Maybury, M. T. (1990). Generating natural language definitions from classification. Proceedings of the 1st ASIS SIG/CR Classification Research Workshop, 101-108. doi: 10.7152/acro.v1i1.12469

Generating Natural Language Definitions from Classification Hierarchies

Mark T. Maybury

The MITRE Corporation, Bedford, MA 01730, USA

(This work was performed at Rome Air Development Center under project number 55812730.)

ABSTRACT

In interactions with users, knowledge based systems are often called upon to define their terms or concepts [Maybury, 1989]. These terms and concepts usually comprise classes within some classification scheme (e.g., a generalization hierarchy). Beyond simply retrieving the superclass of the to-be-defined class (e.g., "a mammal is a vertebrate") a more sophisticated definition also requires selection of distinguishing features or characteristics of this class (e.g., "a mammal is a vertebrate that gives live birth to and nurses its offspring"). To do this, we have refined and extended set theoretic, feature-based models of object similarity and prototypicality, and developed an algorithm that selects the most distinguishing set of attributes and attribute-value pairs of a class in the context of a taxonomy of classes and their properties based on notions of prototypicality and discriminatory power. In this paper, we illustrate a classificatory representation using objects and attribute-value pairs in a test domain of vertebrates; describe our algorithm for computing prototypicality, discriminatory power, and distinctive power, based on this sample representation; and show how this algorithm is implemented to generate definitions of object classes in this representation.

INTRODUCTION

A common method of describing an entity (i.e., an object, action, event, process, or state) is to define it. Perhaps the oldest form of definition is the *logical* (also called *formal*) approach first espoused by Greek orators such as Plato, Aristotle, and Cicero. Logical definition consists of identifying an entity (*species*) by its class (*genus*) and its distinguishing characteristics (*differentia*). Consider: "A parallelogram (*species*) is a quadrilateral (*genus*) whose opposite sides are parallel (*differentia*)." The order of elements of a logical definition is variable: "A polygon of three sides is a triangle."*

Since parents of entities are generally explicitly encoded in a generalization hierarchy found in most knowledge based systems, the genus of an entity can be easily retrieved. Differentia are more complex. In current systems, distinguishing features of the entity (e.g., a brain is unique from other organs because of its function and location) are hand-encoded in the knowledge base [McKeown, 1985]. In contrast to this labor-intensive approach, the algorithm presented in this paper automatically generates an entity's unique characteristics in a domain-independent manner by reasoning about attributes and values of the entity as well as those of closely related entities. Because differentia selection is "on-line", it can be modulated by context and perspective (e.g., we can emphasize structural, functional, or other types of properties depending upon the current context in which the definition is formulated). This differentia algorithm is currently used in

* Logical definition is so common that one model under consideration for the lexeme template of the proposed *Oxford Machine-Readable Dictionary* defines a lexical entry with respect to a genus and a number of differentia [Atkins, 1989].

ISSN: 2324-9773

Maybury, M. T. (1990). Generating natural language definitions from classification. Proceedings of the 1st ASIS SIG/CR Classification Research Workshop, 101-108. doi: 10.7152/acro.v1i1.12469

logical entity definition but can also be applied to referent identification.

The remainder of this paper presents a sample classificatory representation for vertebrates, describes the differentia algorithm showing how it uses this classification, and then illustrates its application by the text planner TEXPLAN [Maybury, 1990], a natural language generation system which produces English definitions of entities in several application domains.

CLASSIFICATORY REPRESENTATION

The vertebrate classification implemented to test and illustrate application of the differentia algorithm was motivated by the set theoretic approach to object similarity in [Tversky, 1977] as well as the psycholinguistic examples of Collins and Quillian [1969] and Rosch [1976 and 1977] who measured human subjects' response time in verifying statements classifying natural objects. A portion of this small generalization hierarchy is as follows:

```
entity
  event
    action
    process
    transition
  state
  object
    animate
      invertebrate
        arachnid
        crustacean
        myriapod
      vertebrate
        amphibian
        bird
          bluebird
          canary
          cardinal
          crow
          hummingbird
          penguin
          robin
          sparrow
        fish
        mammal
          humans
          livestock
          whale
        reptile
    inanimate
```

Each entity is encoded using the notation of *objects* (also known as *frames*), *attributes* (also known as *slots*), and *values* of attributes. An attribute together with its value is known as an *attribute-value pair*. The following is a typical object, illustrating *bird* and its associated attribute-value pairs:

ISSN: 2324-9773

Maybury, M. T. (1990). Generating natural language definitions from classification. Proceedings of the 1st ASIS SIG/CR Classification Research Workshop, 101-108. doi: 10.7152/acro.v1i1.12469

bird	
nervous-system	yes
segmented-spinal-column	yes
reproduction	eggs
respiratory-system	lungs
blood	warm
natural-habitat	terrestrial
covering	feathers
propellers	wings
subparts	(crest crown bill thorax tail)
domesticated	no
movement	flies
eats	seeds

Object classes may inherit attribute-value pairs from their superclasses, or they may have the same attributes with overriding values, or they may have new attribute-value pairs. For example, the *hummingbird* object, a child of the *bird* object class, may inherit from the above *bird* object all the attribute-value pairs from *nervous-system yes* through *movement flies*. Furthermore, it may have the following local attribute-value pairs:

hummingbird	
eats	flower-nectar
size	(3.5 inches)
date-named	1637
origin	North-America
makes-sound	humming
speed	fast
tongue	very-extensile
wing-span	narrow
bill-size	slender
color	brilliant

It is these local attribute-value pairs that in general are the most unique features of the object. We next detail numerical measures that guide the selection of the distinguishing features of an object.

DIFFERENTIA ALGORITHM

The differentia algorithm is based on two numerical measures which are used to select distinguishing attributes and values (i.e., differentia) of a given object. A third measure is derived from these. The range for each of these three measures is [0,1]. The first measure, *P*, indicates the prototypicality of a given attribute or attribute-value pair (i.e., its commonness). The second measure, *D*, indicates the discriminatory power of a given attribute or attribute-value pair (i.e., its uniqueness). Both measures are dependent upon the context of related objects in a generalization hierarchy (e.g., if some feature f, is characteristic of some entity e, as well as of all its siblings, then f is not very discriminating of e). A composite of prototypicality and discriminating power yields the distinctive power, *DP*, of an attribute or attribute-value pair of an object. Using this third measure, distinctive features of an object -- its differentia -- can be selected.

ISSN: 2324-9773

Maybury, M. T. (1990). Generating natural language definitions from classification. Proceedings of the 1st ASIS SIG/CR Classification Research Workshop, 101-108. doi: 10.7152/acro.v1i1.12469

Prototypicality, P, is measured in relation to an object's children. An attribute or attribute-value pair is prototypical of an object if it is found in each of this object's children; if this is the case, then P equals 1. Conversely, if the attribute or attribute-value pair is found in none of this object's children, then it is not at all prototypical, and P equals 0. Prototypicality of an attribute *<a>* with respect to an object's children *c* may be expressed as follows:

$$P(<a>,c) = \frac{\text{no. of c which have <a>}}{\text{total no. of c}}$$

Similarly, prototypicality of an attribute-value pair *<a,v>* with respect to an object's children *c* may be expressed as follows:

$$P(<a,v>,c) = \frac{\text{no. of c which have <a,v>}}{\text{total no. of c}}$$

In contrast to prototypicality, which is measured with respect to an object's children, *discriminatory power, D*, is measured in relation to an object's siblings. An attribute or attribute-value pair has maximum discriminatory power for an object if it is not found in any of this object's siblings; if this is the case, D equals 1. Conversely, if the attribute or attribute-value pair is found in each sibling of this object, then it has no discriminatory power, and D equals 0. Discriminatory power of an attribute *<a>* with respect to an object's siblings *s* may be expressed as follows:

$$D(<a>,s) = \frac{\text{total no. of s - no. of s which have <a>}}{\text{total no. of s}}$$

Similarly, discriminatory power of an attribute-value pair *<a,v>* with respect to an objects's siblings *s* may be expressed as follows:

$$D(<a,v>,s) = \frac{\text{total no. of s - no. of s which have <a,v>}}{\text{total no. of s}}$$

Thus, based on the above measures, we can order attributes or attribute-value pairs according to their prototypicality for some object, as well as according to their discriminatory power for that object. Finally, a composite of prototypicality and discriminating power yields the *distinctive power, DP*, of an attribute or attribute-value pair with respect to an object. DP may be expressed as follows for attribute *<a>* and attribute-value pair *<av>*, respectively:

$$DP(<a>) = \frac{P(<a>,c) + D(<a>,s)}{2}$$

$$DP(<a,v>) = \frac{P(<a,v>,c) + D(<a,v>,s)}{2}$$

Consider the object *bird* and attribute-value pairs *covering feathers*, *movement flies*, and *blood warm* from the *bird* object in the preceding section. Assume that *bird* has four siblings, as in the classification in the preceding section, and 250 children, and that these attributes — *covering*, *movement*, and *blood* — hold for all children and siblings of *bird*. Therefore, since *no. of children which have the attribute = total no. of children*, and *no. of siblings which have the attribute = total no. of siblings*, each measure — *P(<a>,c)*, *D(<a>,s)*, and *DP(<a>)* — would be the same, respectively, regardless of which of these three attributes is considered. In particular, with respect to only attribute, the following equalities hold:

ISSN: 2324-9773

Maybury, M. T. (1990). Generating natural language definitions from classification. Proceedings of the 1st ASIS SIG/CR Classification Research Workshop, 101-108. doi: 10.7152/acro.v1i1.12469

$$P(\text{covering},c) = P(\text{movement},c) = P(\text{blood},c) = \frac{250}{250} = 1$$

$$D(\text{covering},s) = D(\text{movement},s) = D(\text{blood},s) = \frac{4-4}{4} = 0$$

$$DP(\text{covering}) = DP(\text{movement}) = DP(\text{blood}) = \frac{1+0}{2} = .5$$

That is, each attribute is absolutely prototypical for the object *bird* considering its children, but has absolutely no discriminatory power with respect to its siblings; hence, the distinctive power for each attribute is one-half.

Measures for attribute-value pairs are more interesting in this case. In addition to previous assumptions for object *bird*, assume that six children represent species of flightless bird. Then, for object class *bird*, prototypicality based on these three attribute-value pairs would be as follows, in descending order (the first two attribute-value pairs are tied):

$$P(<\text{covering,feathers}>,c) = \frac{\text{no. of c which have <covering,feathers>}}{\text{total no. of c}} = \frac{250}{250} = 1$$

$$P(<\text{blood,warm}>,c) = \frac{\text{no. of c which have <blood,warm>}}{\text{total no. of c}} = \frac{250}{250} = 1$$

$$P(<\text{movement,flies}>,c) = \frac{\text{no. of c which have <movement,flies>}}{\text{total no. of c}} = \frac{244}{250} = .936$$

Discriminatory power for object class *bird*, based on these same attribute-value pairs, would be as follows, in descending order (the first two attribute-value pairs are tied):

$$D(<\text{covering,feathers}>,s) = \frac{\text{total no. of s - no. of s which have <covering,feathers>}}{\text{total no. of s}} = \frac{4-0}{4} = 1$$

$$D(<\text{movement,flies}>,s) = \frac{\text{total no. of s - no. of s which have <movement,flies>}}{\text{total no. of s}} = \frac{4-0}{4} = 1$$

$$D(<\text{blood,warm}>,s) = \frac{\text{total no. of s - no. of s which have <blood,warm>}}{\text{total no. of s}} = \frac{4-1}{4} = .75$$

Finally, distinctive power for object class *bird*, based on these attribute-value pairs, would be as follows, in descending order:

$$DP(<\text{covering,feathers}>) = \frac{P(<\text{covering,feathers}>,c) + D(<\text{covering,feathers}>,s)}{2} = \frac{1+1}{2} = 1$$

$$DP(<\text{movement,flies}>) = \frac{P(<\text{movement,flies}>,c) + D(<\text{movement,flies}>,s)}{2} = \frac{.936+1}{2} = .968$$

$$DP(<\text{blood,warm}>) = \frac{P(<\text{blood,warm}>,c) + D(<\text{blood,warm}>,s)}{2} = \frac{1+.75}{2} = .875$$

Thus, for object class *bird*, being covered with feathers has greater distinctive power than movement by flying, which in turn has greater distinctive power than being warm-blooded.

DEFINITION GENERATION

Definitions of object classes in knowledge bases are generated by TEXPLAN, which uses the differentia algorithm to select the propositional content of a logical definition by (1) retrieving the parent(s) of the object to be defined and (2) selecting the characteristics with maximum DP.

ISSN: 2324-9773

Maybury, M. T. (1990). Generating natural language definitions from classification. Proceedings of the 1st ASIS SIG/CR Classification Research Workshop, 101-108. doi: 10.7152/acro.v1i1.12469

For example, calculating the distinctive features of the object class *vertebrates*, the algorithm uses measures of prototypicality (P), discriminatory power (D), and distinctive power (DP), as described in the preceding section, to collect features common to its children (e.g., all vertebrates have a nervous system and a segmented spinal column) and then determine which features are unique with respect to its siblings (e.g., invertebrates don't have spinal columns).

ATTRIBUTE-VALUE PAIRS	*P*	*D*	*DP*
(movement flies)	1.0	1.0	1.0
(propellors wings)	1.0	1.0	1.0
(covering feathers)	1.0	1.0	1.0
(eats seeds)	.88	1.0	.94
(blood warm)	1.0	.75	.88
(subparts (crest crown bill tail ...))	1.0	0	.5
...			
(segmented-spinal-column t)	0	0	0

TABLE 1. Prototypicality (P), Discriminatory power (D), and Distinctive Power (DP) of attribute-value pairs of object class *bird.*

TABLE 1 shows the calculated values of P, D, and DP for attribute-value pairs of the object class *bird*. Using the DP value we can select the most distinctive features of a bird: its flying motion, wings, feathers, and seed-eating characteristics. Similarly, the most distinctive features of a canary are that it is a yellow, domesticated, singing bird from the Canary Islands. This corresponds to the logical definition generated by the system in [Maybury, forthcoming]:

> A canary is a yellow bird with a Canary Islands origin, that sings, and is domesticated.

The most distinctive characteristic(s) can be given a prominent surface position (or intonation), such as modifying the head noun in the object position. For example, notice above how "yellow", the most salient property of a canary, modifies the subject "bird".

The differentia algorithm has been applied in a number of other domain knowledge bases, including neuropsychology [Maybury and Weiss, 1987], mission planning [Dawson et al., 1987], and knowledge based battle simulation [Anken, 1989]. Logical definitions generated by the system using these various domain knowledge bases include the following:

> An optical lens is a component for focusing located in a camera.
>
> A brain is an organ located in the skull consisting of gray nerve tissue and white nerve fibers.
>
> An A-10 is a fighter for air-to-ground interdiction.

ISSN: 2324-9773

Maybury, M. T. (1990). Generating natural language definitions from classification. Proceedings of the 1st ASIS SIG/CR Classification Research Workshop, 101-108. doi: 10.7152/acro.v1i1.12469

A variation on logical definition would include other types of information, not normally considered differentia, such as subparts or purpose. Use of purpose in place of differentia is illustrated by the last example given above. The distinguishing characteristics of an A-10 are computed by recognizing that other classes of aircraft (e.g., tankers/cargo, reconnaissance, etc.) have similar attributes (e.g., speed, range, empty and loaded weights, etc.) and only slightly differing values. However, they do have unique tactical roles or purposes, and this is what distinguishes the A-10 from them. For use of subparts in definitions, consider:

> A bicycle is a light vehicle having two wheels, one behind the other, a steering handle, a saddle seat(s), and pedals by which it is propelled. (Webster's New Collegiate Dictionary, 1957)

Unlike computing differentia, non-differentia such as subparts or purpose are usually simply looked up in the underlying knowledge base [McKeown, 1985]. To generate these types of definition, however, requires more abstract knowledge representation schemes which distinguish between structural, functional, and other types of information.

One interesting area for further work involves considering how differentia may be affected by different user perspectives. For example, if the system believes that the user has a goal, then the system can identify the purpose of an entity if it can help the user achieve their goal. Hence the system can tailor its definitions to its model of the user.

CONCLUSION

This paper reports on a differentia algorithm that computes the most distinctive characteristics of a given object class by reasoning about its features and those of its relatives in a generalization hierarchy. This algorithm is expoited to automatically generate English logical definitions of object classes. The algorithm has been implemented. It is currently used to generate logical definitions from several domain knowledge bases. What remains to be done is to test and validate the differentia algorithm with psycholinguistic evidence and to extend its application to other types of entities such as actions, events, and states.

REFERENCES

Anken, C. S. June, 1989. "LACE: Land Air Combat in Eric." *Rome Air Development Center DL-9-0043.*

Atkins, B. T. S. 1989. "Building a Lexicon: Reconciling Anisomorphic Sense Differentiations in Machine-readable Dictionaries." *Proceedings of the BBN Natural Language Symposium, BBN, Cambridge, MA, 29 November - 1 December, 1989.*

Collins A. M. and Quillian M. R. 1969. "Retrieval Time from Semantic Memory." *Journal of Verbal Learning and Verbal Behavior* 8(1969):240-247.

Dawson, B., R. Brown, C. Kalish and S. Goldkind. May, 1987. "Knowledge-based Replanning System." *Rome Air Development Center TR 87-60.*

ISSN: 2324-9773

Maybury, M. T. (1990). Generating natural language definitions from classification. Proceedings of the 1st ASIS SIG/CR Classification Research Workshop, 101-108. doi: 10.7152/acro.v1i1.12469

Maybury, M. T. and C. Weiss. 1987. "The NEUROPSYCHOLOGIST Expert System Prototype." *Intelligence Integration, Proceedings of the Canadian Information Processing Society, Edmonton, Alberta, Canada, 16-19 November, 1987.*

Maybury, M. T. 1989. "GENNY: A Knowledge-Based Text Generation System." *Information Processing & Management* 25(2):137-150.

Maybury, M. T. 1990. "Custom Explanations: Exploiting User Models to Plan Multisentential Text." *Proceedings of the Second International Workshop on User Models, University of Honolulu, Hawaii, 30 March - 1 April, 1990.*

Maybury, M. T. forthcoming. *Generating Multisentential English Text Using Communicative Acts.* Ph.D. Dissertation. Cambridge University, England.

McKeown, K. 1985. *Text Generation.* Cambridge University Press.

Rosch, E., C. B. Mervis, W. D. Gray, D. M. Johnson and P. Boyes- Braem. 1976. "Basic Objects in Natural Categories." *Cognitive Psychology* 8(1976):382-439.

Rosch, E. 1977. "Classification of Real World Objects: Origins and Representations in Cognition." *Thinking: Readings in Cognitive Science,* editors P. N. Johnson-Laird and P. C. Wason. Cambridge: Cambridge University Press.

Tversky, A. 1977. "Features of Similarity." *Psychological Review* 84(1977):327-352.

ISSN: 2324-9773