

KIND TYPES IN KNOWLEDGE REPRESENTATION

K. Dahlgren
IBM Los Angeles Scientific Center
11601 Wilshire Bl.
Los Angeles, California 90025

J. McDowell
Department of Linguistics
University of Southern California
Los Angeles, California 90089

Abstract This paper describes Kind Types (KT), a system which uses commonsense knowledge to reason about natural language text. KT encodes some of the knowledge underlying natural language understanding, including category distinctions and descriptions differentiating real-world objects, states and events. It embeds an ontology reflecting the ordinary person's top-level cognitive model of real-world distinctions and a database of prototype descriptions of real-world entities. KT is transportable, empirically-based and constrained for efficient reasoning in ways similar to human reasoning processes.

I. The problem A model of the semantic knowledge of concepts underlying natural language is definitional rather than assertional in that it contains general descriptions of objects and their relations, as opposed to facts about specific objects (Levesque 84). Part of competence in English is the knowledge that an elephant is an animal, and therefore it moves on its own. Competence also involves knowing particular things about elephants, such as that they have trunks. This general description of the elephant concept is part of commonsense knowledge and belief. We will call this the cognitive model. In order to implement it, a computer system must represent what speakers of a language believe about the world and their named concepts, rather than represent the actual world. A complete computer model of the cognitive model would represent the commonsense conceptual scheme presupposed by a particular culture and language, in this case, urban American English.

Knowledge of a natural language implies knowledge of a kind of theory of the environment used by a culture. In learning a language a child learns the category cuts recognized in that theory (Berlin 72) (Dougherty 78). Assuming that knowledge of word meaning is not differently represented than other kinds of knowledge (Tarnawsky 82), KT is designed to encode the world view embodied in natural language as ordinary knowledge, while retaining the autonomy of combinatorial semantics and of syntax. KT does not provide all the meaning, but it yields an interesting and transportable portion of it.

The work reported here addresses two major problems. First, the knowledge associated with a concept does not always give necessary and sufficient conditions for deciding whether an object falls under the concept name. The problem is to find a systematic way of predicting which concepts can be reasoned about using first order logic directly and simply (as a conjunction of predicates) and which ones require default logic. Second, the cognitive model models the actual world, which is open and continuous (Hayes 85). The potential concepts and relations between them are infinite. Nevertheless, humans manage to reason without cognitive overload. Can the computer model be systematically constrained using predictions paralleling those employed by humans in reasoning about the actual world?

Consider the following database of facts (assertions) and axioms (definitions). The language used is irrelevant; any system isomorphic to first order logic will have the same deficiency.

1) FACTS

Human(mary).
Teacher(mary).
Studentof(mary,john).
Teacherof(john, mary).

AXIOMS

$\forall x \text{ Teacher}(x) \rightarrow \text{Human}(x)$.
 $\forall x \exists y \text{ Teacher}(x) \rightarrow \text{Teacherof}(y,x)$.

A question-answering system with the above database of facts and

axioms can respond easily to questions such as (2) but would be unable to answer (3).

- 2) Does Mary have a student?
Who is Mary's student?
Is Mary human?
- 3) Is Mary touchable?
Can Mary move himself about?
What does Mary do?
Are John and Mary part of an institution?

The questions in (3) can be answered by a system which has a taxonomic hierarchy with features at the nodes, such as KL-ONE (Brachman and Schmolze 1985). If Mary is human, Mary is physical object, which has the feature "touchable". Similarly, since Mary is an animal, she can move herself about. KT employs such a taxonomy, and it is called an ontology to reflect the fact that KT reasons with such information as though it were true and complete, in contrast to generic information which is probabilistic. The ontology is unique to KT and is based upon results in cognitive psychology, linguistics and philosophy.

Another deficiency of the database in (1) is that it knows nothing about John, Mary and their relationship, even though English speakers share descriptions of the typical objects in the sets defined by the predicates Human, Teacher and Student. For example, it would be desirable if the system could respond as follows:

- 4) Is Mary intelligent? --Probably so.
Is Mary articulate? --Probably so.
Does John listen to Mary? --Probably so.
Is Mary educated? --Inherently so.
What does Mary do? --Inherently, teaches.

The questions in (4) reflect the kind of things that average people think of when confronted with the predicates (1) (Dahlgren 85). Why not have the AI system infer similarly? In order for such information to be useful, the system needs to know that "intelligent" is a probabilistic feature associated with the predicate Teacher. Therefore, if told $\neg \text{Intelligent}(\text{mary})$ it should be able to reason that (5) while reasoning that (6) is definitely inconsistent.

- 5) $\neg \text{Intelligent}(\text{mary}) \wedge \text{Teacher}(\text{mary})$
- 6) $(\text{Remainder}(X/2) = 0) \wedge \text{Oddnumber}(X)$

A system needs the capacity to reason with prototype information associated with concepts. But the vastness of such information is an obstacle to its use in commonsense reasoning systems. The strategy employed in the KT system is to take advantage of the high degree of structure in prototype information in order to constrain it. Different types of kinds, such as artifacts, natural kinds and persons, are associated with predictably different types of information, and KT exploits these constraints.

II. Diversity in the Lexicon The task of representing meaning for sorts (common nouns) and predicates (verbs and adjectives) themselves, has been impeded by several philosophical problems which are yet to be resolved. The traditional approach, decomposition into conjunctions of other predicates, is notoriously defective. There is no principled way to select or limit the number of other predicates. Suppose the meaning of apple is represented as (7).

- 7) $\text{Apple} \approx \text{Fruit} (\text{Red} \vee \text{Green}) \wedge \text{Round} \wedge \text{Size10}$

Why not add Growsontrees? The proposal to justify the addition of further predicates by the contrast with the meaning of other words has been rejected on a number of grounds (Dowty 79).

Predicate meaning representation is difficult because the domain of the cognitive model is the actual world, which is both open and unknown to a large extent. Humans can never be totally expert about the actual world. And, the knowledge of predicates used by speakers of a natural language varies with expertise, how precise the predicate itself is, and context. Some psychologists maintain that the inherent openness of the actual world is dealt with cognitively by making clear (though possibly inaccurate) category cuts, and then reasoning about categories of objects, including the unclear cases, using prototypes. (Rosch, et al 76) (Smith and Medin 80). This view implies diversity of representations of predicate meanings across the lexicon. Some types of predicates will have criterial features, ODD NUMBER, others, such as names and natural kinds, LEMON, will not.

Because it represents sort and predicate meaning with prototypes, and because it uses first order logic, KT differs in theory and results from systems such as KL-ONE. In KL-ONE, concepts are defined by their roles (descriptive elements) and their subsuming concepts (those concepts superordinate to them in the taxonomy). The concept ELEPHANT is defined by rolesets describing facts such as "has 4 legs", and by its attachment to MAMMAL. The claim is that all and any instantiation of the ELEPHANT concept has 4 legs. In contrast, descriptions in KT are probabilistic. The system accepts elephants with 3 legs, though it knows that elephants inherently have 4 legs. It accepts eggs which are brown, even though it knows that eggs are prototypically white. Further, in KL-ONE, since the descriptions are meant to be defining, non-defining associated information is not encoded. By contrast, KT encodes a great deal of information usually associated with a concept, without the implicit claim that it applies to all instantiations of the concept. ELEPHANT can have features "forgetful", "lumbering" and so forth, without claiming that all elephants have those features.

Another implication of the prototype model is that the content of features is seen as essentially limitless. In contrast, the semantic net model assumes that there is a manageable set of primitive concepts whose size is much smaller than that of the English lexicon, that these are explicitly connected. In KT, only ontological relationships are stated as rules. The relationships between specific descriptions can be derived through problem-solving, but is not encoded. For example, in KL-ONE, the fact that both clouds and eggs are white is directly stated by a link from both CLOUD and EGG to WHITE. In KT, that both have a color is stated in the kind type PHYSICAL OBJECT, but that they both have the same color is reasoned at run time.

The diversity of information KT accepts is constrained by kind types, which predict that associated with ELEPHANT are features describing parts, because ELEPHANT is in the kind type PHYSICAL OBJECT. On the other hand, ELEPHANT does not have features describing its mode of construction because it is not in the kind type ARTIFACT. Thus, the KT system predicts limitless numbers of possible descriptions which are constrained by types deriving from correlational constraints of the actual world.

The KT system differs from most other representations of the commonsense knowledge underlying natural language in taking the content of descriptions from psycholinguistic studies. Because of its empirical basis, KT responds to queries in a natural and human-like way. Though other formalisms could be used to represent empirically-derived models of human commonsense knowledge, KT lends itself to representing the diversity of information found in the data because it allows a virtually unlimited number of features, while organizing them with the kind types.

III. The Kind Types System KT reads geography text, and shows its understanding of the text by answering questions. Text understanding demonstrates the usefulness of the system, but many interesting problems in that area of research are not addressed by this work. KT is written in VM/PROLOG. It uses a parser, a first-order logic translator and a metainterpreter developed by Stabler and Tarnawsky (1985). It employs a set of databases which represent the commonsense ontology, the generic features for sorts, type information for the generic features, and kind types for the ontology. Below is a sample text representative of the English KT understands.

Sample Text John is a miner who lives in a mountain town. His wife raises a chicken who lays brown eggs. The company-owned clinic is near the mine. The nurse monitors the health of the miners. She approves of John's diet.

III.1 The Ontological Schema To capture ontological constraints, KT employs a top-level conceptual schema, some of which appears in Figure 1. It is intended to mirror the average English-speaker's beliefs about what the major category cuts of the environment are, that is, a commonsense ontology.

Figure 1 The Ontological Schema


```

ENTITY → (ABSTRACT ∨ REAL) & (INDIVIDUAL ∨ COLLECTIVE)
ABSTRACT → IDEAL ∨ PROPOSITIONAL ∨ QUANTITY ∨ IRREAL
REAL → (PHYSICAL ∨ TEMPORAL ∨ SENTIENT)
      & (NATURAL ∨ SOCIAL)
PHYSICAL → (STATIONARY ∨ NONSTATIONARY)
      & (ANIMATE ∨ INANIMATE)
NONSTATIONARY → SELFMOVING ∨ NONSELFMOVING
COLLECTIVE → MASS ∨ SET ∨ STRUCTURE
STATIONARY → ¬ MOVEABLE
TEMPORAL → STATIVE ∨ NONSTATIVE
NONSTATIVE → (GOAL ∨ NONGOAL)
      & (PROCESS ∨ ACTIVITY ∨ MOTION)
PROCESS → POSITIVE ∨ NEGATIVE
ACTIVITY → OCCUPATIONAL ∨ INTERACTIONAL
OCCUPATIONAL → AGRICULTURAL ∨ MININGMANU
      ∨ TRADE ∨ SERVICE ∨ EDUCATION
INTERACTIONAL → POSSESSIVE ∨ ASSISTIVE ∨ CONTACTUAL
      ∨ CONFRONTATIONAL
MOTION → (FAST ∨ SLOW) & (TOWARD ∨ AWAY)

```

The goal is to encode an ontology which is consistent with an empirically verifiable cognitive model. As much evidence as possible was derived from psychological research. The schema was developed to handle the predicates found in 4100 words of geography text drawn from textbooks. Despite the complexity of constructing a computer model of the ontology, two commonly-used simplifications, binary trees and planar branching in trees, were rejected. First, though binary trees have simplifying mathematical properties, they are not likely to be psychologically real. People easily think in terms of more than two branches, such as FISH vs BIRD vs MAMMAL, and so on, off of the VERTEBRATE node.

Secondly, most representations assume that each node has a unique parent. But cross-classification is needed since commonsense reasoning uses it. People understand, for example, that entities cross-classify as individuals or sets and real or abstract. This means that at each node, more than one plane might be needed for branching. Cross-classification is handled as in (McCord 85). A type hierarchy is generated which permits each node to be cross-classified in *n* ways. In the top level rule of Figure 1 each entity must be classified both ways, as either ABSTRACT or REAL, and as either INDIVIDUAL or COLLECTIVE. This corresponds to the claim that cognitively there is essentially a parallel ontological schema for collectives. For example, people know that herds consist of animals, so that herds are real and concrete. Thus we have the parallel ontology fragments in (8).



Inheritance of properties works differently for the collectives than it does for individuals. Because cow is under ANIMAL, "cow is a kind of animal" is true. In contrast, herd attaches to ANIMAL, but "a herd is a kind of animal" is not true. A herd consists of animals. We have found that though there are gaps among the collectives, a surprising number of types of entities have collective names in English. For example, prop-

sitions come in collectives (discourse, theme, book). Another important cross-classification involves SOCIAL vs NATURAL. Entities (or events) which come into being (or take place) naturally must be distinguished from those which arise through some sort of social intervention. ARTIFACT is one of the SOCIAL nodes. The distinction needs to be made high up in the ontology because it affects most kind types. For example, events may either be SOCIAL (party) or NATURAL (earthquake). (Section IV expands upon the justifications for the ontology).

The ontology also assumes the possibility of multiple attachments of instantiations to nodes. Thus the representation is actually a lattice rather than a tree. For example, an entity, John, is both a HUMAN with the physical properties of a mammal, and is also a PERSON who thinks. The latter makes John very similar to other sentients such as institutions and social roles. Instead of loading all of that complexity into a single HUMAN node, we make the SENTIENT/NON-SENTIENT distinction high up in the hierarchy. There is ample philosophical (Strawson 53) and psychological (Gelman and Spelke 81) support for this decision. Any actual person is attached to both the HUMAN and PERSON nodes in the ontology.

III.2. Generic Information In the generic features database, each sort is represented as a predicate with two arguments. The first is a list of prototype features and the second is a list of inherent features. A prototype feature is typically associated with a sort or predicate. Most entities have more prototypical features than inherent features. From our sample, a miner is typically "male"; a nurse is typically "female"; a town typically has "houses", "a square", "a fountain", and so on. Inherent features are rationally unrevisable properties of a sort or predicate. Thus, a man is inherently "male", a wife is inherently "married", a house is inherently "house-sized". From our sample, a miner inherently "works in a mine", a nurse inherently is "educated", a town inherently contains "buildings", and so on.

III.3. Feature Types The prototype features are represented by the same set of predicates used to represent the inherent features, thus achieving some economy in the rules. Nevertheless, the number of predicates needed to encode the inherent and prototype features is theoretically limitless. Fortunately, a small and manageable set of 33 feature types encodes a great deal of information, although not exhaustively. The features themselves were chosen empirically to correspond with psycholinguistic data gathered by Rosch et al (1976), Ashcraft (1976) and Dahlgren (1985a) When asked to list prototypical features of various concrete objects, subjects tend to name features which fall into a small number of types such as SIZE, COLOR, SHAPE, and FUNCTION. Similarly, a few types of features such as STATUS, SEX, INTERNAL TRAIT AND RELATION are named for social roles.

Notice that a feature type such as SIZE or COLOR may be inherent for one sort but only prototypical for another. For instance, while blood inherently has COLOR "red", a brick is only prototypically "red". While a brick inherently has SHAPE "rectangular parallelopiped", bread is only prototypically "loaf-shaped". In some cases, a sort has a feature type both inherently and prototypically. For example, a doctor has the inherent FUNCTION "treats sick people" and the prototypical FUNCTION "consoles sick people".

III.4. Kind Types as Metasorts Most knowledge representation systems permit any combination of the features in descriptions. KT limits these combinations by taking advantage of several important ontological constraints affecting the possible real-world objects and therefore possible combinations of features in commonsense knowledge. Objects fall into kinds. In particular, natural kinds exist because their members share some underlying trait, while artifacts and social kinds exist because of social convention Schwartz(1979) Dahlgren (1985b). We call classifications of kinds KIND TYPES, so that NATURAL KIND constitutes one kind type, ARTIFACT another, and so on. Kind types constrain the commonsense knowledge base in several ways. First, each kind type is understood in terms of certain predictable feature types. NATURAL KIND is conceived primarily in terms of perceptual features, while ARTIFACT adds functional features. Second, there is a correlational structure to the features of real-world objects. Given that an object is a

mammal, certain features will be found (eg. "fur") and others will be absent (eg. "feathers").

Associated with each node in the ontology is kind type information encoding feature types entities attached at that node may have. Entities may be described by features falling into a some or all of these feature types, and no others. Inheritance up the tree ensures that any lower node has all the feature types of higher nodes on any path to ENTITY. For instance, any node under PHYSICAL may have certain feature types, and any node under ARTIFACT may have those inherited from PHYSICAL, as well as further feature types, as below:

```
PHYSICAL - Shape, Size, Color, Material
           Texture, Odor, Hasparts, Partof
ARTIFACT - {PHYSICAL}, Function, Operation,
           Construction, Owner
```

At each node, only certain feature types are applicable. Conversely, each feature known to KT is classified by type as a COLOR, SIZE, FUNCTION, INTERNAL TRAIT or other. Cohn (1985) describes the economy of the use of sorts in logic programming. In the KT system, sorts and predicates appear at the terminal nodes of the ontology. In addition, the kind types employed by the system represent metasorts, in that they constrain the possible types of sorts recognized by the system.

III.5. Encoding the Common Sense Knowledge The representations described above will be illustrated with the sort nurse. Nurse is attached to the ontology in axiom

```
9) nurse(X) → role(X).
```

From this axiom nurse inherits SENTIENT, SOCIAL, PHYSICAL, REAL, INDIVIDUAL and ENTITY from the ontology. In the generic database, the axiom (10) lists the prototype and inherent features of nurse.

```
10) nurse ({caring,female}, {educated,assistant,
    help(X,Y) & person(Y) & sick(Y)}).
```

Notice that the last inherent feature is in the form of a PROLOG clause. This makes it possible to use the whole complex feature as input to the English grammar in order to formulate an English response to a question such as "What does the nurse do?", or "Does the nurse help people?". The feature typing database classifies the features as follows:

```
relation(assistant).
internaltrait(caring).
internaltrait(educated).
sex(female).
function(help(*,*)).
```

The kind types predict that as a ROLE, nurse will have certain types of features. Inherited from the SENTIENT kind type are feature types INTERNAL TRAIT ("caring") and GOAL ("tries to help"). Inherited from the SOCIAL kind type are feature types FUNCTION ("takes care of patients") and REQUIREMENT ("license"). In addition, RELATION type features ("assistant") are predicted with a ROLE.

IV. The Inference Mechanism Built into the natural language component by Stabler and Tarnawsky is a metainterpreter which solves queries of all axioms active in the system. This permits us to query ontological and generic information as well as textual information. The translation of the first sentence of Sample Text is as in (11).

```
11) miner(john) & town(town22())
```

The problem solver derives the answers to queries as in (12). matching logic translations of the queries, which are in the form of Prolog goals, to the database.

```
12) Is John a miner? -- Yes
    Does John live in a town? -- Yes
```

In addition, KT is able to make a number of inferences from the text which are not directly stated there. The inferences are drawn from various aspects of the common-sense knowledge built into KT.

IV.1 Inheritance Using the ontological database and the same problem solver, the KT system deduces taxonomically inherited information about the entities mentioned in the text, as in (13)-(14).

- 13) What is a miner?
--A miner is a role, sentient, concrete, social, individual and an entity.
- 14) What does a miner do?
--A miner digs for minerals.
What is digging?
--a goal-oriented, natural, nonmental, real, temporal activity

If an entity has dual attachment, for example as a human and as a role, or as a place and as an institution, then KT explains inheritance relations along both paths of the ontology. A clinic is both a social place and an institution, and so when asked (15),

- 15) What is the clinic?

KT replies both that "A clinic is an institution, sentient, physical, real, collective, structure." and that "A clinic is a social place, place, inanimate, physical, stationary, social, real, individual." Direct ontological questions such as (16) are also answered:

- 16) Is the clinic a social place? --Yes
Is the clinic collective? --Yes

The inheritance path is followed in answering such questions, so that the system can answer not only queries of node attachments at to the terminal nodes of the ontology, but at all higher levels.

IV.2 Complete and Incomplete Knowledge In reasoning with this schema, the system knows which valid inferences it can derive ontologically, and therefore definitively, and which knowledge is incomplete. For example, KT knows that it knows the following for certain:

- $\forall x \text{ Human}(x) \rightarrow \text{Thinks}(x)$
 $\forall x \text{ Teacher}(x) \rightarrow \text{Human}(x)$

It also knows that if something is HUMAN, it is not ABSTRACT. When asked "Is the teacher abstract?" it answers "No". Thus it handles the exclusivity of sets called for by Hendrix (1979) and Tenenbaum(1985). On the other hand, it knows which information is incomplete. With generic descriptions, KT knows that it only knows at the probabilistic level. If asked, "Is Mary intelligent?" it responds "Probably so." This reflects the fact that most English speakers share a prototype of teachers as intelligent. The logic works this way. If a question is ontological, KT gives definitive (yes/no) answers. If the question is generic, the answer is qualified as either prototype or inherent. If no answer can be derived to a non-ontological question, KT responds "I don't know." Thus KT makes the open world assumption except with regard to ontological classifications. This ability to reason about incomplete definitions is similar to Levesque's proposal for incomplete databases (Levesque 84).

IV.3 Prototype and Inherent Features KT answers queries concerning features of the entities in the text, both directly and by types of features. Direct feature queries are of the form (17). The form of the answer depends upon whether the feature is prototypical or inherent.

- 17) Is the miner rugged? --Probably so.
Is the clinic a place? --Inherently so.
Does the chicken lay eggs? --Inherently so.
Are the eggs white? --Probably so.
How is digging done? --Probably with a shovel.
Where is digging done? --Probably in the earth.

IV.4 Overriding Features Generic information is handled differently from ontological information. First, it is tentatively inferred, and checked against the current knowledge base of information built up from the reading the text. If anything in the textual database conflicts with a generic inference, the latter is overridden. KT takes the text as the authority, and if the text says that an entity has a feature contradicting those in its commonsense knowledge of the entity, the text's claim comes first. For example, Sample Text says that the eggs are "brown", which overrides the prototypical generic feature "white" which is listed for egg, as in (18).

- 18) Are the eggs brown? --The text says so.

The cancellation takes place simply by matching to the textual database first. Similarly, if a text said that an elephant had three legs, the KT system would reason that it had three legs, and not the inherent four that elephants have. By overriding inherent features, KT gets around the cancellation problem which arises when features are viewed as logically necessary. If "has four legs" is taken to be a logically necessary feature, any three-legged elephant forces a contradiction, or special processing for exceptions (Brachman and Schmolze 1985). The KT system accepts both facts as true, with no contradiction. This particular elephant has three legs, and elephants inherently have four legs.

In attempting to match to both the textual and generic databases, the possibility of infinite recursion arises. This is true in principle for the human reasoner, as well. KT prevents infinite recursion by limiting inferences to a depth of 5.

Because of the feature typing, KT can answer queries as in (19).

- 19) What color are the eggs?
What function does the clinic have?

Feature typing classifies "brown" as of type COLOR. When KT looks first at the translation of the text to see whether it contains an assertion which states a color for the eggs, it must distinguish the facts in the text which are relevant to the feature type queried. With respect to Sample Text, in order for KT to answer "What color are the eggs?", KT must know that "brown" is a COLOR. Without feature types, KT would not contrast "white" with "brown".

KT deduces sets of facts, as well as individual facts. When queried for a type of feature, such as FUNCTION, KT responds with all functions listed for a sort. For example, clinics prototypically have both outpatient and emergency functions, and KT lists both when queried for function. For sorts which are structural, that is, concrete objects and institutions, KT is able to describe the structure. If asked "What structure does the clinic have?", KT answers that typically it has a hierarchy of head-assistant-clientele and has roles of doctor for head, nurse for assistant and patient for clientele. Similarly, if asked "What structure does the fish have?", KT answers, inherently it has these parts: fins, 1 tail, 1 head, 2 eyes, scales. When KT lists parts, bare plurals mean an unspecified number greater than one.

IV.5 Kind Types The kind types are useful in both parsing and inferencing for text understanding. In the parsing phase, kind types can be used four ways. First, verb sense ambiguity can be resolved by the kind types of subject and object head nouns. In a sentence with the verb take, knowing that the subject is a vehicle forces the choice of the one sense of the verb, and knowing that it is a human forces another. Secondly, KT can reason the other way around, and use selection restrictions on verbs to infer the kind types entities referred to in the sentence. Consider sentence (20).

- 20) ABC sued the man.

Using kind types for selection restrictions, KT infers that the entity named ABC is a SENTIENT. Given the further information in (21), KT infers that ABC is an INSTITUTION and not a PERSON, because the verb join requires an INSTITUTION as object.

21) The man had joined ABC illegally.

Thirdly, certain anaphoric references can be resolved using kind types. When verb selection restrictions classify the entity referred to by a pronoun as in a certain kind type, then possible antecedents are correspondingly constrained. Consider the relationships in (22). It corefers with milk because, when intransitive, spill requires a LIQUID as subject.

22) The cat drank the milk. It spilled.

Fourthly, attachment ambiguities for prepositional phrases can be resolved using kind types. Consider sentence (23).

23) John bought the lock in the afternoon.

It is syntactically possible for the prepositional phrase in the afternoon to modify the lock, the verb phrase or the whole sentence. Since afternoon is in the kind type TEMPORAL, KT can resolve this syntactic ambiguity, and attach the prepositional phrase so that it modifies the whole sentence (Dahlgren and McDowell, 1986).

IV.6 Summary of Inference Mechanism In summary, predications used to derive inferences in the text are found in five aspects of common-sense knowledge:

- a) Ontological Schema
- b) Verb and adjective selection restrictions
- c) Generic Information
- d) Typing Information
- e) Kind types

In using KT, queries drive these inferences. After a text such as the Sample Text has been read, KT can respond to queries and seem to understand the text in a more human-like way using the various aspects of knowledge indicated above. Below are listed some queries and responses.

Q: Who is John?

A: The man who lives in the town.

---Prototypical town has people living in it.

---Prototypical male person is a man (not a boy).

Q: Was the town built?

A: Yes.

---By ontology of artifacts

Q: Who built the town.

A: People.

---By ontology of artifacts

Q: Does John wear pants?

A: Probably so.

---By prototype database.

Q: Does John eat eggs.

A: Yes.

---Because eggs are food.

Q: What does health think?

A: Health doesn't think.

---By kind types.

Q: Does John look like a clinic?

A: No.

---By ontology database.

Q: Does John live in a tent?

A: Probably not.

---By prototype of town

Q: Does John have a function?

A: Yes.

---By kind types

IV. Basis for the Commonsense Knowledge Results in linguistic research underline the importance of category distinctions, such as those between abstract and concrete objects, and persons as opposed to other objects. These actively affect sentence interpretation and generation. The sentence "The rock read the book" must either be interpreted as anomalous or metaphorical because only persons read. These constraints provide an empirical basis for the ontology. Cognitive psychological re-

search provides a further basis for the ontology. Keil's work on ontological categorization in cognitive development was consulted in constructing the schema (Keil 79). Gelman and Spelke's results suggested placing SENTIENT higher in the schema (Gelman and Spelke 81). Graesser and Clark's studies were the basis of the verb ontology (Graesser & Clark, 1985). Psycholinguistic research in the prototype theory provided descriptions of the actual prototypes shared by English-speakers for a number of these categories (Rosch, et al 76) (Dahlgren 85).

The ontological schema was developed in two steps. First, the verbs from the corpus of geography texts were classified according to selectional restrictions (SRs) on subjects and objects. Second, the minimal categories needed to accommodate these SRs were arranged in a hierarchical schema. Certain SRs, such as HUMAN, ANIMATE, CONCRETE, were expected. Others were surprises. Some verbs required complements that were marked for PLACE, and others required either subjects or objects to have certain moveability features. These are summarized below.

STATIONARY: normally immobile, attached to the earth, moved only at great effort.

SELFMOVING: normally in motion or designed for motion, in some cases with no apparent initial source.

NONSELFMOVING: normally immobile but can be moved with slight effort. A source for the motion is expected, usually something SELFMOVING.

One other interesting result from this stage of the project is that a number of verbs take either a PROPOSITIONAL or a SENTIENT subject. Both a book and a person can say something.

Once the set of categories had been established, the next stage was fitting them into a hierarchy from which inheritance of features could be computed by KT. There were several constraints guiding this process. First, we wanted the ontology to be as compact as possible. Second, we wished to minimize nonexistent leaf nodes. Third, we preferred that the system infer too little than too much. During this process it was also necessary to decide which of the SRs represented true category cuts in an ontological schema and which were merely features on individual lexical items. The guiding principle here was that if the distinction under examination (i.e. ANIMATE/INANIMATE) pervaded some subtree, then it was assigned to a branching point. But if some distinction was needed in isolated parts of the tree, then it was represented as a feature. For instance, we found that the INDIVIDUAL/COLLECTIVE distinction pervades the lexicon and must be a primary cut in the ontology. Many verbs select only INDIVIDUAL (mingle) or only COLLECTIVE (stampede) subjects or objects. Properties which were assigned feature status were items like EDIBLE and SIZE.

The NATURAL/SOCIAL distinction was placed high on the tree because human intervention pervades the world. All abstract entities are products of the human mind, but every category of real entities, including events and states, contains dozens of examples of the products of society. We therefore reserved the term ARTIFACT for inanimate man-made objects to distinguish them from natural inanimate objects. The SENTIENT/PHYSICAL distinction is also fairly high. SENTIENT is often placed as a subordinate of ANIMATE, but in commonsense reasoning, the properties of people and things are very different. The NATURAL/SOCIAL distinction applies to SENTIENT just as it does to PHYSICAL. A NATURAL, SENTIENT entity is a PERSON, that is a man, woman, boy, girl, whereas a SOCIAL, SENTIENT entity is a ROLE, secretary, miner, president. A collection of PERSON is a BODY, crowd, mob. A collection of ROLE is an INSTITUTION, hospital, school. The INDIVIDUAL/COLLECTIVE cut had to be made at the level of ENTITY (the highest level) at the same place as the ABSTRACT/REAL cut. This was not the only place where multiple distinctions applied (see Figure 1).

Our term COLLECTIVE applies to all collections of entities, classified into three subgroups. True collectives are sets in which each member of the set is identical to all the others (herd, mob, crowd, fleet). Masses are collections whose members are referred to only in terms of measurable

units (sand, water). Finally, there are structures where the members have specified relations, such as in institutions (school, company, village).

It was consideration of both the constraints listed above, and the assignments of SRs to feature or node status that led us to abandon both binary branching and planar trees as useful representational devices. While it was possible to model some distinctions as binary, others required more than two branches. For example, ABSTRACT entities which divided into IDEAL, PROPOSITIONAL, QUANTITY, and IRREAL, all of which have equivalent status as SRs.

Figure III Terminal Nodes in the Schema

Example	Rule
bush	PLANT ← REAL & INDIVIDUAL & PHYSICAL & NATURAL & ANIMATE & STATIONARY
bear	ANIMAL ← REAL & INDIVIDUAL & PHYSICAL & NATURAL & ANIMATE & NONSTATIONARY & SELFMOVING
mountain	PLACE ← REAL & INDIVIDUAL & PHYSICAL & STATIONARY & INANIMATE
mountain	NATURAL_PLACE ← NATURAL & PLACE
village	SOCIAL_PLACE ← PLACE & SOCIAL
stone	MINERAL ← REAL & INDIVIDUAL & PHYSICAL & NATURAL & INANIMATE & NONSTATIONARY & NONSELFMOVING
Santos	PERSON ← REAL & INDIVIDUAL & NATURAL & SENTIENT
car	VEHICLE ← ARTIFACT & SELFMOVING & PHYSICAL & SELFMOVING
radio	ARTIFACT ← ARTIFACT & NONSELFMOVING
secretary	ROLE ← REAL & INDIVIDUAL & SENTIENT & SOCIAL
book	DISCOURSE ← ABSTRACT & COLLECTIVE & PROPOSITIONAL

We were still faced with the fact that many entities still seemed to straddle the hierarchy. Is an individual human a PRIMATE or a PERSON, or both? Is a hospital an INSTITUTION or a PLACE, or both? If we were to establish a hierarchy which would reflect these differences, we would end up with a very large and unwieldy schema with huge gaps. Therefore, we decided on multiple attachment for those entities which required it. This decision was justified as well by examination of the texts which revealed that a human being was generally dealt with in a context as either a person or a physiological being, but rarely as both at the same time. Figure III gives examples of some nouns, their assignment to categories and rules by which terminal nodes in the schema are generated from higher-level nodes. Figure III shows only a few examples of terminal nodes in the schema. However, every path through the ontology results in a terminal node which is named and which represents a unique class defined by inheritance of features up the tree. Terminal node names distinguish the individuals from the collectives. For instance, the collective node corresponding to PLANT is FLORA. The individual node corresponding to DISCOURSE is PROPOSITION. Similarly, STUFF is the collective of MINERAL, INSTITUTION is the collective of ROLE, and BODY is the collective of PERSON, etc.

The types of features which occurred in the data at each node in the ontology were the basis of the kind types. It is an empirical fact that feature types are correlated in relation to ontological classifications. At each node in the ontology is a kind type encoding certain sets of properties that any entity classified at that node may have. Inheritance up the tree ensures that any lower node has all the properties of higher nodes on a single path to ENTITY. For each property at a node, a set of values applies. While the values for items such as COLOR are fairly obvious, we have had to construct value ranges elsewhere. For SIZE, we have started with the set {microscopic, tiny, small, handleable, medium, large, huge, building-sized, skyscraper-sized, mountainous, region-sized}, which is a reality-oriented scale to be applied loosely. The kind types were extracted empirically from the generic data after all the features were typed, by inspection of types of features associated with sorts and predicates at each node of the ontology.

The texts in the corpus describe lifestyle and industry in various countries. Generic descriptions of the nouns in the text were drawn from the psycholinguistic literature, to the extent possible. ((Rosch 76); (Ashcraft 76); (Dahlgren 85)). For ROLE, we used generic descriptions of social roles collected by Dahlgren and partially published in (Dahlgren 85). For PHYSICAL we used generic descriptions from (Ashcraft 76). For those

nouns where no data existed, generic descriptions were created conforming to the types of information generated by subjects for similar nouns. We do not consider this a defect of our system, since we are not trying to argue for the psychological reality of any particular generic description, but merely for the efficacy of a reasoning system which uses them. The decision to place features in the prototype list or the inherent list for a sort or predicate was decided by two judges. It is a research goal to verify judgments experimentally.

Conclusion In conclusion, KT encodes an ontology which models the top level of typical English speaker's cognitive model of the actual world. It employs several different types of information to reason in human-like ways about text that it reads. In addition to the ontology, it uses verb selection restrictions and generic information associated with concepts. By employing systematic constraints in the form of kind types associated with nodes in the ontology, KT reasons efficiently. All of the information KT uses is drawn from empirical studies of human cognitive psychology, linguistics or the corpus of text which KT reads. Because of this empirical basis, and the breadth of the ontology, KT is a transportable system which is potentially useful for understanding any text of a general, literal nature.

REFERENCES

- Ashcraft, M.H. Property norms for typical and atypical items from 17 categories. *Memory and Cognition*, 6: 227-32, 1978.
- Brachman, R.J. and J.G. Schmolze. 1985. An overview of the KL-ONE knowledge representation system. *Cognitive Science* 9:171-216.
- Berlin, B. 1972. Speculations on the growth of ethnobotanical nomenclature. *Language and Society* 1:41-86.
- Dahlgren, K. 1985a. The Structure of Social Categories. *Cognitive Science* 9:379-398.
- Dahlgren, K. 1985b. Kind types in lexical representation. To appear.
- Dahlgren, K. and J. McDowell. 1986. Using commonsense knowledge to disambiguate. To appear.
- Dougherty, J.W.D. 1978. Salience and relativity in classification. *American Ethnologist*, 5:66-80.
- Dowty, David R. 1979. *Word Meaning and Montague Grammar*. Dordrecht, Holland: D. Reidel Publishing Company
- Gelman, R. and E. Spelke. 1981. Thoughts About Animate and Inanimate Objects, in *Social Cognitive Development*, eds. J.H. Flavell and L. Ross, p. 43-81.
- Graesser, A. and L. Clark. 1985. *Structures and Procedures of Implicit Knowledge*. Norwood, New Jersey: Ablex.
- Hayes, P.J. 1985. The second naive physics manifesto. In *Formal Theories of the Commonsense World*, eds. J.R. Hobbs and R.C. Moore. Norwood, N.J.: Ablex.
- Hendrix, G.G. 1979. Encoding Knowledge in Partitioned Networks, in *Associative Networks*, ed. N.J. Findler, p.51-92.
- Keil, F. C. 1979. *Semantic and Conceptual Development*. Harvard U Press.
- Levesque, H. 1984. The logic of incomplete knowledge bases, in *On Conceptual Modeling*, eds. M. L. Brodie, J. Mylopoulos and J.W. Schmidt, New York: Springer-Verlag.
- McCord, M. 1985. The lexical base for semantic interpretation in a prolog parser. Workshop on the Lexicon, Parsing and Semantic Interpretation. CUNY Graduate Center.
- Rosch, E., Mervis, C.B., Gray, W.D., Johnson, D.M. & Boyes-Braem, P. 1976. Basic objects in natural categories. *Cognitive Psychology* 8:382-439.
- Smith, Edward E. and Medin, Douglas L. 1985. *Categories and Concepts*. Harvard U.
- Stabler, E.P., Jr., and G.O. Tarnawsky. 1985. NLP Prolog--- A prolog-based natural language facility. To appear.
- Strawson, P.C. 1953. *Individuals*. London: Methuen.
- Tarnawsky, G.O. 1982. *Knowledge Semantics*. Unpublished NYU Dissertation.
- Tenenbaum, J.D. 1985. Taxonomic reasoning. *Proc. IJCAI*.