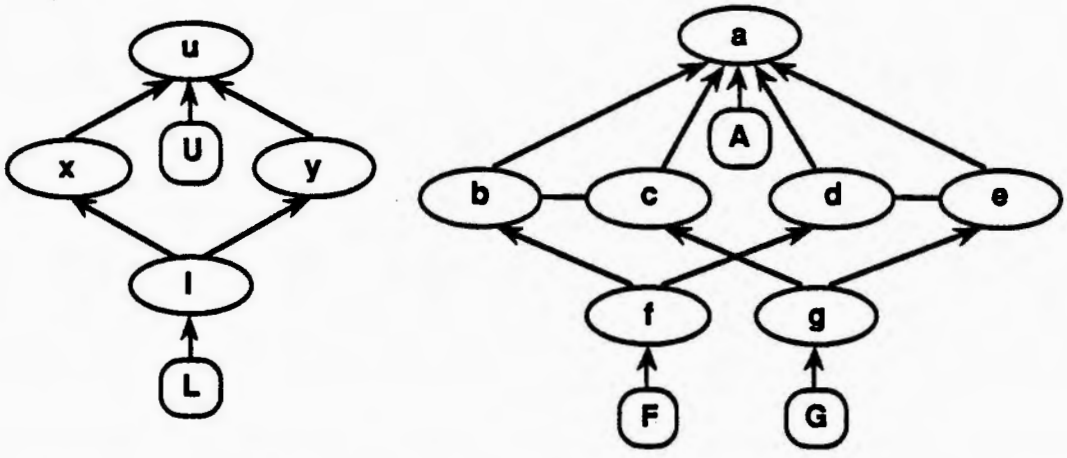


Fig.8 Calculation of distance measures between concepts and between constructs

The distance measure defined readily extends to constructs through the comparison of poles as shown on the right of Figure 8:

"b—c distance d—e" $d(b—c, d—e) = CA - CF - CG$ (7)

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.

What of trichtomous constructs such as those shown in Figure 3? The construction of Figure 8 readily extends to the categorical case on the left but does not capture the scalar semantics on the right. The construction on the left of figure 9 shows how the intermediate poles may be taken into account by also examining distances between the upper concepts that represent the scale:

$$d(b—d—c, g—l—h) = (CA - CF - CG - CG) + (CA - CO - CP) \quad (8)$$

This measure counts one for an individual at the extreme of one construct that falls under the center pole of the other, and two if it under the opposite extreme. 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).

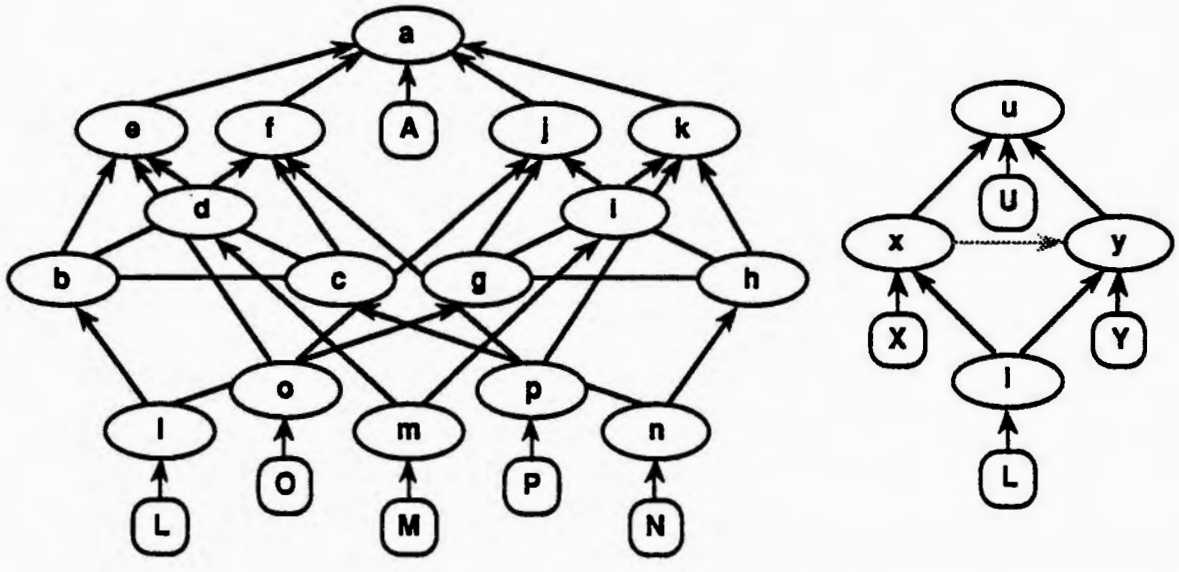


Fig.9 Scalar measures between constructs and induction of rules between concepts

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, fo 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).

The measures used in the induction of a rule linking to concepts are also readily derived as shown on the right of Figure 9. $\mathcal{CX}$ is the number of "predictions" made by concept x as the left hand side of a rule, and $\mathcal{CL}$ is the number which are correct. Thus, the measures of the validity of inducing the rule, $x \rightarrow y$, are:

"prior probability of y" $p(y) = \mathcal{CY}/\mathcal{CU}$ (9)

"probability correct $x \rightarrow\rightarrow y$" $p(x \rightarrow\rightarrow y) = \mathcal{CL}/\mathcal{CX}$ (10)

"probability by chance $x \rightarrow\rightarrow y$" $c(x \rightarrow\rightarrow y) = I_{p(y)}(\mathcal{CX}\text{-}\mathcal{CL}, \mathcal{CL}+1)$ (11)

where I is the incomplete beta function summing a binomial distribution tail. These measures are precisely those used by Induct (Gaines, 1989) inducing rules from datasets.

PRACTICAL APPLICATIONS

The knowledge representation primitives developed in this paper are not just theoretical constructs but have immediate practical application in knowledge acquisition tools. In particular, the diagrams form the basis of graphic knowledge editors providing a visual, formal language for entering and editing knowledge structures. For example, Figure 10 shows the minimal solution to Cendrowska's (1987) contact lens problem graphed in the graphic editor associated with KRS, our knowledge representation server (Gaines, 1990).

The three bipolar constructs and three tripolar constructs involved are graphed at the top of Figure 10. The five default rules are graphed in the center, and two individual cases are graphed at the bottom. The convention in the grapher is that concepts are represented by ovals, indviduals by rectangles, roles (which are not used in this example) by unboxed text, and special expressions (such as cardinality constraints) by rounded-corner rectangles. The lines and arrows have exactly the interpretation given in this paper, a plain line denotes a disjoint relation between primitive concepts, an arrow from one concept to another denotes that the second concept subsumes the first, the double arrow for a rule is broken with an expression "Rule" for perspicuity. Primitive concepts are indicated by small interior lines on the left and right sides which line up with the line indicating a disjoint relation in a construct.

The graphic editor itself operates in a very similar way to MacDraw with additional features taking advantage of the defined syntax and semantics of the knowledge structures to simplify data entry, for example text is automatically wrapped to as rectangular a shape as possible, boxes are automatically fitted to text, lines reconnect when items are dragged, and arcs between boxes cannot be entered if they violate the syntax. Knowledge structures in KSS0 grids and rules (Gaines & Shaw, 1987), or in KRS knowledge bases (Gaines, 1990), can be output to the graphic editor, edited and returned to KSS0 or KRS.

Since the knowledge structures in the graphic editor conform exactly to the formal schema defined in this paper it is possible to compile them directly to operational knowledge structures that will run as knowledge-based systems in a suitable shell. Thus the graphic editor provides a *visual programming language* for knowledge-based systems, one with a clear inheritance from semantic nets but now having a simple and well-defined formal semantics.

Figure 11 shows the output when the graph in Figure 11 is compiled for KRS (which provides a KL-ONE/CLASSIC-like language interface. Primitives are defined as a list of the previously defined disjoint concepts followed by the name of the new primitive to be defined. Non-primitive concepts and individuals are defined by the name followed by the defining concepts. Rules are defined by the name of the concept defining the recognition of the left-hand side followed by the concepts to be asserted on the right-hand side. The knowledge base output in Figure 11 runs directly in KRS and infers the lens recommendation for the individuals specified.

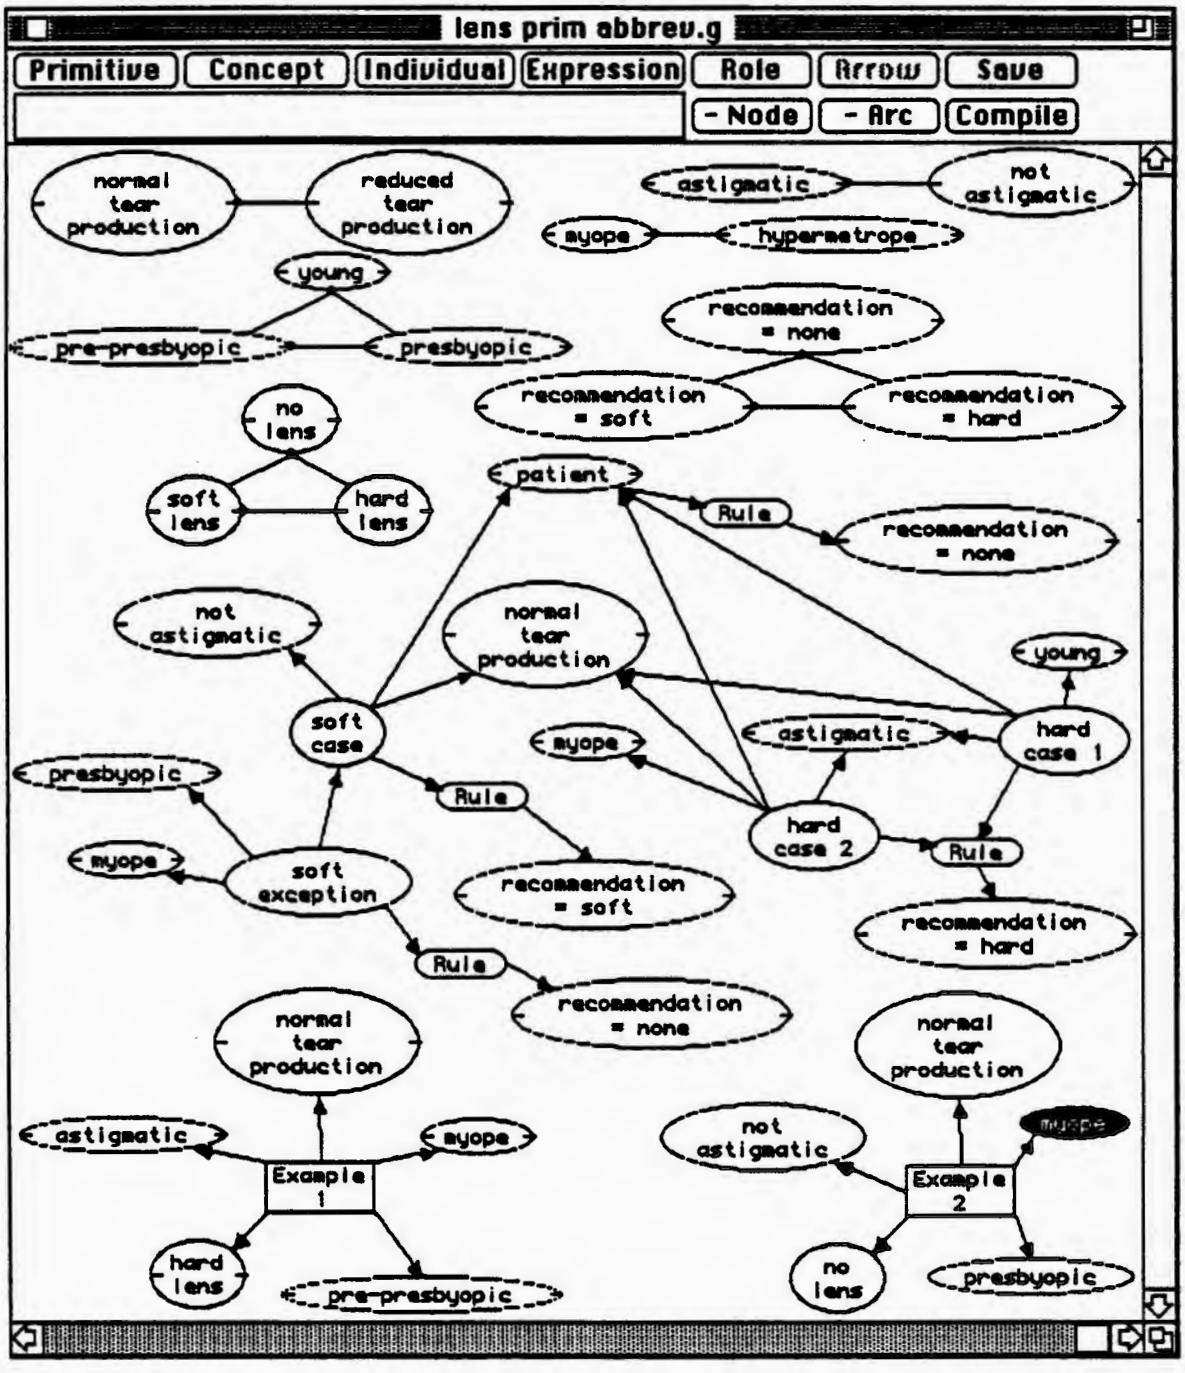


Fig.10 Graphic knowledge editor

Primitive(patient)
Primitive(normal tear production)
Primitive(normal tear production, reduced tear production)
Primitive(astigmatic)
Primitive(astigmatic, not astigmatic)
Primitive(myope)
Primitive(myope, hypermetrope)
Primitive(young)
Primitive(young, pre-presbyopic)
Primitive(young, pre-presbyopic, presbyopic)
Primitive(no lens)
Primitive(no lens, soft lens, hard lens)
Primitive(no lens, soft lens)
Primitive(recommendation = none)
Primitive(recommendation = none, recommendation = soft)
Primitive(recommendation = none, recommendation = soft, recommendation = hard)
Concept(hard case 1, patient, normal tear production, astigmatic, young)
Concept(hard case 2, patient, normal tear production, astigmatic, myope)
Concept(soft case, patient, normal tear production, not astigmatic)
Concept(soft exception, presbyopic, soft case, myope)
Rule(patient, recommendation = none)
Rule(hard case 1, recommendation = hard)
Rule(hard case 2, recommendation = hard)
Rule(soft case, recommendation = soft)
Rule(soft exception, recommendation = none)
Individual(Example 1, normal tear production, astigmatic, myope, pre-presbyopic, hard lens)
Individual(Example 2, normal tear production, not astigmatic, myope, presbyopic, no lens)

Fig.11 Visual knowledge base in the graphic editor compiled for use in the knowledge representation server

Note in Figure 11 the way in which the top-down specification of formal knowledge structures leads to a definitional assymetry in the disjoint relation; **reduced tear production** is defined as disjoint to the already-defined **normal tear production**. No logical assymetry is introduced and definition (2) still applies, but a natural psychological assymetry is replicated: that we define opposition in terms of an existing concept. Although the disjoint relation is logically symmetric, it is, at least initially, psychologically assymetric. In terms of Hegel's (1812) *Science of Logic*, the initial concept is the *thesis*, its disjoint opposition is the *antithesis*, and the resultant role representing the dimension of variation is the *synthesis*.

The compiler checks for circular definitions and reports them as errors which must be rectified before the knowledge structure is used. It automatically sorts the definitions according to dependencies so that the concepts are specified in top-down form for a knowledge-base. This is, in itself, a significant support for knowledge elicitation since the visual knowledge structures may be developed in whatever way is natural, top-down, bottom-up or middle-out, and the top-down definitional requirement that is a major problem in knowledge base development (Morik, 1987) is no longer a burden.

In addition, the compiler operates in the same way over collections of graphs allowing pre-existent concept definitions to be simply included. Thus a library of general and specific domain and task ontologies may be supported.

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. Figure 12 shows the primary techniques currently in use:

- **Domain ontology** techniques cited in the introduction elicit knowledge directly as an ordinal structure of concepts defined in terms of collections of roles and constraints. This is an efficient approach to the development of knowledge structures if the knowledge is already overtly available as it is in many technical domains. As shown above, the graphic representation of the subsumption, disjunction and rule relations between concepts, and the internal roles and constraints defining them, is a natural way to represent a knowledge structure, whether directly or indirectly elicited. Some form of graphic editing tool for presenting and editing domain ontologies is an essential component for an integrated knowledge acquisition tool.
- **Repertory grid** techniques cited in the introduction elicit knowledge indirectly by prompting the experts for critical cases and relevant constructs in a coherent sub-domain. The constructs define a domain concept that acts as a templet for the cases which are expressed as individuals recognized by the domain concept. Figure 13 shows how the grid relates to the conceptual structure of a domain, a grid is that structure which results from combining a concept with the individuals that it distinguishes. Thus, there is a grid corresponding to each domain concept or sub-domain. This corresponds to the application of grid tools in rapid prototyping to characterize a number of significant sub-domains in detail. The significance of grid techniques is that they provide a method of indirectly developing domain concepts through examples *when direct methods based on overt knowledge are not available.*
- **Empirical Induction** techniques such as C4.5 (Quinlan, 1987), Prism (Cendrowska, 1987) and Induct (Gaines, 1989) take a database of cases and derive constraints upon it that correspond to rules in knowledge representation schema. Each side of a rule defines a concept in terms of roles and constraints and these form a natural sub-domain concept structure as shown at the bottom of Figure 8. There are no fundamental differences between the application of induction techniques to a small database of correctly described critical cases expressed in terms of relevant attributes in a grid, and to a large database of lower quality data containing errors and irrelevant attributes. There is a strong trade-off between the quality of data and the amount of data required to induce the same quality of knowledge with knowledge acquisition tools emphasizing the elicitation of high-quality data, and empirical induction tools emphasizing the analysis of large datasets (Gaines, 1989).

The three techniques described come together naturally to give an integrated system as shown in Figure 14.

Stage ① consists of developing a basic domain ontology focusing on determining the major sub-domains rather than the detailed elicitation of the complete domain. Unstructured interviews, group synectics, text and protocol analysis techniques are appropriate to this stage in developing a broad and complete picture of the major features of a domain (Shaw & Woodward, 1989).

Stage ② consists of eliciting grids for each major sub-domain in order to characterize them in detail in terms of relevant constructs and critical cases. This enables the templet concept for the sub-domain to be developed together with cases for further analysis.

Stage ③ consists of using inductive analysis of the cases to derive constraints providing a conceptual structure and rules for the sub-domain. This can be evaluated by a sub-domain expert as an overt knowledge structure and in terms of its performance on additional cases, and edited and further developed as necessary.

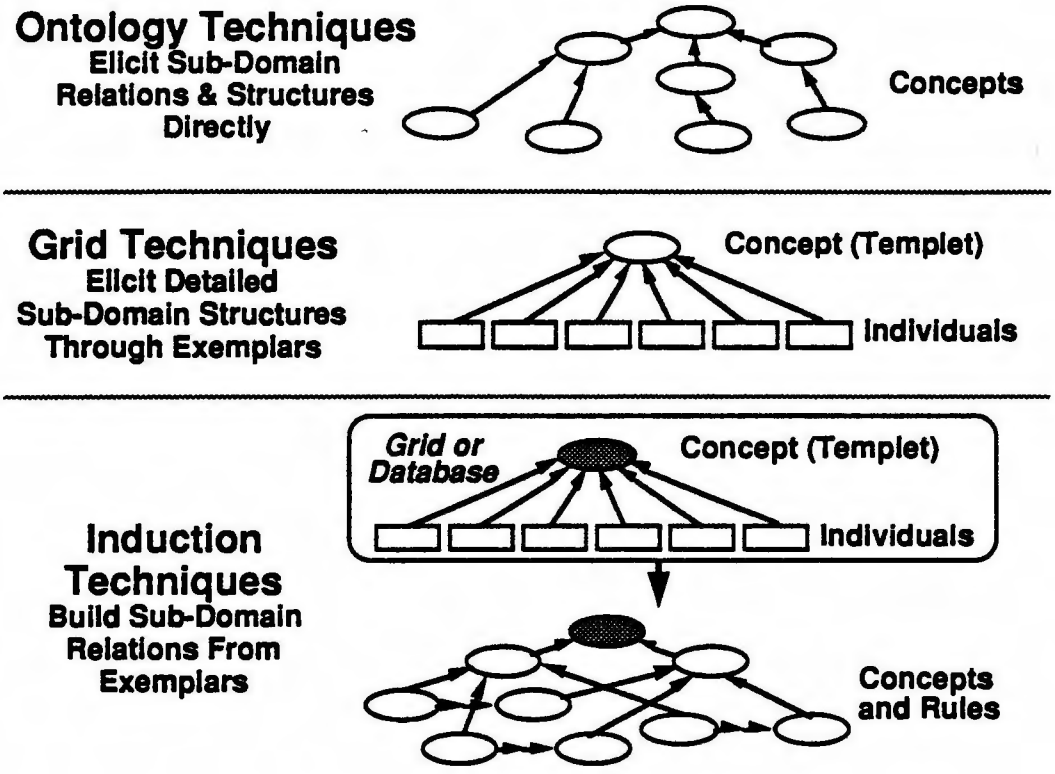


Fig.12 Knowledge acquisition techniques

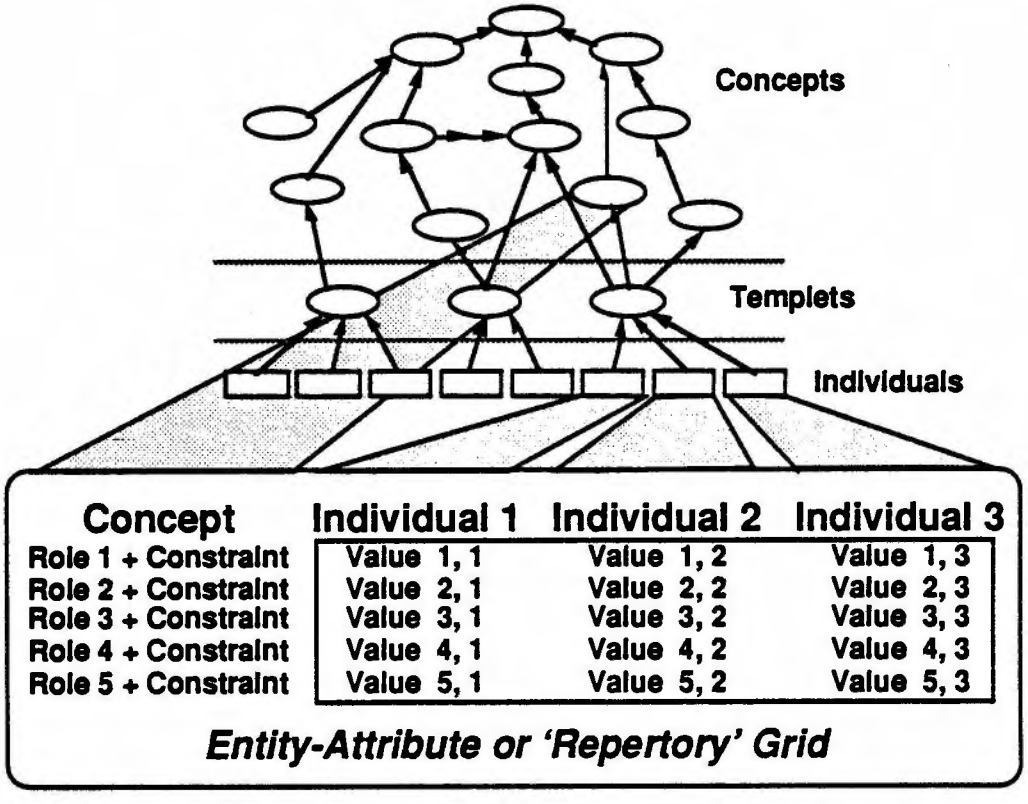


Fig.13 The repertory grid: individuals in the coordinate system of a concept

Stage ④ consists of combining the sub-domain knowledge structures to form an overall knowledge-based system for the domain. This will usually consist of linking the sub-domains through individuals in one sub-domain that are values of roles in others.

Stage ⑤ consists of testing the overall knowledge base and iterating back through any of the earlier stages in order to develop and refine it further. Note that the integrative, homogeneous framework for the overall process makes this a reasonable exercise.

This knowledge acquisition schema is an extrapolation and integration of existing, well-developed techniques for knowledge acquisition. It is particularly significant in providing means for the integration of acquisition and performance, supporting ongoing development of knowledge-based systems as part of their application. Each case for which the system is used provides new examples within the sub-domains that may be incorporated within the grids and used for further development. There is no intrinsic distinction between acquisition and performance.

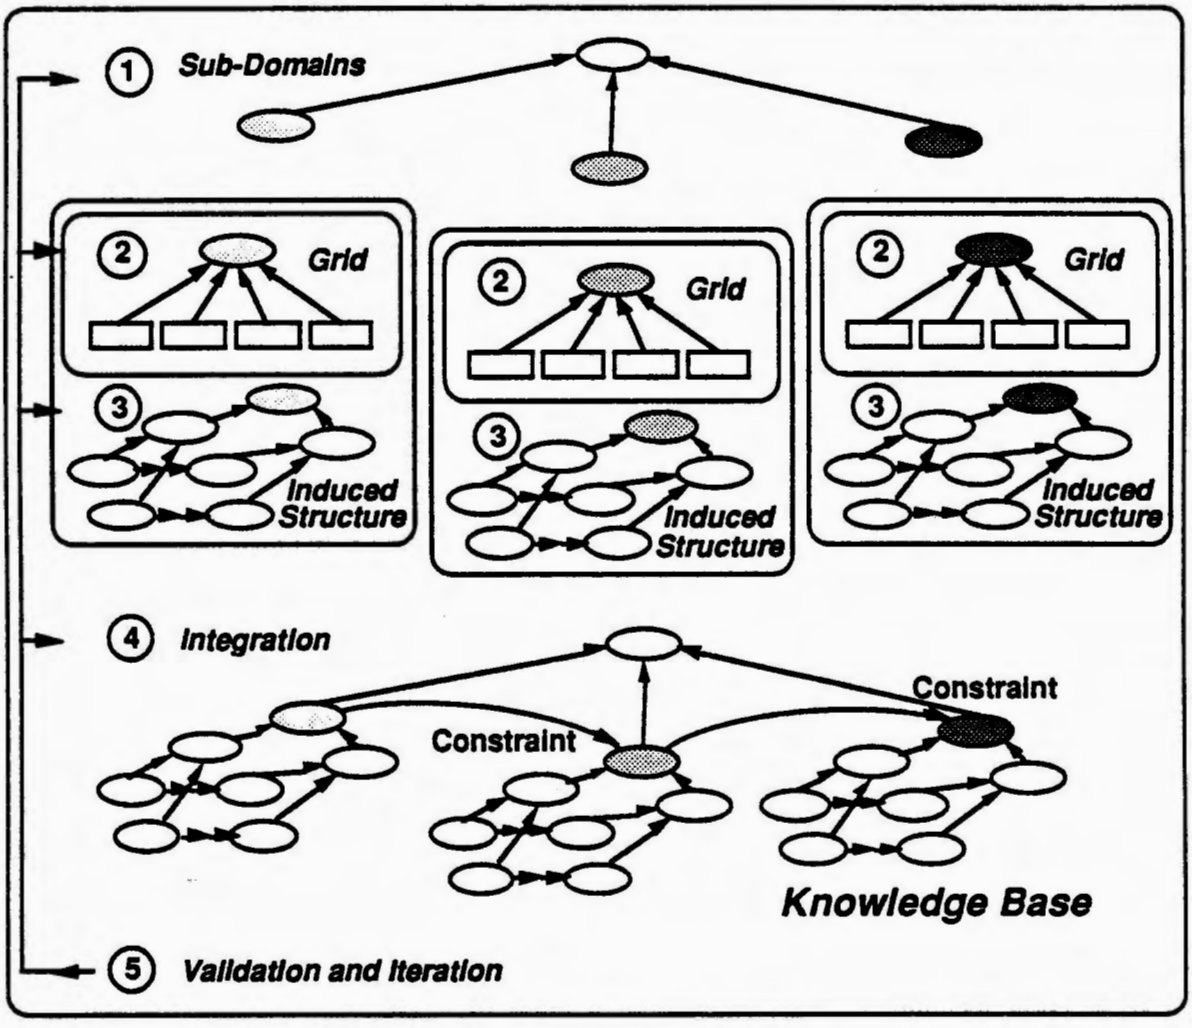


Fig.14 Integration of techniques

Other tools have a natural place in the systems of Figures 12 and 14. Conceptual clustering techniques such as FOCUS and PrinCom (Shaw & Gaines, 1987) are an alternative way of developing concepts by grouping related constructs together. They are attractive in allowing the expert to visualize conceptual relationships and see how they derive from the cases.

Techniques for comparing the knowledge structures of different experts, such as Socio (Shaw & Gaines, 1989), are important in highlighting conflicts in the use of distinctions and terminology and in enriching concepts by providing additional distinctions and alternative ways of expressing them. This analysis is related to Kelly's *commonality corollary*, that similar construction systems generate similar psychological processes.

The individual and attribute hierarchies introduced in AQUINAS (Boose & Bradshaw, 1987) are intrinsically available within the knowledge representation framework developed in this paper. They correspond to the extension of the grid in Figure 13 to encompass multiple related concepts.

Abstract individuals in the grid, such as an *ideal entity*, cannot be adequately represented as concrete individuals. They are basically a system of constraints not values and correspond to a concept rather than an individual. However, it is possible to provide a user interface for the definition of such abstract individuals in terms of constraints which is very similar to that for the entry of data about existent individuals in terms of values.

Such abstract individuals also provide the mechanism for decision-making based on pattern-matching as used in AQUINAS' inbuilt performance engine. Cases are generalized into concepts which can be used to recognize similar cases in future.

CONCLUSIONS

The main features of personal construct psychology have been presented. 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." Figure 15 recapitulates the main definitions derived from this primitive.

Distinction	Primitive notion
Subsumes	Asserting one distinction commits one to asserting another
Disjoint	Asserting one distinction commits one to not asserting another
Individual	Initial distinction that cannot be asserted of others
Concept	Mental entity imputed to agent to enable it to make a distinction
Construct	Triple of two disjoint concepts subsumed by a third *(dual* to concept)
Role	Internalization of a construct within a concept
Value	Concept in role applicable to an individual
Non-Primitive	Concept defined in terms of roles and applicable concepts
Recognize	Individual conforming to definition of a non-primitive concept
Rule	Commitment to assert an additional concept of a recognized individual

Fig.15 Main definitions in paper

The primitives derived in Figure 15 have been used to provide cognitive and logical foundations for existing knowledge acquisition tools and techniques, and for the design of integrated knowledge acquisition systems. A companion paper (Gaines, 1990) describes a practical implementation of such a system.

It is interesting to consider in conclusion what this paper leaves out, the *sources* of distinctions. Two paradigms not based on overt knowledge representation are both relevant. Situated action (Suchman, 1987; Clancey, 1989) explicates behavior in terms of a total systemic relation between the environment and the individual, including the inheritance of distinctions embedded in the built environment. *Neural networks* (Smolensky, 1986) provide a computational mechanism for creating new major distinctions at the symbolic level from what is virtually a continuum of distinctions at the sub-symbolic level. The framework presented in this paper has natural extensions to include these other paradigms and, potentially, give a complete account of human knowledge processes.

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, Bill Clancey, Gordon Pask, Alain Rappaport, Doug Skuce, Brian Woodward and other colleagues at the Knowledge Acquisition Workshops world-wide.

REFERENCES

Anglin, J.M. (1977) **Word, Object, and Conceptual Development.** Norton: New York.

Bannister, D. & Mair, J.M.M. (1968) **The Evaluation of Personal Constructs.** London: Academic Press.

Boose, J.H. (1984). Personal construct theory and the transfer of human expertise. **Proceedings AAAI-84**, 27-33. California: American Association for Artificial Intelligence.

Boose, J.H. (1986). **Expertise Transfer for Expert Systems.** Amsterdam: Elsevier.

Boose, J.H. & Bradshaw, J.M. (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.

Borgida, A., Brachman, R.J., McGuiness, D.L. & Resnick, L.A. (1989). CLASSIC: a structural data model for objects. Clifford, J., Lindsay, B. & Maier, D., Eds. **Proceedings of 1989 ACM SIGMOD International Conference on the Management of Data.** pp.58-67. New York: ACM Press.

Brachman, R.J. (1983). What IS-A is and isn't. **IEEE Computer, 16**(10) 30-36 (October).

Brachman, R.J. & Schmolze, J. (1985). An overview of the KL-ONE knowledge representation system. **Cognitive Science, 9**(2) 171-216 (April-June).

Bradshaw, J.M., Boose, J.H., Covington, S.P. & Russo, P.J., (1988) How to do with grids what people say you can't. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop.** pp.3-1-3-15. Banff (November).

Breuker, J. & Wielinga, B. (1987) Use of models in the interpretation of verbal data. Kidd, A. (Ed) **Knowledge Elicitation for Expert Systems: A Practical Handbook.** Plenum.

Brown, G.S. (1969) **Laws of Form.** London: George Allen & Unwin.

Bruck, R.H. (1971) **A Survey of Binary Systems.** Berlin: Springer-Verlag.

Cendrowska, J. (1987) An algorithm for inducing modular rules. **International Journal of Man-Machine Studies 27** (4), 349-370 (October).

Clancey, W.J. (1990) The frame of reference problem in the design of intelligent machines. van Lehn, K.& Newell A. (Eds.) **Architectures for Intelligence: The Twenty-Second Carnegie Symposium on Cognition.** Hillsdale: Lawrence Erlbaum.

Diederich, J., Ruhmann, I. & May, M. (1987). KRITON: A knowledge acquisition tool for expert systems. **International Journal of Man-Machine Studies, 26**(1), 29-40.

Eshelman, L., Ehret, D., McDermott, J. & Tan, M. (1987). MOLE: A tenacious knowledge acquisition tool. **International Journal of Man-Machine Studies, 26**(1), 41-54

Ford, K.M., Stahl, H., Adams-Webber, J., Novak, J. & Jones, J.C. (1990) An integrated constructivist knowledge acquisition tool. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Fifth AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop.** pp.7-1-7-20. Banff (November).

Frege, G. (1892). On sense and reference. Geach, P. & Black, M., Eds. **Translations from the Philosophical Writings of Gottlieb Frege.** pp.56-78. Oxford: Basil Blackwell, 1970.

Gaines, B.R. (1976). Foundations of fuzzy reasoning. **International Journal Man-Machine Studies, 8,** 623-668.

Gaines, B.R. (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. **Proceedings of 6th International Workshop on Machine Learning,** pp.156-159. San Mateo, California: Morgan Kaufmann (June).

Gaines, B.R. (1990) An architecture for integrated knowledge acquisition systems. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Fifth AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop.** pp.8-1-8-22. Banff (November).

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

Gaines, B.R. & Shaw, M.L.G. (1981). A programme for the development of a systems methodology of knowledge and action. Reckmeyer, W.J., Ed. **General Systems Research and Design: Precursors and Futures.** pp. 255-264. Society for General Systems Research.

Gaines, B.R. & Shaw, M.L.G. (1984). Hierarchies of distinctions as generators of system theories. Smith, A.W., Ed. **Proceedings of the Society for General Systems Research International Conference.** pp. 559-566. California: Intersystems.

Gaines, B.R. & Shaw, M.L.G. (1986). Induction of inference rules for expert systems. **Fuzzy Sets and Systems, 18**(3), 315-328 (April).

Gaines, B.R. & Shaw, M.L.G. (1987). Knowledge support systems. **ACM MCC-University Research Symposium.** Austin, Texas: MCC. 47-66.

Gracia, J.J.E. (1983) Individuals as instances. **Reviews of Metaphysics, 37**(1), 37-59 (September).

Garg-Janardan, C. & Salvendy, G. (1987). A conceptual framework for knowledge elicitation. **International Journal of Man-Machine Studies, 26**(4), 521-531.

Grossmann, R. (1983) **The Categorical Structure of the World.** Bloomington: Indiana University Press.

Hegel, W.F. (1812) **Science of Logic.** Johnston, W.H. & Struthers, L.G., Trans. (1929) London: George Allen & Unwin.

Hinkle, D.N. (1965). The change of personal constructs from the viewpoint of a theory of implications. PhD Thesis, Ohio State University.

Jones, W.P. (1989) Bringing corporate knowledge into focus with CAMEO. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Fourth AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop.** pp.17-1-17-23. Banff (October).

Kelly, G.A. (1955). **The Psychology of Personal Constructs.** Norton: New York.
Klinker, G., Bentolila, J., Genetet, S., Grimes, M. & McDermott, J. (1987) KNACK—report-driven knowledge acquisition. **International Journal of Man-Machine Studies, 26**(1), 65-79.
Lloyd, G.E.R. (1966) **Polarity and Analogy.** Cambridge: Cambridge University Press.
Mac Lane, S. (1971) **Categories for the Working Mathematician.** New York: Springer-Verlag
Mac Lane, S. (1986) **Mathematics Form and Function.** New York: Springer-Verlag
Marcus, R.B. (1962). Interpreting quantification. **Inquiry, 5**, 252-259.
Marcus, S. (1987). Taking backtracking with a grain of SALT. **International Journal of Man-Machine Studies, 26**(4), 383-398.
Maybury-Lewis, D. & Almagor, U. (1989). (Eds) **The Attraction of Opposites.** Ann Arbor: University of Michigan Press.
Meinong, A. (1910). **On Assumptions.** Berkeley: University of California Press, 1983.
Monarch, I. & Nirenburg, S. (1987) The role of ontology in concept acquisition for knowledge-based systems. **Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems (EKAW'87).** Reading University, UK (September).
Morik, K. (1987) Sloppy modeling. Morik, K. Ed. **Knowledge Representation and Organization in Machine Learning.** pp.107-134. Berlin: Springer-Verlag.
Motta, E., Eisenstadt, M., Pitman, K. & West, M. (1988) Support for knowledge acquisition in the knowledge engineer's assistant (KEATS). **Expert Systems 5**(1), 6-28.
Pask, G. (1975). **Conversation, Cognition and Learning,** Elsevier, Amsterdam.
Quinlan, J.R. (1987) Simplifying decision trees. **International Journal of Man-Machine Studies 27** (3), 221-234 (September).
Russell.B. (1905). On denoting. **Mind, 14**, 479-493 (October).
Sanford, D.H. (1975). Borderline logic. **American Philosophical Quarterly, 12**, 29-39.
Shaw, M.L.G. (1979). Conversational heuristics for enhancing personal understanding of the world. Gaines, B.R., Ed. **General Systems Research: A Science, A Methodology, A Technology.** pp. 270-277. Kentucky: Society for General Systems Research.
Shaw, M.L.G. (1980). **On Becoming A Personal Scientist: Interactive Computer Elicitation of Personal Models Of The World.** London: Academic Press.
Shaw, M.L.G. & Gaines, B.R. (1979). Externalizing the personal world: computer aids to epistemology. **Improving The Human Condition: Quality and Stability in Social Systems.** pp. 136-145. Kentucky: Society for General Systems Research.
Shaw, M.L.G. & Gaines, B.R. (1983). A computer aid to knowledge engineering. **Proceedings of British Computer Society Conference on Expert Systems,** 263-271 (December). Cambridge.
Shaw, M.L.G. & Gaines, B.R. (1987). KITTEN: Knowledge Initiation & Transfer Tools for Experts & Novices. **International Journal of Man-Machine Studies, 27**, 251-280.
Shaw, M.L.G. & Gaines, B.R. (1989) A methodology for recognizing conflict, correspondence, consensus and contrast in a knowledge acquisition system. **Knowledge Acquisition 1**(4), 341-363 (December).
Shaw, M.L.G. & Woodward, J.B. (1989) Mental models in the knowledge acquisition process. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Fourth AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop.** pp.29-1-29-24. Banff (October).
Skuce, D., Shenkang, W. & Beauville, Y. (1989) A generic Knowledge Acquisition Environment for conceptual and ontological analysis. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Fourth AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop.** pp.31-1-31-20. Banff (October).

Smolensky, P. (1986). Information processing indynamical systems: foundations of harmony theory. Rumerhart & D.E., McClelland, J.L., Eds. **Parallel Distributed Processing 1**. pp.194-263. Cambridge, Massachusetts: MIT Press.

Suchman, L.A. (1987). **Plans and Situated Action.** Cambridge: University Press.

Woods, J. (1974). **The Logic of Fiction.** The Hague: Mouton.

Woods, W.A. (1975). What's in a link: foundations for semantic networks. Bobrow, D.G. & Collins, A.M. **Representation and Understanding: Studies in Cognitive Science.** pp.35-82. New York: Academic Press.

Woodward, B. (1990) Knowledge engineering at the front-end: defining the domain. **Knowledge Acquisition 2**(1), 73-94 (March).

Zadeh, L.A. (1964) The concept of state in system theory. Mesarovic, M.D. (Ed.) **Views on General Systems Theory: Proceedings of the Second Systems Symposium at Case Institute of Technology**. New York: John Wiley.

Zadeh, L.A. (1965). Fuzzy sets. **Information and Control, 8**, 338-353.

Zalta, E.N. (1983). **Abstract Objects.** Dordrecht: Reidel.