Computers Math. Applic. Vol. 23, No. 6-9, pp. 543-555, 1992
Printed in Great Britain. All rights reserved
0097-4943/92 $5.00 + 0.00
Copyright© 1992 Pergamon Press plc

PROTOTYPES IN A HYBRID LANGUAGE WITH PRIMITIVE DESCRIPTIONS

ENRICO FRANCONI, BERNARDO MAGNINI AND OLIVIERO STOCK
IRST, I-38050 Povo TN, Italy

Abstract—In this paper, the problem of instance recognition within an extended hybrid knowledge representation system is addressed. Structural aspects of concepts are represented at two separate levels, the terminological (using strict definitions) and the prototypical; individuals are expressed in a frame-based assertional component. Primitive descriptions in the terminological language are defined and discussed. The hybrid reasoning mechanism recognizes the type of the individuals with respect to the terminology, making use of reasoning with prototypes associated to primitive concepts. This approach shares basic ideas with the "Dual Theory" about the mental representation of concepts; concepts have a twofold representation: a "core description," useful for compositional meaning, and a further "identification procedure," containing information for recognition of typical instances of a concept. Within the identification procedure a "similarity model" is proposed that describes the probability rating that an object belongs to a class, supported by the similarity that the object shares with the prototype of that class.

1. INTRODUCTION AND MOTIVATIONS

This paper addresses the relation among individuals, concepts, primitive descriptions and prototypes within an hybrid approach *á la* KL-ONE [1-4]. The typical problem we want to solve is the recognition of an instance member of a "natural kind" class: is **Tweety** a `bird`? A `bird` is strictly defined using only the necessary properties that every instance of `bird` must have, while sufficient conditions are missing (i.e., it is a primitive concept, cfr. Brachman and Schmolze [1]). Therefore we can never conclude undeniably that **Tweety** is a `bird`. Unless an explicitly stated membership in the class `bird` is present in the information describing **Tweety**, this fact cannot be derived from the terminological knowledge. Sufficient conditions must instead be represented in the prototypical part of knowledge regarding `birds`. The reasoning part of the prototypical component derives the type for **Tweety** with a similarity mechanism, comparing the description of the instance with the prototype. If that similarity match succeeds (i.e., it reaches a given threshold), it is possible to assume that **Tweety** is a `bird`, since it is similar to a typical `bird`, even if we cannot be one hundred percent confident.

Brachman in [5,6] provided the basis for a clear distinction in a conceptual hierarchy between the "concept as definition" interpretation and the "concept as prototype" interpretation. Two important research works which have addressed that problem within a semantic network framework are Padgham's [7,8] and Shastri's [9] (but see also [10,11]). In Padgham's system concepts are described as having a "type core," for necessary information, and a "type default," for typical information. An extension operation allows one to make a nonmonotonic jump from the core to the default of the same concept, insofar as default information does not conflict with already known information. The system is able to derive inheritance paths with defaults, over which a consistency relation is defined. While core and defaults are qualitatively separated, the lack of a weighted representation for default properties does not allow any decision in multiple choice situations. Shastri, on the other hand, suggests an evidential reasoning formalism based on a semantic network. Nodes represent prototypes or properties, and links are weighted by means of values derived from the object frequency in the knowledge domain. The model provides an

This work was supported in part by the Italian National Research Council (CNR), project "Progetto Finalizzato Sistemi Informatici e Calcolo Parallelo." We thank the anonymous referees for the helpful and incisive comments on an earlier version of this paper.

Typeset by $\mathcal{A}_{\mathcal{M}}\mathcal{S}$-TEX

evidential formalization of inheritance and recognition. However, the lack of a qualitative account does not allow any distinctions between membership and typicality judgements.

In contrast with both Padgham's and Shastri's works our own work is based upon two coexisting basic properties that characterize the concept/prototype relation: (i) given the fact that they represent qualitatively different information, we have two qualitatively different inference mechanisms cooperating in an hybrid system; (ii) assuming the prototype theory is well founded, it is necessary that prototype recognition gives rise to graduated results: this fact imposes a weighted representation for prototypes.

Within the *hybrid* representation paradigm [12–15], different subsystems with a unified semantics provide a convenient way to express—and to reason about—the different kinds of knowledge involved (e.g., terminological in the TBox and assertional in the ABox). Pursuing the ideas presented in [16,17], we propose a further division, such that structural (as opposed to assertional) aspects of concepts are represented at two separate levels, the terminological (strict definitional) and the prototypical; individuals (a.k.a. instances) and sets are expressed in a frame-based assertional ABox component. This hybrid reasoning mechanism recognizes the type of the individuals on the basis of the terminological and prototypical knowledge.

In any realistic knowledge base, incomplete definitions are an important phenomenon that must be treated. We attack the problem of the general relation between terminology and prototypes through the case of primitive descriptions, both for its relevance and because consequences of this extended hybridity are well circumscribed.

Basic ideas are shared with the so called Dual Theory [18,19] about the mental representation of concepts. Within this theory concepts have a twofold representation: a "core description," useful for compositional meaning, and an "identification procedure" for typical instance recognition. Our own realization of such a distinction is that the core strictly defines the necessary and sufficient properties for the concepts (only the necessary ones in the case of primitive concepts), while the identification procedure is a similarity mechanism that works over a collection of perceptual and functional properties. We call such a collection the *prototype* for that concept. Some concepts (in particular primitive concepts such as natural kinds like `bird`) are customarily referred to through the prototype features (a `bird` typically flies and has wings), and whose core description depends on necessary properties (a `bird` is, among other things, an `animal` which *reproduces-through* an `egg`. On the other hand there are concepts that have a core description defining them completely and whose prototype expresses unnecessary but useful typical information in the definition. For example, a `horseman` is precisely a `person` who *rides* a `horse` and has a prototype represented by a set of typical features (e.g., `wears spurs`, and *has-behaviour* `courageous`).

This work starts with the issue of representations of primitive descriptions in a terminological language. We shall introduce a new component in the hybrid reasoning process, based on prototypical knowledge bound to the terminological one through primitive concepts. The central part of our work focuses on the relations between the different representation and reasoning components within the hybrid representation paradigm. We will discuss how this view accommodates a cognitively plausible model for the similarity mechanism, based upon the prototypical component. Results and consequences are maintained in the ABox, as shown in the conclusion of the paper.

2. PRIMITIVE DESCRIPTIONS IN A HYBRID LANGUAGE

In this section, we shall introduce primitive descriptions in the framework of a hybrid language.[1] Regarding the terminological component, we shall show how a definition (i.e., the relation between a name and a terminological description) is to be interpreted for primitive concepts and roles, and how atomic descriptions are related to primitive descriptions. Finally the assertional component is introduced.

[1] Although the assumptions made in our work are mainly independent from the particular hybrid TBox/ABox language, we will refer implicitly to the YAK system [17,18]. Specifically, YAK is quite similar to Classic [2] and Loom [3]. YAK has been developed as the knowledge representation system of the ALFresco natural language system, a multimodal dialogue prototype for the exploration of Italian art history.

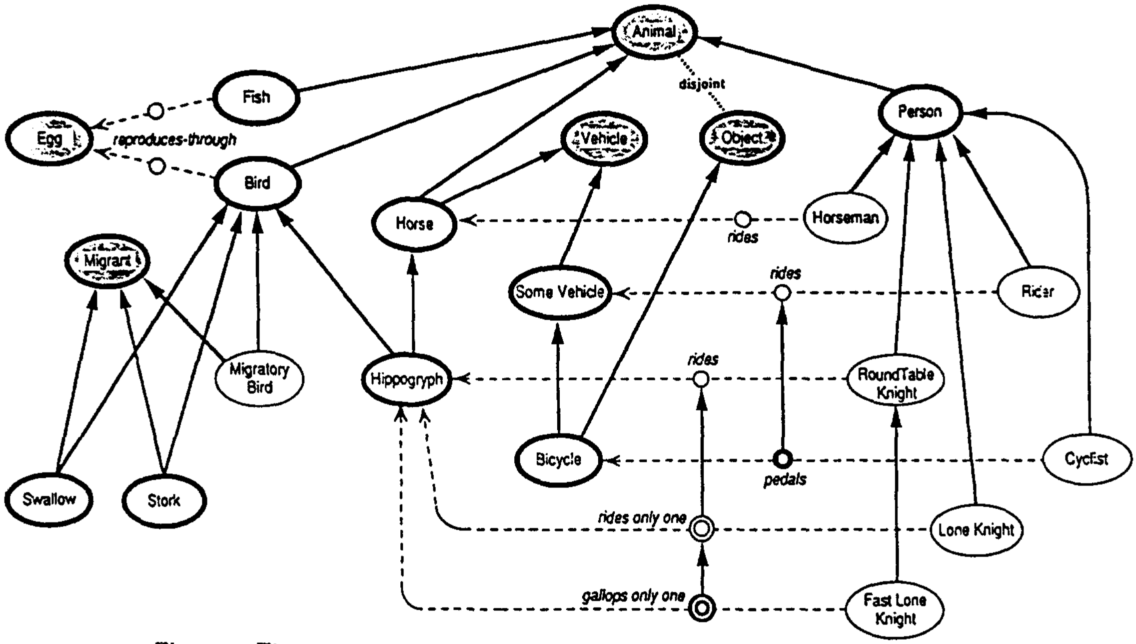

Figure 1. The terminological definitions of the knowledge base referred to throughout the article are depicted here, following the graphical paradigm of KL-ONE. The equivalent linear form representation of this terminology is expressed in Figure 2, using the YAK language. Concepts are represented by the bigger ellipses, and the links emerging from them form their definition. Dashed links denote roles (the small circles) and their restrictions, while solid lines denote the definitional "specializes" relations. A primitive concept, i.e., a concept with an incomplete definition, is depicted with bold ellipses. Roles are represented by the small circles; they can be defined as specialization of other roles (through the solid links), as single-valued (double circles), as primitive (bold circles). Ground atomic concepts, i.e., concepts with no associated definition, are filled with a grey pattern.

TBox descriptions (terms, i.e., concepts and roles) are specified in terms of other descriptions. A concept can be defined as being precisely the conjunction of other concepts (for example[2] a Hippogryph is both a horse and a bird, among other possible properties), or as having a role restricted to another concept (e.g., a horseman is a person who *rides* a horse).

2.1. Incomplete Definitions

The possibility of expressing primitive concepts in the terminological language is introduced to account for the incompleteness of a definition. This is appropriate when the description associated with a concept is merely necessary but not sufficient to define the concept [1]. In analogy with the Dual Theory we call *core description* the description associated to a primitive concept; the core description expresses the necessary conditions an individual must satisfy to belong to the class represented by the primitive concept.

For example, to define completely the concept person it is not sufficient to state that a person is an animal (this is the core description); that is why we introduce that concept as primitive (Figures 1 and 2). With such a primitive concept we do not say what distinguishes a person from an animal; we just say that there is some specialization which makes person a subclass of animal, and we give it a name. As in the case of non-primitive—a.k.a. *definite*—concepts, in such a definition we have associated to a name (person) a description (it is a generic specialization of animal), which expresses exactly the knowledge that defines that term.

In systems in which the objects of predications are concepts—as in McAllester [20]—the association between a name and a concept description (its definition) is indeed an equality between terms. So in every conceptual expression containing a name we can substitute its definition. This

[2] The examples we use throughout the paper refer to the terminological definitions in Figures 1 and 2. In the text concepts appear in plain font, roles in italic font, and individuals in bold font.

is a nice property, because it simplifies the algorithms and the formal study of the language. In the case of primitive concepts the association between a concept name and its description is a set inclusion between terms. In fact, from an extensional point of view, the set of all individuals satisfying the core description associated to a primitive concept is a *generic* subset of the extension of the same description associated to a definite concept. So it is no longer possible to substitute names with their associated conceptual descriptions, if they are primitive concepts. Substitutivity is lost essentially because the name *is* part of the definition of a primitive concept. Informally speaking, we can say that the name of a primitive concept fulfils the part of knowledge which is not specified in the (partial) definition; in other words, the name is part of the semantics of a primitive concept. Formal semantics of primitive concepts is addressed, for example, in [21,22].

```
> (DEFCONCEPT PERSON (PRIMITIVE ANIMAL))                        ; Concept definitions
(PRIMITIVE  ANIMAL  {P}1)
> (DEFCONCEPT HORSEMAN (AND PERSON (EVERY RIDES HORSE)))
(AND  PERSON  (EVERY  RIDES  HORSE))
> (DEFCONCEPT RIDER (AND PERSON (EVERY RIDES SOME-VEHICLE)))
(AND  PERSON  (EVERY  RIDES  SOME-VEHICLE))
> (DEFCONCEPT ROUND-TABLE-KNIGHT (AND PERSON (EVERY RIDES HIPPOGRYPH)))
(AND  PERSON  (EVERY  RIDES  HIPPOGRYPH))
> (DEFCONCEPT CYCLIST (AND PERSON (EVERY PEDALS BICYCLE)))
(AND  PERSON  (EVERY  PEDALS  BICYCLE))
> (DEFCONCEPT LONE-KNIGHT (AND PERSON (EVERY RIDES-ONLY-ONE HIPPOGRYPH)))
(AND  PERSON  (EVERY  RIDES-ONLY-ONE  HIPPOGRYPH))
> (DEFCONCEPT FAST-LONE-KNIGHT (AND ROUND-TABLE-KNIGHT
                                    (EVERY GALLOPS-ONLY-ONE HIPPOGRYPH)))
(AND  ROUND-TABLE-KNIGHT  (EVERY  GALLOPS-ONLY-ONE  HIPPOGRYPH))
> (DEFCONCEPT SOME-VEHICLE (PRIMITIVE VEHICLE))
(PRIMITIVE  VEHICLE  {P}2)
> (DEFCONCEPT BICYCLE (PRIMITIVE (AND SOME-VEHICLE OBJECT)))
(PRIMITIVE  (AND  SOME-VEHICLE  OBJECT)  {P}3)
> (DEFCONCEPT HORSE (PRIMITIVE (AND ANIMAL VEHICLE)))
(PRIMITIVE  (AND  ANIMAL  VEHICLE)  {P}4)
> (DEFCONCEPT HIPPOGRYPH (PRIMITIVE (AND HORSE BIRD)))
(PRIMITIVE  (AND  HORSE  BIRD)  {P}5)
> (DEFCONCEPT FISH (PRIMITIVE (AND ANIMAL (EVERY REPRODUCES-THROUGH EGG))))
(PRIMITIVE  (AND  ANIMAL  (EVERY  REPRODUCES-THROUGH  EGG))  {P}6)
> (DEFCONCEPT BIRD (PRIMITIVE (AND ANIMAL (EVERY REPRODUCES-THROUGH EGG))))
(PRIMITIVE  (AND  ANIMAL  (EVERY  REPRODUCES-THROUGH  EGG))  {P}7)
> (DEFCONCEPT MIGRATORY-BIRD (AND BIRD MIGRANT))
(AND  BIRD  MIGRANT)
> (DEFCONCEPT SWALLOW (PRIMITIVE (AND BIRD MIGRANT)))
(PRIMITIVE  (AND  BIRD  MIGRANT)  {P}8)
> (DEFCONCEPT STORK (PRIMITIVE (AND BIRD MIGRANT)))
(PRIMITIVE  (AND  BIRD  MIGRANT)  {P}9)
> (DEFRELATION PEDALS (PRIMITIVE RIDES))                            ; Role definitions
(PRIMITIVE  RIDES  {P}10)
> (DEFRELATION RIDES-ONLY-ONE (SINGLE-VALUED RIDES))
(SINGLE-VALUED  RIDES)
> (DEFRELATION GALLOPS-ONLY-ONE (PRIMITIVE RIDES-ONLY-ONE))
(PRIMITIVE  RIDES-ONLY-ONE  {P}11)
> (CLASSIFY-ALL)                                            ; Classify the whole KB
OK
> (GRAPH-CONCEPTS)                              ; Draw induced taxonomy (Figure 3)
OK
```

Figure 2. An actual snap of the interaction with the YAK system, while introducing the same terminology expressed in Figure 1. Comparing those definitions with the corresponding definitions expressed in graphical form (Figure 1), the reader can understand the meaning of such a linear syntax. The answers from the system appear in bold font, and denote the internal representations of the concept definitions. Finally the whole knowledge base is classified, i.e., all the valid subsumptions between any two concepts of the terminology (reflected in Figure 3) are computed.

It is still possible to have substitutivity between names and primitive descriptions, if the representation language is augmented with the possibility of marking primitive descriptions with an *index*, which relates a specific name with a unique description. For example (see Figure 2), if fish is a primitive concept defined through the necessary condition of being an animal which *reproduces-through* an egg, by marking it with the index {P}6 it is possible to distinguish it from a bird, which is defined through the same necessary condition but is marked {P}7. In this way unnecessary names are avoided, if they are used only to identify primitives, resulting in an economy for the taxonomy.[3] However, this method of indexing primitives lacks intuitivity within a graphic paradigm, in which the difference between primitives with an equal description is stated simply by drawing them as (spatially) different concepts (as for example fish and bird).

In analogy with concepts, roles also can be specified in terms of other roles, and can be in a taxonomic hierarchy. For example the relation *pedals* is defined as a primitive specialization of the relation *rides*.

2.2. Atomic Terms

An atomic concept (or role) is the simplest kind of concept (or role), the only information associated to it being its name. An atomic concept is equivalent to a primitive concept having the same name and a core description stating that it is a specialization of anything—i.e., the most general term representing the whole domain, at the top of the taxonomy. In fact nothing else can be stated about an atomic concept, except that its extension represents a non-specified subset of the domain.

We make the unique name assumption, thus it is never the case that the model-theoretic interpretations of two atomic concepts with different names are equal in every valid model of the knowledge base (the *theory*). For example, an animal is "different" from a vehicle, even if both these concepts have no information attached to them and appear to have the same syntactical definition except for their names. This means that there may be some valid interpretations of the theory in which animals and vehicles are equivalent, or interpretations in which they are in a hierarchical relation, or interpretations in which they are disjoint, or possibly interpretations in which there is simply an intersection. If we explicitly put a disjointedness constraint between two atomic concepts—as for (inanimate) object and animal—we restrict the number of possible interpretations, such that it is possible to conclude not only that they are not in hierarchical relation, but also that they are disjoint.

2.3. Classification with Primitive Descriptions

The *terminology*, i.e., the set of all the names related to some definition, includes both primitive and definite concepts. The reasoning mechanism—the classifier—deduces all the valid subsumptions between a terminological description and the terminology, i.e., it derives the set of most specialized concepts in the terminology subsuming the description to be classified. The classifier is based on the subsumption procedure, which is able to conclude whether a relation of subsumption holds between two concepts—either definite or primitive. The taxonomic hierarchy is induced from a terminology by classifying every named term (Figure 3): thus there is a unique taxonomy involving both definite and primitive concepts.

It is now shown through some examples how classification is performed in a language allowing primitive descriptions. As usual we refer to the terminology defined in Figures 1 and 2; the complete taxonomy induced by the classifier is shown in Figure 3.

Looking at the case of definite concepts, the classifier deduces that a Round-Table-knight specializes a horseman because he *rides* a Hippogryph.

Primitive concepts are assumed to be "different" unless they are labelled with the same name, even if they have the same core description associated: a bird is different from a fish even if both are defined as generic specializations of animal reproducing through eggs.

[3]Names are not only a syntactic sugar to label complex descriptions, but they are the elements on which the classification service (see below) is based.

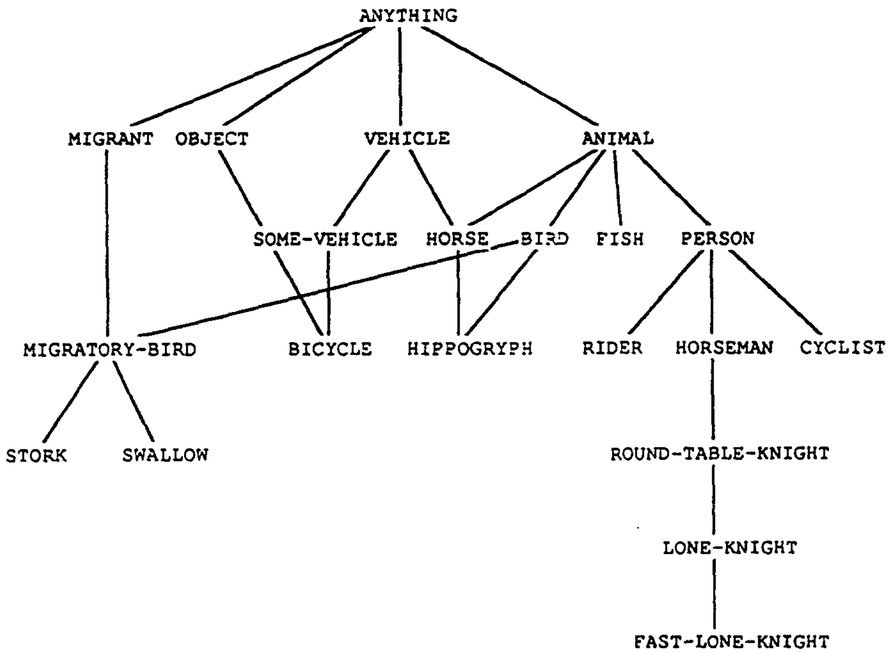

Figure 3. The taxonomy induced by the definitions of the knowledge base shown in Figures 1 and 2. The links represent all the valid subsumption relations between any two concepts, in an optimal way (considering that subsumption is a reflexive and transitive relation).

The classifier will not conclude that a primitive concept subsumes another (definite or primitive) concept, unless the latter mentions this fact explicitly in its definition (i.e., refers to the name of the primitive concept). A `horse` is classified under `animal` because this is stated in its definition: a `horse` is by definition, among other things, an `animal` and a `vehicle`. In the same way the classifier deduces that a `horseman` is a `person`.

However, the classifier cannot deduce that, given the concept `some-vehicle` defined as a generic specialization of `vehicle`, a `horseman` specializes a `rider`, defined as a `person` who *rides* `some-vehicle`. This is false even if a `horse` is a `vehicle`, because a valid interpretation for `rider` could be for example a `cyclist` (bicycle rider).

It is also not true that a `cyclist` is a `rider`, even if intuition suggests that a `person` who *pedals* a `bicycle` should be also a `person` who *rides* `some-vehicle`, given the fact that the relation *pedals* is a specialization of the relation *rides*. This fact derives essentially from the universal condition imposed on the fillers of a relation defining a concept: every `cyclist` *pedals* necessarily a `bicycle`, if he *rides* `some-vehicle`. By restricting the domain and the range of a relation (like the relation *pedals* with respect to the relation *rides*), it is possible that in some interpretation of the theory the extension of the universally defined concept increases (there is some `cyclist` who is not a `rider`). However in the case of functional roles [23]—i.e., relations constrained to have only one filler—this possibility is ruled out and the intuitively expected deduction is satisfied: this is the case of a `fast-lone-knight` who is deduced to be a specialization of a `lone-knight`.

It is also worth noting that, from the disjointedness constraint between `animal` and (inanimate) `object`, a `cyclist` can never be a `horseman` in any interpretation of our theory.

Although the classifier seems to be limited by the incompleteness of knowledge in the specification of a primitive concept [24,25], it can still draw all the correct deductions. The classifier provided within the YAK system has a sound and complete algorithm: it deduces all (and only) the true subsumption relations [21]. The algorithm works in polynomial time if there are no names in the terminology—yet Nebel [26] has proved that any terminological language with naming is inherently intractable. We argue that, thanks to a caching mechanism and to an intelligent name expansion scheduling, the general algorithm is still tractable in the average case.

2.4. The ABox and the Recognizer

Individuals in the assertional component (ABox) are related with the terms representing their *type*. A type represents the set of all the most specific concepts, from the given terminology, which describes the individual; it is obtained by classifying in the terminology the complete "abstraction" of the individual, i.e., a concept describing only that individual. Types for individuals are maintained by a procedure called *recognizer* [15,27,28], which is the basic reasoning mechanism of the ABox.

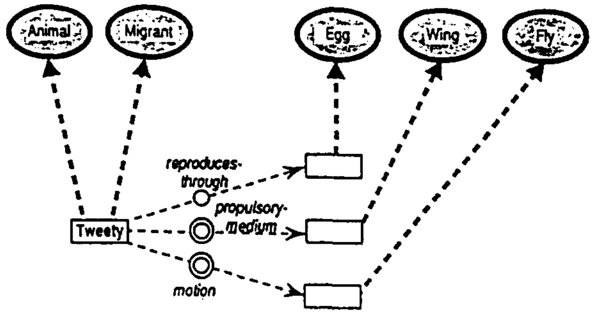

Figure 4. Assertion of an individual in the ABox: the system is informed of the existence of an individual called Tweety, which (referring to the terms defined in the TBox) is an Animal and a Migrant, and has the properties of reproducing through some Egg, moving using with a kind of Wing and flying. The system will deduce, from the terminological knowledge, that the type of Tweety is nothing other than the union of the concepts Animal and Migrant; the system will not deduce that Tweety is a bird without using prototypical knowledge.

For example, looking at Figure 4, the individual **Tweety** is asserted as `migrant` and `animal` which has the properties of reproducing through some `egg`, moving using with a kind of `wing` and flying. The abstraction of this individual is a concept which is classified under the concepts `migrant` and `animal`. In other words the recognizer identifies the individual with respect to the known terms. Since the recognizer is based on the terminological classifier, nothing can be said to be a specialization of a primitive concept (like `bird`) without referring to it directly (for the reasons explained above). That is why the recognizer is not able to conclude nor to refute definitely that the type for **Tweety** is `migratory-bird`, even if we would say that it has the typical properties of a migratory bird.

Another mechanism is necessary to refine the type of an individual with respect to definitorial knowledge: a reasoning mechanism based on prototypical knowledge. A prototype captures sufficient conditions to recognize an individual, whereas in the terminological language a definite concept captures both the necessary and sufficient conditions and a primitive concept only the necessary conditions. In the following we consider prototypes as expressing the knowledge that allows an individual to be recognized as an instance of a primitive concept. Prototypical knowledge is linked to appropriate names in the terminology through primitive concepts.

3. PROTOTYPES

The prototype theory, as developed by Rosch [29,30], suggests some plausible hypotheses for an explanation of typicality effects. One of the basic assumptions is that members of the same category are not equivalent, but they show a degree of representativeness (e.g., a robin is a better exemplar of bird than a penguin) though the ordering may not be total. The typicality value plays a part in much cognitive behaviour, such as differences in speed of processing, free production of exemplars, natural language use of category terms, asymmetries in similarity relationships between category exemplars, learning etc.

Among approaches to prototype representation compatible with Rosch's theory, the feature approach, claiming that prototypes are mentally represented as a collection of relevant features [19] is the cornerstone of our work. Upon this basis we accomodate a cognitive model resulting from the work of Smith and Osherson [31–33]. They suggest a *similarity model* that describes the probability rating that an object belongs to a class supported by the similarity that the object shares with the prototype of that class. The model comprises three components: (a) a prototype representation; (b) a procedure for prototype modification; (c) a rule for determining the similarity between an object and a prototype. We shall introduce in Section 3.1 the representation of prototypes in our framework. Point (c) is addressed in Section 3.2.

The procedure of point (b) is intended to enable prototype combination (e.g., the conjunction of two prototypes). We have not yet investigated this problem, nor more generally the problem of the relations among prototypes.

3.1. Prototype Representation

With respect to our feature-based prototype representation resulting from the Smith and Osherson model, we shall give an account of prototype formation from individual entities and relations within the terminological component.

Prototypes are represented by means of a labelled collection of attribute/value pairs. Attributes can only be roles defined in the terminology, and values are concepts. Each attribute/value pair for a given prototype is provided with two weights: the ***diagnosticity*** of the attribute and the *salience* of the value. How these two measures are determined depends on the prototype formation process.

Prototype formation. We consider a prototype as an abstraction over a set of individual entities, such that it embodies the salient attribute/value pairs. Prototypes refer to primitive concepts in the terminology, so the prototype formation problem is: how to determine sufficient conditions for a given primitive concept? While at this stage we set aside the problems related to learning, in general these can be handled by considering the set of all the instances of the primitive concept and letting an abstraction process work over this set. For the purpose of this paper we can consider the abstraction mechanism as a way of determining both the measures of the diagnosticity of an attribute and the salience of a value. So far we have just based this mechanism on frequencies[4] over the instance set; diagnosticity and salience are defined in the following way:

- attribute diagnosticity: Given a prototype P derived from a primitive concept C, we define the diagnosticity of an attribute A associated to P, as the probability that A describes an instance X, given the fact that X is a member of C.
- value salience: given P as the prototype derived from C, we define the salience of the attribute/value pair A/V as the probability that the A/V pair describes an instance X, given the fact that X is a member of C.

Figure 5 provides a graphical representation for the prototype referring to the primitive concept `bird`.

Relations to the terminology. Though attribute/value pairs correspond to roles and concepts in the terminology, we need a stronger relation between primitive descriptions and prototypes in order to guarantee consistency. The basic idea to this end is to consider the prototype definition within the terminology and then to test its relations to the associated primitive description. We shall assume two steps: first a ***terminological equivalent*** for a prototype is built, making use of the attribute-role and value-concept correspondences; then non-contradiction is verified between the terminological equivalent and the primitive description associated to the prototype.

3.2. Reasoning with Prototypes

Point (c) of the Smith and Osherson model indicates the necessity of a rule for determining the similarity between an object and a prototype. Our prototypical component adopts a version of the *contrast model* developed by Tversky [34] and Osherson [35]. Instance categorisation in this model succeeds when a given threshold on the similarity function is reached.

The contrast model represents the similarity between an instance I and a prototype P depending on three feature sets. The three feature sets are: (1) $I \cap P$, the features common to I and P; (2) $P-I$, the prototype distinctive features; (3) $I-P$, the instance distinctive features. The additive version of the contrast model we adopt is expressed by the equation:

[4] Frequencies constitute a clear framework but they do not account for relevant pragmatic, cultural or social aspects that are essential in a learning model.

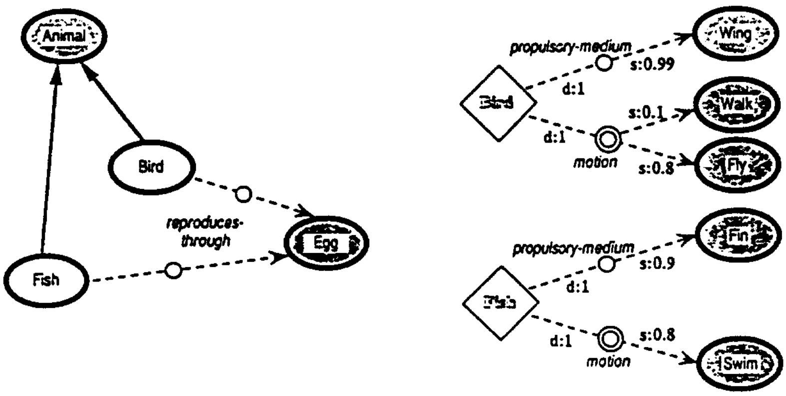

Figure 5. The prototypical representations of the terms Bird and Fish, compared with the corresponding terminological definitions. A primitive definition of a concept corresponds to each prototype. In this example a bird has the necessary property of reproducing through eggs, but not of flying (because not every bird flies), whereas the property of flying can be considered typical for a bird - 80% of birds fly. To walk is less typical for the motion of a bird, but it is still significant.

$$SIM(P,I) = \sum_i [af_i(P \cap I) - bf_i(P - I) - cf_i(I - P)].$$

Where i indexes the relevant attributes, f_i is assumed to multiply the number of votes (salience) for attribute i, and a, b, c are parameters that determine the relative contribution of each attribute set to the overall similarity. As an example consider the way a prototypical type for a certain instance is inferred. Let **Tweety** (Figure 4) have been typed as `animal` and `migrant` by means of a terminological classification, and let the "bird" and the "fish" prototypes (Figure 5) be considered in order to refine the typing on a prototypical basis. The similarity between the instance and the two prototypes is measured by two applications of the contrast model and the best result, provided that it reaches a given threshold, is chosen. Assuming the parameters a, b, c all equal to 1, we obtain:

$$SIM(BIRD, TWEETY) = (0.99 - 0.01 - 0.01) + (0.8 - 0.2 - 0.2) = 1.37$$
$$SIM(FISH, TWEETY) = (-0.9 - 0.1) + (-0.8 - 0.2) = -2.$$

This assigns a clear preference for the "bird" interpretation of **Tweety**.

While this formulation is only based on a sum over the features for single feature dependent functions, it can well be extended to a more sophisticated one for which one can take account of functions that depend on multiple features. In this case relations could be accounted for by a view quite consistent with the rationale that led to the introduction of structural descriptions involving relations in the terminology.

The contrast model respects the following properties that influence human similarity judgements: focus hypothesis (i.e., common features have a greater weight than distinctive features), asymmetry (i.e., the subject has a larger import than the referent of the judgement), context dependency (i.e., context dependent features are more relevant).

3.3. A comment on Probability and Similarity

A large cognitive science literature investigates the relation between typicality and similarity, usually coming to the conclusion that object categorisation is mainly based on similarity judgements—but see Rips [36]. On the other hand, computational models give particular attention to the decision making problem within limited information contexts (i.e., define a calculus that provides a rating for the match between an instance and a prototype). Such a calculus is usually based on some kind of frequency combination among the domain objects, and object categorisation is defined as the best choice among a set of given candidates.

However, while probability theory (at least in the standard Bayesian formulation) is a well known formalism supporting extensional semantics, it has been shown that human behaviour does not strictly follow the probability rules. Tversky and Kahneman [37] point out that at least two pieces of evidence contradict Bayes theorem. In one case they show that people are inclined to violate the conjunction rule when asked to estimate the probability that an object belongs to a compound category. They give the example of Linda, described as "outspoken and concerned with issues of social justice," who is judged more likely to be a feminist bank teller than a bank teller. They explain this behaviour by claiming that people's estimates are based on similarity rather than true probabilities. Moreover in [37] it is claimed that people, when asked to estimate the probability of a hypothesis given some evidence, often ignore the prior probability of the hypothesis. The authors produce examples in which people consider only the representativeness or similarity of the evidence to the hypothesis.

Though the similarity hypothesis allows a better explanation for human behaviour, a deep comprehension of similarity is still just beginning. Two aspects used to define similarity are: the surface-deep differentiation, and the global-dimensional one. Surface similarity is often viewed as bound to perceptual aspects, while deep is bound to functional descriptions. A better version of this might view surface similarity as that depending on readily accessible components of concepts (and therefore often, but not necessarily, limited to perceptual characteristics). On the global-dimension aspect, global similarity is seen as a holistic attitude, while dimensional similarity is conceived of as the ability to differentiate and to make actual use of different dimensions (e.g., for perceptual properties: part identity, colour, size or relations such as "greater than"). With regard to the surface-deep aspect, our approach permits the core definition to be maintained separately from surface aspects, that could, in principle, be characterized also with the level of accessibility. Regarding the second point our approach is not global in the specified sense, but synthetic, with the possibility of restricted application along with a collection of privileged dimensions.

4. HYBRID REASONING WITH PROTOTYPES

We have seen why the type of an individual should be derived not only from definitional properties, but also from another source of knowledge: the prototype. We will see now how this can be accomplished within the hybrid representation paradigm.

The hybrid recognizing process is guided by the terminological classifier, which uses the description of the individual to be recognized, enriched with hypotheses drawn from consulting, in a coherent way, the prototypical knowledge.

At first the individual is recognized as such, according to the definitions present in the knowledge base, resulting in the most precise, universally true, type assignment. For example (Figure 6) the individual **Tweety** is recognized as a `migrant animal`. **Tweety** is not recognized as a `bird` because this fact is not universally true: it is not possible to distinguish it from a `fish`—we cannot exclude, as far as the terms are defined, the existence in some interpretation of a `migrant fish` who has wings and flies.

```
> (CLOSE-ASSERT TWEETY                                              ; Assert Tweety
      (INDIVIDUAL (AND ANIMAL MIGRANT)
            (FILLS REPRODUCES-THROUGH (INDIVIDUAL EGG))
            (FILLS (SINGLE-VALUED MOTION) (INDIVIDUAL FLY))
            (FILLS (SINGLE-VALUED PROPULSORY-MEDIUM) (INDIVIDUAL WING))))
TWEETY
> (INDIVIDUAL-TYPE TWEETY)                                          ; Recognize Tweety
(ANIMAL  MIGRANT)
> (CLOSE-ASSERT TWEETY (INDIVIDUAL BIRD))         ; New assertion (from protot. kb)
TWEETY
> (INDIVIDUAL-TYPE TWEETY)                                    ; Recognize Tweety again
(MIGRATORY-BIRD)
```

Figure 6. An actual snap of the interaction with the YAK system, while asserting and recognizing an individual. The asserted individual Tweety is recognized as both an Animal and a Migrant, with respect to the terminological knowledge base (see also Figure 4). The type of Tweety can be refined consulting the prototypical knowledge—Tweety being similar to a Bird and not very similar to the other known animals. Such a new type is asserted again to specialize the knowledge about the individual. In this way the system can deduce that Tweety is a Migratory-bird.

To discriminate between possible alternative interpretations, the recognizing process verifies whether some hypotheses about the individual can be assumed from the non-definitional part of knowledge (i.e., the prototypical). The prototypes associated to primitive concepts specializing the already computed type are taken into consideration, in order to check their similarity with the individual to be recognized.

Before activating the similarity matching, the compatibility between the individual description and the core description of the primitive concept associated to the prototype is verified. An individual is compatible with the core description of a primitive concept if such a description subsumes the abstraction of the individual. If there is incompatibility (for example a disjointedness constraint or a cardinality constraint is violated) the similarity matching immediately fails, because the preconditions are not satisfied. This is true, for example, of the (in)famous three-legged elephant [6]: if the terminological definition of an elephant states that an elephant must have four legs (as minimum and maximum), an individual similar to an elephant and with three legs will never be recognized as an elephant, and the prototype for elephant should not have any knowledge about the number of legs of a typical elephant (because every elephant has four legs, by definition). On the other hand if the term elephant is described as having possibly from zero to four legs in the TBox, and having typically four legs in the prototypical component, a three-legged animal similar to an elephant can be recognized as elephant without generating a contradiction.

If the individual to be recognized is compatible with the core description, it will be checked if the value of the similarity function, applied on the individual and the prototype, reaches a predetermined threshold. If the match succeeds for more than one prototype, the type of the individual is the conjunction of the separate results only if the corresponding terms are not disjoint. In the latter case the responses are mutually exclusive and the highest value of the similarity function reveals the type to prefer. The success of the similarity matching implies that the individual satisfies the (sufficient) conditions imposed by the prototypical description. The individual is estimated to be a "good" exemplar of such concepts, and its definition is refined by specifying its membership in those classes. At this point the recognition process is activated again and a more specific type for the individual is computed. The general mechanism requires that the terminological classifier/similarity reasoning cycle is recursively activated over and over again in a process of gradual refinement of the individual type assignment.

The importance of verifying the compatibility with the core description before applying the similarity match guarantees the "soundness" of the reasoning based on prototypes with respect to the terminological knowledge. We claim that the similarity mechanism is responsible only for the instance's fuzzy membership of the effective extension of the primitive term, given that the instance belongs to the extension defined by the core description, and it is not responsible for an all-or-none membership judgement. We contrast this with a conceivably different mechanism that requires an instance only to be not in contradiction with the core description (and does not require that it is compatible with it). In this case if the subsequent similarity match succeeds, the instance would "inherit" the core description from the primitive concept associated with the type being asserted for it. We believe that in this latter way the task of similarity reasoning is distorted: it is not only the graduated response about membership, but also a sort of "decision" about the necessity of the inherited description.

Coming back to our example (Figure 6), **`Tweety`** is estimated to be a good exemplar of `bird`, because the property of proceeding with wings makes it more similar to the prototype of `bird` than to the prototype of `fish` (Figure 5). Moreover **`Tweety`** is compatible with the core description of `bird` (**`Tweety`** is an animal which *`reproduces-through`* an egg). The fact of being a `bird` is added to the definition of **`Tweety`** and the type is computed again, resulting in `migratory-bird` (Figure 7). Now this process can be activated once more, to conclude that **`Tweety`** is a `swallow` or a `stork` or neither of the two.

It is worth noting that the type computed using prototypical knowledge has of course a lesser degree of import than the definitory one. A mechanism of belief revision should handle those individuals with a different priority, when subsequently asserted knowledge generates a contradiction in the knowledge base.

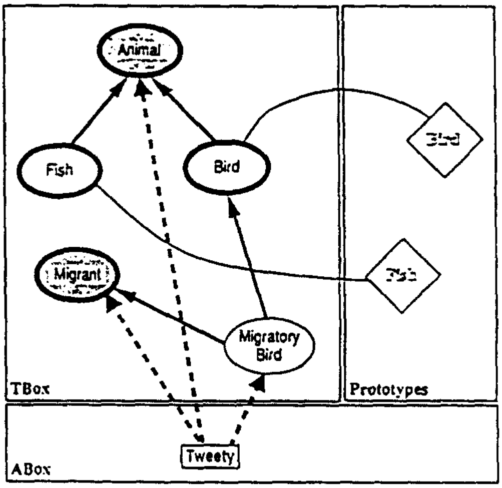

Figure 7. A summary of the hybrid recognizing process for the individual Tweety. As a result of the first classification step, Tweety is recognized to be Migrant and Animal. The individual is therefore compared with the prototypes associated to two concepts specializing Animal: Bird and Fish. Since the similarity matching succeeds for the prototype of Bird, the individual is augmented with the information of being a Bird, and it is recognized again. The final result is that Tweety is recognized to be a Migratory-bird.

5. CONCLUSIONS

In this paper, we have addressed the problem of instance recognition within an extended hybrid architecture. We have focused our attention on relationships between primitive concepts in the terminological component and prototypes in the prototype component. The hybrid reasoning mechanism we propose extends the recognizing process of individuals in the assertional component. It makes use of the terminological knowledge to derive a first type assignment for the individual. This attribution is successively improved by comparing the description of the individual (via a similarity mechanism) to prototypes stored in the prototypical component. The apparatus distinguishes between qualitatively different information and yet can deal with the problem of preferences among the results of similarity-based reasoning.

Within the framework of a complex hybrid architecture supporting multiple reasoning modalities, several aspects must still be addressed. A first extension to our work involves the term-prototype association, in the sense that also well defined concepts should map to a prototype, as stated in the Dual Theory. For example the system should be able to represent prototypical knowledge for the concept horseman, even if it is terminologically defined in a precise way as, in our example, a person who *rides* a horse. In this case the core description is itself sufficient in order to recognize individuals. It is evident that a nontrivial extension of the hybrid reasoner is required.

Another important issue concerns the inference control procedure in the assertional component, as far as a single conclusion can arise out of different modalities which have different import. In these cases a belief revision mechanism is necessary to manage nonmonotonic effects.

REFERENCES

1. R.J. Brachman and J.G. Schmolze, An overview of the Kl-one knowledge representation system, *Cognitive Science* 9 (2), 171–276 (1985).
2. R.J. Brachman, A. Borgida, D.L. McGuiness and L.A. Resnick, The CLASSIC knowledge representation system, or, KL-ONE: The next generation, Presented at the *Workshop on Formal Aspects of Semantic Networks*, Santa Catalina Island, CA, (1989).
3. R. MacGregor, The evolving technology of the KL-ONE family of knowledge representation systems, Presented at the *Workshop on Formal Aspects of Semantic Networks*, Santa Catalina Island, CA, (1989).
4. W. Woods and J. Schmolze, The KL-ONE family, (this volume).
5. R.J. Brachman, What IS-A is and isn't: An analysis of taxonomic links in semantic networks, *IEEE Computer* 16 (10), 30–36 (1983).

6. R.J. Brachman, Defaults and definitions in knowledge representation, *AI Magazine Fall 1985* 6 (3), 80–93 (1985).
7. L. Padgham, A model and representation for type information and its use in reasoning with defaults, In *Proceedings of the AAAI-88*, St. Paul, MN, pp. 409–414, (1988).
8. L. Padgham, Defeasible inheritance: A lattice based approach, (this volume).
9. L. Shastri, *Semantic Networks: An Evidential Formalization and Its Connectionist Realization*, Morgan Kaufmann, San Mateo, CA, (1988).
10. H.W. Beck, Upgrading terminological reasoners to account for current theories of categorization, Statements of interests of *Workshop on Term Subsumption Languages*, Thorn Hill, NH, (1989).
11. B. Pfahringer, Integrating definitions and defaults, Statements of interests of *Workshop on Term Subsumption Language*, Thorn Hill, NH, (1989).
12. R.J. Brachman, V.P. Gilbert and H.J. Levesque, An essential hybrid reasoning system: Knowledge and symbol level accounts of Krypton, In *Proceedings of the IJCAI-85*, Los Angeles, CA, pp. 532–539, (1985).
13. M. Vilain, The restricted language architecture of a hybrid representation system, In *Proceedings of the IJCAI-85*, Los Angeles, CA, pp. 547–551, (1985).
14. P.F. Patel-Schneider, A hybrid, decidable, logic-based knowledge representation system, *Computational Intelligence* 3 (2), 64–77 (1987).
15. B. Nebel, Reasoning and revision in hybrid representation systems, *Lecture Notes in Artificial Intelligence Series*, Springer-Verlag, Berlin, (1990).
16. C. Castelfranchi and O. Stock, Concept-class-prototype: unum an trinum?, In *Proceedings of the 8th Europ. Meeting on Cybernetics and Systems Research*, Vienna, (1986).
17. E. Franconi, B. Magnini and O. Stock, Towards an integration of definitional and prototypical representation of concepts, IRST—Technical Report 8907-04, Trento, Italy, (1989).
18. S.L. Armstrong, L.R. Gleitman and H. Gleitman, What some concepts might not be, *Cognition* 13 (1), 263–308 (1983).
19. E.E. Smith and D.L. Medin, *Categories and Concepts*, Harvard University Press, Cambridge, MA, (1981).
20. D. McAllester, B. Givan and T. Fatima, Taxonomic syntax for first order inference, In *Proceedings of the First Conference on Principle of Knowledge Representazion and Reasoning*, Toronto, Ontario, pp. 289–300, (1989).
21. R. Cattoni and E. Franconi, Walking through the semantics of frame-based description languages: A case study, In *Proceedings of the Fifth International Symposium on Methodologies for Intelligent Systems*, Knoxville, TN, North-Holland, New York, (1990).
22. K. von Luck and B. Owsnicki-Klewe, New AI formalisms for knowledge representation: A case study, *Report TU-Berlin Project KIT-BACK* Vol. 57, Berlin, (1987).
23. E. Franconi, The YAK (Yet Another Krapfen) manual, IRST—Manual 9003-01, Trento, Italy, (1990).
24. I.J. Israel, Interpreting network formalisms, *Comp. and Maths. with Appls.* 9 (1), 1-13 (1983).
25. E. Decio, P. Petrin and L. Spampinato, Pushing the terminological barrier, In *Proceedings of the Workshop on Inheritance Hierarchies in Knowledge Representation and Programming Languages* Viareggio, Italy, Wiley, Chichester, U.K., pp. 107–123, (1989).
26. B. Nebel, Terminological reasoning is inherently intractable, *Artificial Intelligence* 43 (2), 235–249 (1990).
27. R. MacGregor, A deductive pattern matcher, In *Proceedings of the AAAI-88*, St. Paul, MN, pp. 403–408, (1988).
28. F.M. Donini, M. Lenzerini and D. Nardi, An efficient method for hybrid deduction, In *Proceedings of the ECAI-90*, Stockholm, (1990).
29. E. Rosch, Cognitive representation of semantic categories, *Journal of Experimental Psychology: General* 104 (3), 192–233 (1975).
30. C.B. Mervis and E. Rosch, Categorization of natural objects, *Annual Review of Psychology* 32 (1), 89–115 (1981).
31. E.E. Smith, D.N. Osherson, L.J. Rips and M. Keane, Combining concepts: A selective modification model, *Cognitive Science* 12 (4), 485–527 (1988).
32. E.E. Smith and D.N. Osherson, Conceptual combination with prototype concepts, *Cognitive Science* 8 (1), 337–361 (1984).
33. E.E. Smith and D.N. Osherson, Similarity and decision making, *Similarity and Analogical Reasoning* (Edited by A. Ortony and S. Vosniadou), Cambridge University Press, Cambridge, MA, pp. 60–75, (1989).
34. A. Tversky, Features of similarity, *Psychological Review* 84 (4), 327–352 (1977).
35. D.N. Osherson, New axioms for the contrast model of similarity, *Journal of Mathematical Psychology* 31 (1), 93–103 (1987).
36. L.J. Rips, Similarity, typicality, and categorization, *Similarity and Analogical Reasoning* (A. Ortony and S. Vosniadou, Editors), Cambridge University Press, pp. 21–59, (1989).
37. A. Tversky and D. Kahneman, Extensional versus intuitive reasoning: The conjunctive fallacy in probability judgement, *Psychological Review* 90 (4), 293–315 (1983).