

THE ROLE OF PERSPECTIVE IN RESPONDING TO PROPERTY MISCONCEPTIONS'

Kathleen F. McCoy
Department of Computer & Information Science
University of Pennsylvania
Philadelphia, PA 19104

Abstract

In order to adequately respond to misconceptions involving an object's properties, we must have a context-sensitive method for determining object similarity. Such a method is introduced here. Some of the necessary contextual information is captured by a new notion of *object perspective*. It is shown how object perspective can be used to account for different responses to a given misconception in different contexts.

1. Introduction

As a user interacts with a database or an expert system, s/he may attribute a property or property value to an object that that object does not have. For instance, imagine the following query to a database.

U. Give me the HULL-NO of all DESTROYERS whose MAST-HEIGHT is above 190.

If a system were to evaluate such a query, it might find that there are no such ships in the database. The reason for this is that the user has queried a value of the property MAST-HEIGHT that it cannot have for the object DESTROYER. I term this error a *property misconception*. Upon encountering such a query, even a very cooperative system could only respond:

S. There are no DESTROYERS in the database with a MAST-HEIGHT above 190. Would you like to try again?

In most cases, however, this is not the way a human would respond. A study of human/human transcripts reveals that a human conversational partner often tries to get at the cause of the misconception and offer additional information to correct the wrong information. The additional information often takes the form of a correct query that is a possible alternative to the user's query. In this paper I describe some of the knowledge and reasoning that are necessary for a natural language interface to a database or expert system to mimic this human behavior.

In the above query, since there is an object similar to a DESTROYER that has the value of HULL-NO given, the users misconception may result from his/her confusing the two objects. Hence a reasonable response would be:

S. All DESTROYERS in the database have a MAST-HEIGHT between 85 and 90. Were you thinking of an AIRCRAFT-CARRIER?

Notice the strategy used to correct the misconception is to (1) deny (implicitly) the property/value given, (2) give the corresponding correct information, (3) suggest an alternative query containing the object the user may have confused with the misconception object.

In other situations, a reasonable alternative query might involve the same object the user asked about, with a different

property/value pair. This is the case in the following query.

U. Give me the HULL-NO of all DESTROYERS whose MAST-HEIGHT is above 3500.

S. All DESTROYERS in the database have a MAST-HEIGHT between 85 and 90. Were you thinking of the DISPLACEMENT?

This response is similar to the one given above except that the alternative query suggests an attribute rather than an object which may have been confused.

In general, there can be two major reasons why a wrong attribution may occur. Either (1) the user has the wrong object - that is, s/he has confused the object being discussed with a similar object or has reasoned (falsely) by analogy from a similar object; or (2) the user has the wrong attribute - that is, s/he has confused the attribute being discussed with a similar attribute. If one of these two can be seen as likely in a given situation, then a revised query can be suggested which mentions the similar object or the similar attribute.

To propose alternative queries, a system must have a method for determining similarity of objects and attributes. In this paper I will focus on responses involving object confusion; thus I will examine a similarity metric for objects. In the next section such a similarity metric is introduced. The following section introduces a new notion of object perspective which is needed to provide the similarity metric with some necessary contextual information, in particular, attribute salience ratings. Finally, an example of how perspective information and the similarity metric can be used to give reasonable responses to misconceptions involving object properties is given.

2. Object Similarity

As was shown above, in order to respond effectively to property misconceptions, we must have a method for determining object similarity. Object similarity has previously been shown to be important in tasks such as organizing explanations [6], offering cooperative responses to pragmatically ill-formed queries [2], and identifying metaphors [9]. In the above systems the similarity of two objects is based on the distance between the objects in the generalization hierarchy. One problem with this approach is that it is *context invariant*. ** That is, there is no way for contextual information to affect similarity judgments.

However, Tversky [8] proposes a measure of object similarity based on common and disjoint features/properties of the objects involved, which enables contextual information to be taken into account. Tversky's similarity rating for two objects a and b , where A is the set of properties associated with object a and B is the set of properties associated with object b , can be expressed as:

$$s(a,b) = MCA \ 0 \ B) - at \ (A - B) - Of \ (B - A)$$

for some ϕ , a , 0 ≥ 0 . This equation actually defines a family of similarity scales where ϕ , a , and 0 are parameters which alter the importance of each piece of the equation, and f maps over the features and yields a salience rating for each. The equation states

that the similarity of two objects is some function of their common features minus some function of their disjoint features. The importance of each feature involved (determined by the function f) and the importance of each piece of the equation (determined by S , a , and β) may change with context.

Previous work [4, 7] has discussed the effect of 'focus' on the prominence of objects. Focusing algorithms can be adapted to set the values of O , a , and O . For instance, if object a is "in focus" and object b is not, then the features of a should be weighted more heavily than the features of b . Thus we should choose $a > B$ so that the similarity is reduced more by features of a that are not shared by b than vice versa.

The problem then is to determine f . Other work [3, 0] has hand encoded salience values for the attributes of individual objects in the knowledge base, effectively setting the f function once and for all. This approach, however, is not sufficient since salience values must change with context. The following examples in which two objects (Treasury Bills and Money Market Certificates) are compared in two different circumstances, illustrate the importance of context on the similarity rating.

Consider someone calling an expert financial advisor to see if she can better invest her money. She begins by telling the expert where her money is:

- V. We have \$40,000 in money market certificates. One is coming due next week for \$10,000... I was wondering if you think this is a good savings...
- E. Well, I'd like to see you hold that \$10,000 coming due in a money market fund and then get into a longer term money market certificate.
- V. Hm... well I was just wondering, what about a treasury bill instead?
- E. That's not a bad idea but it doesn't replace your money market certificate in any way — it's an exact duplicate. They're almost identical types of instruments - so one, as far as I'm concerned, is about the same as another.

Now consider how the same two objects can be seen quite differently when viewed in a different way. Imagine the following conversation:

- U. I am interested in buying some US Government Securities. Now I was thinking of Money Market Certificates since they are the same as Treasury Bills.
- E. But they're not - they are two very different things. A Treasury Bill is backed by the U.S. Government; you have to get it from the federal reserve. A Money Market Certificate, on the other hand, is backed by the individual bank that issues it. So, one is a Government Security while the other is not.

In the first example the objects are viewed as savings instruments. This view highlights attributes such as interest-rates and maturity-dates that are common to Treasury Bills and Money Market Certificates. This highlighting causes the two instruments to be seen as "identical". In contrast, the second example views the objects as instruments issued by a particular company or organization. In this case attributes such as issuing-company and purchase-place are highlighted. Since these highlighted attributes are different for the two objects, the objects are seen as being quite different.

As the examples illustrate, a context-free metric of similarity is not sufficient; contextual information is needed. A notion of object perspective, introduced below, can capture the needed contextual information. In particular, perspective accounts for how the f function (the assignment of salience values to various attributes) changes with context.

3. Perspective

[4, 1] note that the same object may be viewed from different perspectives. For instance a particular building may be viewed as an architectural work, a home, a thing made with bricks, etc. According to this work, an object viewed from a particular perspective is seen as having one particular superordinate, although in fact it may have many superordinates. The object inherits properties only from the superordinate in perspective. Therefore different perspectives on the same object cause different properties to be highlighted.

Although this notion of perspective is intuitively appealing, in practice its use is rather difficult since it hinges on the use of a limited inheritance mechanism. The problem is that attributes may be inherited from the top of the generalization hierarchy, not just from immediate superordinates. So, an object's perspective involves not just one superordinate but a chain of superordinates. Therefore one must not only determine what perspective a particular object is being viewed from, but also what perspective its superordinate is viewed from, and so on. As one continues up the hierarchy in this fashion, the definition of perspective as viewing an object as a member of a particular superordinate becomes less and less appealing.

In addition, this notion of object perspective says nothing about the density of the generalization hierarchy. That is, in some situations the immediate superordinate of an object (and the properties it contributes) may be ignored. For example, even though a whale is a cetacean (a class of aquatic mammal including whales and porpoises), this classification (and all attributes contributed by the classification) may be ignored in some situations in which the important attributes instead are inherited from a superordinate of cetacean, say, mammal. In other situations, the class "cetacean" may be central. The notion of object perspective outlined above has no way of determining whether or not certain superordinates should be ignored or included.

Here I introduce a new notion of perspective which is able to handle both the assignment of differing salience values and the density problem. In this notion, perspectives sit orthogonal to the generalization hierarchy. Each comprises a set of properties and their salience values. A number of perspectives must be defined *a priori* for the objects in a particular domain. The specification of perspectives, just like the specification of an object taxonomy, must be done by a domain expert. Knowledge of useful perspectives in a domain then, is part of the domain expertise.

With this new notion of perspective, when an object is viewed through a particular perspective, the perspective essentially acts as a filter on the properties which that object inherits from its superordinates. That is, properties are inherited with the salience values given by the perspective. Thus properties of the object which are given a high salience rating by the perspective will be highlighted, while those which are given a low salience value or do not appear in the perspective will be suppressed. The density problem is handled by ignoring those superordinate concepts which contribute only attributes suppressed by the current perspective.

4. Using Perspective to Determine Responses

Perspective information can be used with Tversky's similarity metric to help determine alternative queries to a query containing a misconception. To see how this works, consider a domain containing the following three objects with the attributes shown:

Money Market Certificates

Maturity: 3 months
 Denominations: \$1,000
 Issuer: Commercial Bank
 Penalty for Early Withdrawal: 10%
 Purchase Place: Commercial Bank
 Safety: Medium

Treasury Bills

Maturity: 3 months
 Denominations: \$1,000
 Issuer: US Government
 Purchase Place: Federal Reserve
 Safety: High

Treasury Bond

Maturity: 7 years
 Denominations: \$500
 Issuer: US Government
 Penalty for Early Withdrawal: 20%
 Purchase Place: Federal Reserve
 Safety: High

and the following perspectives:

Savings Instruments

Maturity — high
 Denominations — high
 Safety — medium

Issuing Company

Issuer — high
 Safety - high
 Purchase Place - medium

Notice that the perspective of Savings Instruments highlights Maturity and Denominations, and somewhat highlights Safety. This indicates that when people are discussing securities as savings instruments, they are most interested in how long their money will be tied up and in what denominations they can save their money. The perspective of Issuing Company, on the other hand, highlights different attributes. When securities are discussed from this perspective, things like who the issuer of the security is and how safe a security issued from that company is, become important.

Suppose the perspective is Savings Instruments and the user says:

U. What is the penalty for early withdrawal on a Treasury Bill?

This query indicates that the user has a misconception since s/he has attributed a property to Treasury Bills that they do not have. One reasonable correction to the query would contain an alternative query which replaces Treasury Bills with another object that has the property specified and is similar to Treasury Bills. The system may reason that both Money Market Certificate* and Treasury Bonds have the penalty specified, and so check to see if either of these objects is similar to Treasury Bills. Notice that the Savings Instruments perspective highlights attributes common to Treasury Bills and Money Market Certificates (they have the same Maturity and Denominations), as well as attributes disjoint to Treasury Bills and Treasury Bonds (they have different Maturity and Denominations). Using these salience values, the similarity metric will find that Money Market Certificates are very similar to Treasury Bills while Treasury Bonds are very different. Thus Money Market Certificates will be deemed a probable object of confusion and the following correction may be offered:

S. Treasury Bills do not have a penalty for early withdrawal. Were you thinking of a Money Market Certificate?

Notice that if the perspective had instead been Issuing Company, which highlights attributes common to Treasury Bills and Treasury Bonds and disjoint to Treasury Bills and Money Market Certificates, the most reasonable response would be:

S. Treasury Bills do not have a penalty for early withdrawal. Were you thinking of a Treasury Bond?

Selecting the appropriate perspective is in itself a difficult question which is currently under investigation and will be reported in [5]. Certainly important in the selection procedure will be the attributes that have entered into the conversation so far: these attributes should be of fairly high salience in the selected perspective. Other cues to the selection process include the objects under discussion, the superordinates which contribute the attributes under discussion to these objects, and the current goals of the user.

5. Conclusion

In this paper we have seen that a context-dependent similarity metric is needed in order to respond adequately to misconceptions involving the properties of an object. Such a metric has been suggested and a notion of perspective has been introduced to account for some of the contextual information required by the metric. These notions have been shown to account for differences in the way a particular misconception is best corrected in two different circumstances.

6. Acknowledgements

I would like to thank Julia Hirschberg, Aravind Joshi, Martha Pollack, Ethel Schuster, and Bonnie Webber for their many comments and discussions concerning the direction of this research and the content and style of this paper.

7. References

- [1] Bobrow, D. G. and Winograd, T. "An Overview of KRL, a Knowledge Representation Language." *Cognitive Science* 1, 1 (January 1977), 3-46.
- [2] Carberry, Sandra M. Understanding Pragmatically Ill-Formed Input. 10th International Conference on Computational Linguistics & 22nd Annual Meeting of the Association of Computational Linguistics, Coling84, Stanford University, Ca., July, 1984, pp. 200-206.
- [3] Carbonnell, Jaime R. & Collins, Allan M. Mixed-Initiative Systems For Training and Decision-Aid Applications. Tech. Rept. ESD-TR-70-373, Electronics Systems Division, Laurence G. Hanscom Field, US Air Force, Bedford, Ma., November, 1970.
- [4] Grosz, B. Focusing and Description in Natural Language Dialogues. In *Elements of Discourse Understanding*, A. Joshi, B. Webber & I. Sag, Ed., Cambridge University Press, Cambridge, England, 1981, pp. 85-105.
- [5] McCoy, K.F. Correcting Object-Related Misconceptions. 1985. Forthcoming University of Pennsylvania doctoral thesis
- [6] McKeown, K. . *Generating Natural Language Text in Response to Questions About Database Structure*. Ph.D. Th., University of Pennsylvania, May 1982.
- [7] Sidner, C. L. Focusing in the Comprehension of Definite Anaphora. In *Computational Models of Discourse*, Michael Brady and Robert Berwick, Ed, MIT Press, Cambridge, Ma, 1983, pp. 267-330.
- [8] Tversky, A. ■ Features of Similarity." *Psychological Review* 84 (1977), 327-352.
- [9] Weiner, E. Judith. "A Knowledge Representation Approach to Understanding Metaphors." *Computational Linguistics* 19, 1 (January - March 1984), 1-14.