Int. J. Man–Machine Studies (1990) **33**, 505–520

The meaning triangle as a tool for the acquisition of abstract, conceptual knowledge

STEPHEN REGOCZEI†

Computer Studies, Trent University, Peterborough, Ontario, Canada K9J 7B8

AND

GRAEME HIRST

Department of Computer Science, University of Toronto, Toronto, Ontario, Canada M5S 1A4

(Based on a paper presented at the AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop, Banff, November 1988)

The meaning triangle is presented as a useful diagramming tool for organizing knowledge in the informant-analyst interaction–based, natural language–mediated knowledge acquisition process. In concepts-oriented knowledge acquisition, the knowledge explication phase dominates. During the conceptual analysis process, it is helpful to separate verbal, conceptual, and referent entities. Diagramming these entities on an agent-centered meaning triangle clarifies for both informant and analyst the ontological structure that underlies the discourse and the creation of domains of discourses.

1. Introduction

In the construction of knowledge-based systems, even in cases where the domain of application is fairly concrete and straightforward, a considerable amount of abstract, conceptual knowledge must be acquired. This abstract knowledge is often hidden, tacit, commonsense and "too obvious to see"; hence informants and analysts have difficulty describing it. Often careful distinctions have to be made between verbal, conceptual, and external entitites to bring this knowledge to light. Extensions of the meaning triangle, introduced in this paper, are useful diagramming tools to accomplish this task.

In looking at knowledge acquisition as a *process*, Regoczei and Hirst (1990) divided it into three stages (Figure 1): elicitation, explication, and formalization. *Conceptual analysis* is the main technique in the explication of *any* knowledge, but it plays a special role when the knowledge is not about physical objects but rather about abstract entities, which are social constructs. The objective of the knowledge acquisition process is to create and update a *domain of discourse* that contains the knowledge. This domain of discourse is to be described, in a machine-usable form, in a notation such as conceptual graphs (Sowa, 1984; Regoczei & Plantinga, 1987; Regoczei & Hirst, 1990).

Knowledge acquisition can be viewed as *interaction*. We restrict our attention to the case where the domain expert, the informant, is interviewed by an analyst. The discussion between informant and analyst produces tape-recorded and transcribed

† Use Regoczei's address for correspondence.

0020-7373/90/050505 + 16$03.00/0

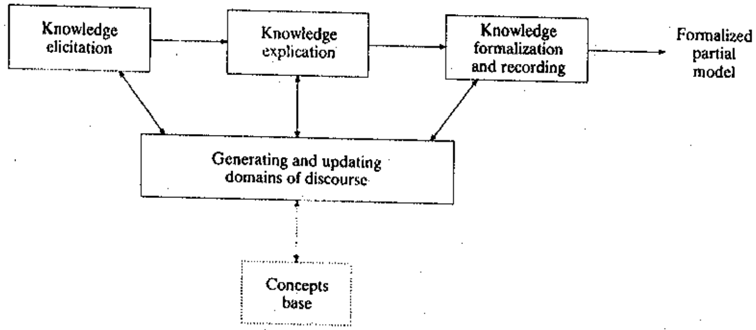

FIGURE 1. Processes in knowledge acquisition. The input to all processes is the mental models of the analyst and the informant. The output is a formalized *partial* model of domains of discourse.

text. This natural language text is the raw material from which the analyst reconstructs the domain knowledge and builds a formal, partial model of those portions of the expertise that are required for the construction of the knowledge-based system.

The informant-analyst–based, natural language–mediated knowledge acquisition process can also be regarded as a process of *mutual attribution*. Adopting the *mental models hypothesis* (Sowa, 1984; Regoczei & Plantinga, 1987; Regoczei & Hirst, 1990), i.e. that people understand the world by forming mental models, we can picture the task of the analyst as building mental models of the informant's mental models. More precisely, the analyst re-creates the informant's knowledge by building mental models that the analyst then "attributes" to the informant, and labels as the model of the mental models of the informant. The correctness of the labeling, and the verisimilitude of the mental models so labeled, is tested by the analyst, who expresses his views, both in terms of colloquial natural language and in more formal terms such as logic and diagrams, and asks for feedback from the informant. Corrective measures are taken as required. In effect, the informant and analyst *negotiate* what counts as (relevant) knowledge. The roles of the informant and analyst are symmetric. Each is building mental models of the other's mental models, and *harmonizing* them through interaction mediated by natural language.

In brief, the analyst does not merely "acquire" something that already exists. Rather, during the knowledge acquisition process, the informant and the analyst cooperatively create a domain of discourse—or perhaps several *distinct* domains—by matching concepts to words.

2. The meaning triangle

The meaning triangle (Figure 2) was prominently featured by Ogden and Richards in their book *The Meaning of Meaning* (1923). It has also been used by Sowa (1984) to motivate his conceptual graphs notation, by Martin (1975) for the analysis of literary texts, and by Umberto Eco (1976) for semiotic analysis. It was introduced in knowledge acquisition by Regoczei and Hirst (1990).

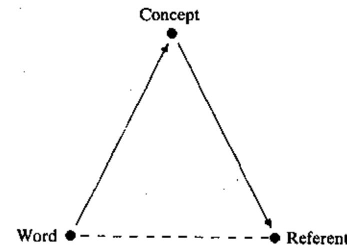

FIGURE 2. The meaning triangle.

The main function of the meaning triangle in knowledge acquisition is to help draw distinctions between *words, concepts,* and *referents*.

- Words, or lexical objects in general, are concrete linguistic entities embodied in speech or in written text. They are pictured as located on the word node of the meaning triangle.
- Concepts are abstract, intangible entities, usually associated with human thinking. There is considerable controversy surrounding the existence and nature of concepts. In this paper, as well as in knowledge acquisition work in general, one takes a rather pragmatic approach to concepts, considering them to be constructs that are used to build mental models of the world. They are tools to aid thinking. Concepts as cognitive constructs are pictured as being located on the concept node of the meaning triangle.
- The referent node is the place for entities such as physical objects that are pictured as being part of the external world.

3. Abstract, conceptual knowledge

The referent node seems simple, but is actually highly problematic. The entities that populate the referent node are usually construed as "objects out there"—hence physical objects are the prototypical example. These are "objects out there" because they are independent of what any person thinks or says. The qualifier "out there" is to capture the intuition that words can be uttered, and thought can be generated, but no amount of thinking or talking is going to bring a table or cat or chair into being.

While this intuition is very strong and clear, and exercises a powerful influence upon our thinking, it is not too difficult to blur the edges. There are socially constructed abstract entities, especially in the fields of law, culture, interpersonal interactions, organizations, and mathematical sciences, which *are* brought about through talking and thinking by individuals within a group. To a single individual, these abstractions may appear as given as the earth and the sky. Yet they are created, and hence may be changed, by people acting together. But the process is not wholly arbitrary; it takes place within the constraints and limitations of the physical and social world. The "software" of such constructs often requires a "hardware" component. For example, armies and corporations are abstract entities with heavy "hardware" components. Although they are abstract, they are much more bound by the physical world, and are much more concrete, than, let us say, the constructs of mathematics. They can be thought of as *hybrid entities*, with lexical, conceptual, and physical components.

In brief, if private, cognitive concepts are to be diagrammed on the meaning triangle by placing them on the concept node, then public, abstract concepts may be placed, with good justification, on the referent node. Acquiring knowledge about these abstract entities, which are essentially social constructs, is the challenging task that we are addressing in this paper. The task is made harder by noting that information or clues about these abstract entities is available mostly through the mental models of individuals and what these individuals utter about their mental models. We may call this difficulty the *perspectivist paradox*.

4. An example: "five chairs"

Esoteric knowledge may have its own intrinsic fascination, but for the building of expert systems, knowledge-based software, or conceptual schemas for databases, the required knowledge is usually quite mundane in nature. Yet even in these simple cases of knowledge acquisition, conceptual knowledge may play a central role.

To take a very basic and simple example, let us assume that we are building a domestic robot with housekeeping expertise. When we interview a resident of the house the robot is to clean, he may utter a statement such as this:

(S1) "There are five chairs in the room."

There are two different types of knowledge being alluded to in this statement. The first type may be characterized as *inventory-like* (Regoczei & Plantinga, 1987; Regoczei & Hirst, 1990). The claim is made that there are five chairs in the room, not four or eight. This kind of propositional, declarative knowledge is typically what is acquired in knowledge acquisition that is oriented towards rules, facts, and propositions. Such knowledge is often acquired only at the verbal level using what is nothing more than uninterpreted character strings.

But there is an even more important kind of knowledge being alluded to by statement S1, and this is pertaining to what exists—*ontological* knowledge (Regoczei & Plantinga, 1987; Regoczei & Hirst, 1990). What is being said, by convention, generates the presupposition, or the claim, that certain things exist. The informant is talking about chairs, and a room, and the concept of containment. He is talking about numbers, in particular the number five, the ability to count objects, and the notion of grouping objects together to be counted.

This gives us a large number of concepts, which we can list using Sowa's conceptual graphs notation for concepts, and we can diagram the proposition "behind" the statement as follows:

[CHAIR: @5]→(LOCATION)→[ROOM: #].

In diagramming this analysis on a meaning triangle (Figure 3), we place statement S1 on the word node, and the various concepts contained in the ontology are located on the concept node. If there is an actual, physical room that is being talked about, and there are actual physical chairs in it, then what is diagrammed on the referent node is the information that pertains to the state of the object domain "out there". The domain of discourse on the concept node should contain not only the general or generic concepts [CHAIR] and [ROOM], but also the specific concepts that refer directly to the particular state-of-affairs within the object domain, i.e. the action defined by the concept [COUNT], as applied to a particular [GROUP-OF-OBJECTS: #], yielding [CHAIR: @5] in a particular room specified by the concept [ROOM: #].

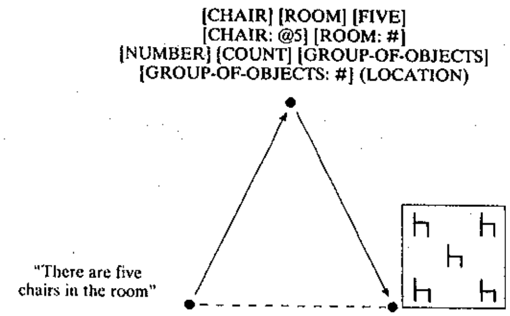

FIGURE 3. Meaning triangle for statement S1.

5. A molecular theory of knowledge

We have been diagramming knowledge on the concept node, not as incoherent, amorphous clouds, but as hard-edged, individual, discrete concepts. This practice needs some explanation, since when people think and talk naturally, they do not "step" carefully from one thought to the next. Their thoughts have a tendency to flow into each other, and be forever changing, like an anthill (Regoczei & Plantinga, 1987).

In the analysis above, we tacitly assumed that knowledge can be "chopped up" into units that are conceptual in nature, that these units can be named in appropriately mnemonic ways, and that these units, recombined into conceptual structures, capture the knowledge required. This assumption can be described as the Molecular Theory of Knowledge.

We certainly do not assume that all knowledge is molecular in nature. On the contrary, much of cognitive activity seems to be sub-conceptual. Nevertheless, we state that in order to build software, we have to find conceptual models for expertise and world knowledge, because, given the state of today's technology, only conceptual models of knowledge are tractable computationally. In discretizing knowledge, we make it easier to handle the mapping between text and concepts, while at the same time avoiding certain controversial theories that claim that all thinking and knowledge is language-based.

The representation of knowledge could be effected in any of several different notations. One would prefer those that emphasize concepts as the basic semantic primitives, rather than logical connectives, predicates, or functions. For this reason,

Item	Notation
word	"fish"
generic concept	[FISH]
concept, pointing to a referent	[FISH: #]
referent	Informal English description, diagram, or a specific concept name such as [FISH: #2356] as a surrogate

FIGURE 4. Summary of conceptual graph notation.

we adopted the notation of Sowa's conceptual graphs (1984). Some of this notation we have already used informally. We summarize the notational conventions of conceptual graphs in Figure 4.

6. Extensions of the meaning triangle

While it is intuitively suggestive, the traditional meaning triangle does not stand up to close scrutiny. The issues that have to be handled in real-life knowledge acquisition interactions between analyst and informant force the extension of the traditional meaning triangle.

6.1. THE AGENT-CENTERED MEANING TRIANGLE

In its traditional form, the meaning triangle provides no guidance as to where to place a public concept. Much of the hesitation in using the meaning triangle for purposes other than literary text analysis came from this shortcoming. The problem,

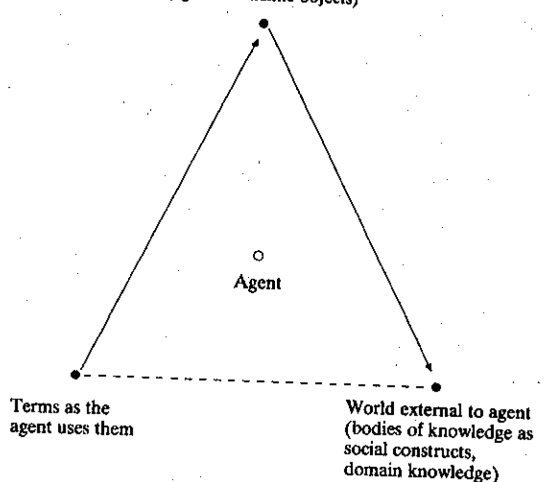

FIGURE 5. An agent-centered meaning triangle.

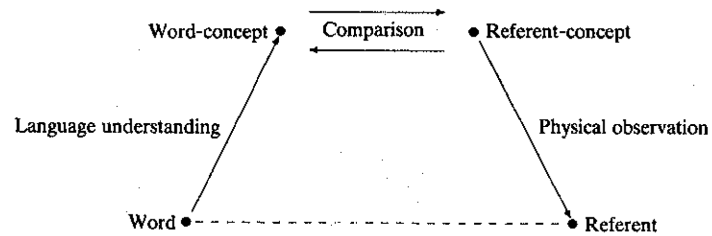

FIGURE 6. The meaning trapezium.

already indicated above, is this: because it is a concept, an abstract public concept should be placed on the concept node. But because a public concept is also a social construct "out there", it should be placed on the referent node. To resolve this deadlock, we introduce the *agent-centered meaning triangle*. The agent-centered meaning triangle has an extra node for specifying the agent relative to whom a given entity is to be classified as verbal, conceptual, or referent in nature (Figure 5).

6.2. THE MEANING TRAPEZIUM

The concept node can be split to accommodate the distinction between forming a concept on the basis of hearing some words, and forming a concept on the basis of observing something in the object domain (Martin, 1975). Splitting the concept node in two (Figure 6) creates the meaning trapezium. All comparisons between a generic concept such as [CAT] and a specific concept such as [CAT: #] or [CAT: #235] take place at the conceptual level. Although we do not claim that there is a sortal distinction between these concepts, drawing such distinctions does enable us to make precise the difference between "speaking in general" and "talking about a specific thing".

7. A frame-like notation for the meaning triangle

Rather than drawing pictures, one can define a frame-like description for meaning triangles. Figure 7 shows a schema for the notation. This notation is especially useful for producing modifications and alternative versions of the same meaning triangle. Changes may be noted without repeating the full information.

Meaning triangle:
- **Name or description:** [*name*]
- **Type:** [*e.g., agent-centered*]
- **Word-node:** [*a description of the entities that populate the word node*]
- **Concept-node:** [*a description of the entities that populate the concept node*]
- **Referent-node:** [*a description of the entities that populate the referent node*]
- **Agent-node:** [*a description of the agent, if applicable*]

FIGURE 7. Frame-like notation for the meaning triangle.

8. An example: the Riemann integral

Abstract, conceptual knowledge is encountered both at the commonsense, everyday level as discussed above, and at the theoretical level in working with abstract theories. We now turn our attention to abstract, public bodies of knowledge, concentrating on the perspectivist paradox stated above, i.e. that information or clues about public concepts are available mostly through the private mental models of individuals.

Mathematics provides the richest area for examples in the acquisition of knowledge about abstract, conceptual entities. The conceptual analysis required to explicate this knowledge is usually difficult. Ontological issues dominate the informant–analyst interaction. One has to able to separate mathematical "reality" out there, usually construed as "platonic entities" or "mathematical truth" out there, from what individuals think about these mathematical entities and theories. Mathematicians are careful to avoid psychologism or mentalism, by objectifying "mathematical truth", picturing it as "eternal", "immutable", and independent of the opinions of mere mortals. Nevertheless, the knowledge acquisition analyst is dealing with a real, live informant, and has nothing more to go by than what this informant thinks and says. These issues are central to the acquisition of abstract, conceptual knowledge in general, and the acquisition of mathematical knowledge is a good test case to study, but, not surprisingly, these issues are seldom discussable within the discourse of mathematics. That is why we so often find people complaining that mathematicians do not know how to explain things.

We shall use the example of an analyst trying to find out about the Riemann integral from a mathematician. On an agent-centered meaning triangle, we can diagram the discourse with the frame-like representation, as shown in Figure 8.

The [RIEMANN-INTEGRAL] is the agent's mental model—possibly consisting of a large, structured *cluster* of concepts—of the public concept of the Riemann integral. In philosophical disputes, the existence of the public concept of the Riemann integral may be called into question. Since in this paper we assume the position of the "reflective practitioner" (Schon, 1983), we have no hesitation in declaring that mathematical knowledge (like medical knowledge, legal knowledge, and so on) exists as a social construct independent of the idiosyncrasies of individuals (but not as a platonic entity). This is the public knowledge that is located on the referent node. The private concept of [RIEMANN-INTEGRAL], however, is relativized to the agent.

Meaning triangle:
- **Name or description:** Meaning triangle for a mathematician trying to talk about the Riemann integral.
- **Type:** Agent-centered.
- **Word-node:** Statements to analyst.
- **Concept-node:** The mathematician's private mental models under the concept label [RIEMANN-INTEGRAL].
- **Referent-node:** The Riemann integral, a public, abstract construct, created as a social construct by the community of mathematicians.
- **Agent-node:** Mathematician acting as the informant.

FIGURE 8. Agent-centered meaning triangle for mathematician.

During the knowledge elicitation and explication stage, the agent may be uttering statements either about the mathematical knowledge "out there" or about his own personal view, or picture, of this mathematical knowledge. This means that a further distinction will have to be drawn. Some statements may be about [RIEMANN-INTEGRAL] on the concept node, and some about the public, abstract concept of the Riemann intergral on the referent node.

This generates a paradox. The public concept of the Riemann integral is accessible to the agent only insofar as he is able to form a mental model of this public concept. So, actually, there may be two domains of discourse, simultaneously located on the concept node, one identified as the "surrogate" domain which "faithfully" reflects the public concept, and the other which corresponds to the more private opinions held by the mathematician. We can update the meaning triangle of Figure 8 as shown in Figure 9.

There is often no effort on the part of the informant to keep separate his private domain of discourse from what we could label as his "surrogate" domain. In fact, he may often deliberately blur the distinction, if he is aware of it at all, in order to get his own personal views on record. Similar blurring may occur in cases when the expert's own way of looking at things differs from the "official" way of doing something. The process may work in reverse in the case of medical diagnosis, where the doctor will quote the medical textbook rather than reveal his own "informal" but quite effective techniques.

Separating and clearly labeling the knowledge about the private domain of discourse and the knowledge about the "surrogate" domain is the responsibility of the analyst. What the informant says, in his own words, usually contains enough clues for the analyst to capture the required distinctions. For example, the mathematician might say:

> "Well, let me tell you how I think about the Riemann integral. You won't find this in books. The books describe things differently. It always bothered me that I couldn't immediately see the difference between the Lebesgue integral and the Riemann integral..."

From this text it is clear that we need two domains of discourse, the first containing the mathematician's private views, and the second his accurate reflection of the "official" story. In the acquisition of expertise, both domains are of interest.

Meaning triangle:
- **Name or description:** Meaning triangle for a mathematician trying to talk about the Riemann integral; Version 2.
- ....
- **Concept-node:**
  - Domain of discourse number 1:
    - Private, containing the concept [RIEMANN-INTEGRAL].
  - Domain of discourse number 2:
    - Private, labeled as the "surrogate" domain, containing mental models of the public entity on the referent node.
  - ....

FIGURE 9. Updated meaning triangle for mathematician.

The informant may continue describing his mental models by saying, for example:

> "For the Riemann integral, the partitions on the domain are too crude, it's like a giant with big fingers not being able to pick up the small pieces, or a net with thick strands and large holes that lets the plankton and small fish slip through as if they were not there.
>
> Point sets with non-zero measure just slip through the fingers... just slip through the net in little bits and pieces.... I also think of the ordinate set as the whale's teeth, straining the plankton, getting something out of a lot of little nothings."

This is not exactly conventional, mathematical prose. In fact, a mathematician may feel rather uneasy talking this way in public. Although the imagery is concrete, the essential concepts "behind" the imagery are abstract. Concepts such as [POINT-SET], [MEASURE], [NON-ZERO-MEASURE], and [ORDINATE-SET] are standard mathematical concepts, but the [RESOLUTION-POWER] concept that the informant implicitly alludes to is not. To capture the informant's private domain of discourse, in addition to the "official" story, one has to accept concepts that have no direct referent in public, abstract, mathematical knowledge.

9. A meaning triangle of domains

Perhaps the most important extension of the meaning triangle, as already indicated above, is formed by placing not only single entities, but entire domains on the vertices. An agent-centered version of such a meaning triangle is shown in Figure 10. The domains are not necessarily formed by the agent himself. The domains may be abstract entities created by an observer who watches the process and diagrams it in the form of a meaning triangle. Thus the various domains—domains of discourse, "surrogate" domain, object domain—may all be creations at a higher level of abstraction. It is at this level that one may draw taxonomic distinctions between processes, events, qualities, attributes, and so on, which are crucial for a smooth flow of the interaction between analyst and informant (Regoczei & Hirst, 1990). At times, the expert seems to be adept at manipulating inconsistent and conflicting knowledge. The contradictions implicit in his statements can often be resolved by drawing distinctions between concepts, and through the creation of separate domains of discourse. For an especially vivid example to illustrate this point, we will use Eddington's paradoxical story of his "two tables" below.

In Figure 10, the various domains of discourse are diagrammed on the concept node. These concepts and conceptual structures are the constructs of the agent. These domains of discourse may contain concepts of possible events, processes, and so on. One or more domains carry special labels as "surrogate" domains reflecting "reality" as it "really" exists "out there".

The referent node contains the object domain that includes both physical objects and abstract entities as social constructs. One may include events and processes that take place in the object domain. What characterizes these is that they unfold in real time. In fact, the referent node is dominated not only by physical space, but also by real physical time.

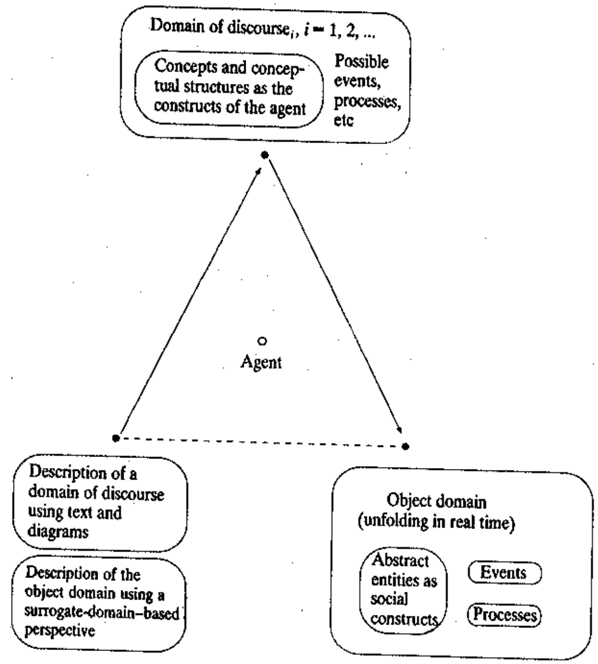

FIGURE 10. Meaning triangle of domains.

The population of the word node provides interesting insight. Essentially, there are two classes of statements:

- descriptions of domains of discourse, using text and diagrams; and
- descriptions of the object domain based upon certain specific perspectives given by the "surrogate" domain.

In ordinary, everyday use of language, we do not distinguish whether we are talking about the population of the concept node or the population of the referent node. Part of the power of natural language derives from this systematic disregard of typing rules through the use of *deliberate category mistakes*. In certain cases, however, during the knowledge acquisition process, one may want to specify, using the distinction on the word node, whether one is talking about concepts in general or concepts directly interpretable within the object domain.

10. Uses of the meaning triangle

As discussed above, the drawing of distinctions between categories of existence such as lexical, conceptual, external referent-like, and agent-like is crucial in the informant–analyst interaction. Extensions of the meaning triangle, introduced in this paper, are helpful in documenting these distinctions. "What are we talking about?" is the basic question that underlies the discussion of the analyst and informant as they try to understand each other. They have to draw distinctions to sort out contradictions and terminological confusion. They have to handle misunderstandings created by semantically ill-formed dialogue, create new conceptual entities as needed, effect appropriate *concept cluster attachments*, and harmonize their mental models.

Detailed examples of the use of the meaning triangle in systems analysis and requirements specification are discussed in *Knowledge acquisition and requirements specifications: uses of the meaning triangle* (Regoczei & Hirst, forthcoming). In this and the next section, we briefly outline some examples from this paper to illustrate the key issues.

10.1. SORTAL NON-SPECIFICITY

We consider the meaning triangle as a diagramming technique on par with structure charts or dataflow diagrams. In talking about an entity X being "on" a node or vertex of the meaning triangle, we are using a spatial metaphor. If X is on vertex A then it cannot be on vertex B. This forces us to "sort" X, determine its ontological status, label it, and place it (by diagramming) on one and only one of the vertices.

During the acquisition interaction, one often encounters sortal vagueness and non-specificity. People insist that X is both a lexical object and a non-lexical object. This category mistake can be handled by using the meaning triangle to construct homomorphic images, linked by *coreference links* and located on different vertices. As an example, consider the well-known safety slogan:

> Accident is just a word until it happens.

This mixes an event (reference type) with the word "accident" (a word type) with people not having a proper conceptual grasp of the notion of an accident (concept type). We may diagram "accident", [ACCIDENT], and the event-of-an-accident or the public-concept-of-an-accident on the word, concept, and referent nodes respectively.

10.2. DISTINGUISHING CONCEPTS AND WORDS

The knowledge acquisition predicament (Regoczei & Hirst, 1990) is this: the analyst is after knowledge, but all he gets is words. Generating concepts under the input stimulus of words is a laborious process that is often not executed. Words are just accepted and used as empty tokens, with no thought to the concepts "behind" the words.

Many existing expert systems contain knowledge only at the verbal level. MYCIN (Shortliffe, 1976), for example, deduces character strings from other character strings. Much of the knowledge that is written in logical notation has no concepts behind it.

Separating concepts and words enables one to implement systems with a two-level architecture: verbal and conceptual. Concepts are implementable on a computer as abstract data types.

10.3. IS THERE A REFERENT?

As the informant is describing his mental models, the most powerful question that the analyst can ask is this: "Is there a referent?" If the informant doesn't understand

the question, the analyst can use the meaning triangle to explicate the distinction between the verbal, the conceptual, and the referent. If there is no referent in the object domain—either as concrete, physical objects, or abstract, conceptual entities—the informant is describing his own, idiosyncratic mental models. His description should not be rejected. Part of his expertise may derive from such private knowledge. But it should be clearly labeled as such, and perhaps contrasted with more "official" versions (section 8).

11. Drawing distinctions

At times, experts seem to be operating with seemingly inconsistent and contradictory knowledge. The contradictions can often be resolved by drawing appropriate distinctions, readjusting concept cluster attachments, and creating new entities through the construction of new concepts.

11.1. SORTALLY COMPOSITE ENTITIES

Eddington, in the introduction to his 1927 Gifford lectures (Eddington, 1928), adopted the startling strategy of talking about his "two tables":

> I have settled down to the task of writing these lectures and have drawn up my chairs to my two tables. Two tables! Yes; there are duplicates of every object about me—two tables, two chairs, two pens. (p. xi)

One table is the familiar, commonplace object of naive physics that everyone is, or should be, familiar with: it has dimension, colour, substance, and some degree of permanence. The other table is the table of "sophisticated" physics, consisting mostly of empty space between nuclei of atoms and electrons moving about with great speed. Yet, despite this strange composition, this scientific table was able to support Eddington's elbows and his writing paper. The puzzling feature is that there is only one referent object, but two conceptual worlds. Is Eddington talking about one "real" table, or two "real" tables?

To resolve the inconsistency, we have to conclude that in talking about "two tables", Eddington is not talking about the table pictured on the referent node. Rather, he is talking about composite entities consisting of a referent physical object and a concept cluster attached to it. This composite entity we could describe as an ordered pair:

⟨physical referent object, concept cluster⟩

Attaching concepts to words and referents creates composite—sortally composite—entities. In many cases, the use of knowledge, the use of language, and the manipulation of objects in the physical world involve composite entities.

A common misunderstanding is this: the expert seems to be talking about the referent object, but in fact he is talking about *several, different* sortally composite entities, with the same first component.

11.2. ATTACHING CONCEPT CLUSTERS

As an example of knowledge acquisition for an expert system for identifying chemical spills and their sources (Buchanan, Barstow, Bechtal, Bennett, Clancey, Kulikowski, Mitchell & Waterman, 1983: 134), the following dialogue fragment is reported:

Knowledge Engineer: Suppose you were told that a spill had been detected in White Oak Creek one mile before it enters White Oak Lake. What would you do to contain the spill?

According to the report, the obliging expert proceeds to consider and describe how to find the source and determine what the spilled material is. But let's imagine a more interesting turn in the conversation, one that a skeptical, more battle-scarred expert might take:

Expert: Well, wait a minute now. Before I "contain the spill" I have to verify that the report is accurate. Just because somebody claims that there's a spill, that's not going to make me jump to conclusions. How does he know? Do I believe him? He detected some foreign substance. Fine. Then that's what he should report.

KE: But suppose it really is a spill. He sees it pouring out of a drum or a tank.

E: I would ask him to turn off the tap. To find out what the substance is, he should check the label on the drum to see what it contains. It might be lemonade.

KE: That doesn't take much expertise.

E: Right. The trick is to figure out what people mean when they say things. People attach all sorts of strange concepts to what they see. As Whorf pointed out some time ago, empty drums can be quite dangerous.

The skeptical expert has a point. The trick is to associate with the empirical observation a concept cluster centered on the concept [SPILL]. The expert is drawing a distinction between a spill (a physical referent, "out there") and the verbal report of a spill. He asks how the person making the report attached concepts to what he experienced or observed (right leg of the meaning trapezium). He is also reluctant to immediately attach the general concept [SPILL] to a verbal report containing the word "spill" without receiving further information (left leg of the meaning trapezium). This point can be made even more strongly by recalling recent events in the Persian Gulf involving the USS *Stark* and the USS *Vincennes*, where in both cases the labels of "friend" or "foe" were wrongly attached to external referents.

11.3. CREATING NEW ENTITIES BY CONSTRUCTING CONCEPTS

Hybrid entities are composed of public concepts, text and hardware. We can create new hybrid entities by constructing concepts, turning them into social constructs, and projecting them onto the physical part of the referent object domain. In the five-chairs example, a group of chairs was created by attaching a concept [GROUP-OF-OBJECTS] to the chairs already in the room. We grouped the chairs by attaching a concept. It is this concept that gave added structure to the contents of the room. We did not create chairs as physical objects.

Creating corporations, academic disciplines, nations, states, and countries as hybrid entities is done by constructing the appropriate concepts and projecting them onto the object domain. Creating sortally composite entities by concept cluster attachment was already discussed in section 11.1.

12. Conclusion

Informant-analyst–based, natural language–mediated knowledge acquisition interaction is facilitated by the use of diagramming techniques such as the meaning triangle that are acceptable to both informant and analyst. These techniques help the largely human-based acquisition process.

It is reasonable to think that future developments of computer-based techniques would be founded on a better understanding of people-based techniques. Thus the use of the meaning triangle described in this paper is not only to help build better knowledge-based systems now, but also to chart a migration path for computer-based knowledge acquisition. Focusing on concepts, rather than propositions, facts, or rules enables meaning triangle techniques to be intergrated with other techniques using personal constructs, the repertory grid, and protocol analysis (Boose, 1984; Gaines, 1986; Shaw & Gaines, 1986).

13. Epilogue

After reading this paper, people may still ask (trying to do knowledge acquisition of sorts): "This is great, and it reads well, but what *is* a concept?" We may demonstrate the techniques outlined in this paper by trying to handle this question.

First, the word "concept" may be looked up in the *Oxford English Dictionary*, because that is how one finds out about words.

Second, the concept of concept, the reader either grasps or does not grasp. Concepts have to be created by the agent; they cannot simply be handed over by another.

Third, the concept as referent is tricky. The authors of this paper testify that they have a fairly clear notion of the concept of concept. In fact, they have quite a few different conceptions of the various concepts of concept. These concepts are on the referent node relative to the reader. As for a public, abstract concept of concept, it is one of the most basic, underlying concepts of Western thought.

This is how a reflective practitioner would handle the question. It is this approach, demonstrated in this epilogue, that helps the reflective practitioner avoid the paralysis generated by sterile disputes unconstrained by the operational necessity of creating working software.

We are grateful to John Burbidge, Roy Hagman, Ed Plantinga and Brian Woodward for discussions, and to an anonymous referee for his or her suggestions. This work was supported by a grant to the first author from the Trent University Committee on Research, and to the second author from the Natural Sciences and Engineering Research Council of Canada.

References

BOOSE, J. H. (1984). Personal construct theory and the transfer of human expertise. *Proceedings of the National Conference on Artificial Intelligence, AAAI-84*, Austin, August 1984, 27–33.

BUCHANAN, B. G., BARSTOW, D., BECHTAL, R., BENNETT, J., CLANCEY, W., KULIKOWSKI, C., MITCHELL, T. & WATERMAN, D. A. (1983). Constructing an expert system. In F. Hayes-Roth, D. A. Waterman & D. B. Lenat, Eds. *Building Expert Systems*, pp. 127–167. (Teknowledge series in knowledge engineering). Reading, MA: Addison-Wesley.

ECO, U. (1976). *A Theory of Semiotics*. Bloomington, IN: Indiana University Press.

EDDINGTON, A. S. (1928). *The Nature of the Physical World*. Cambridge: Cambridge University Press.

GAINES, B. R. (1986). An overview of knowledge acquisition and transfer. *Proceedings of the Knowledge Acquisition for Knowledge-Based Systems Workshop*, Banff, November 1986.

MARTIN, G. D. (1975). *Language, Truth, and Poetry: Notes Towards a Philosophy of Literature*. Edinburgh: Edinburgh University Press.

OGDEN, C. K. & RICHARDS I. A. (1923). *The Meaning of Meaning* (8th edition, 1946). New York: Harcourt, Brace and World.

REGOCZEI, S. & HIRST, G. (1990). Knowledge acquisition as knowledge explication by conceptual analysis. In: J. F. Sowa, N. Foo, & A. Rao, Eds. *Conceptual Graphs for Knowledge Systems*, [to appear].

REGOCZEI, S. & PLANTINGA E. P. O. (1987). Creating the domain of discourse: ontology and inventory. *International Journal of Man–Machine Studies*, **27**, 235–250.

SCHON, D. A. (1983). *The Reflective Practitioner: How Professionals Think in Action*. New York: Basic Books.

SHAW, M. L. G. & GAINES, B. R. (1986). Techniques for knowledge acquisition and transfer. *Proceedings of the Knowledge Acquisition for Knowledge-Based Systems Workshop*, Banff, November 1986.

SHORTLIFFE, E. H. (1976). *Computer-Based Medical Consultations: MYCIN* (Elsevier Computer Science Library; Artificial Intelligence Series 2. New York: Elsevier.

SOWA, J. F. (1984). *Conceptual Structures: Information Processing in Mind and Machine*. Reading, MA: Addison-Wesley.