The ROMPER System: Responding to Object-Related Misconceptions using Perspective¹

Kathleen F. McCoy
Dept. of Computer and Information Sciences
University of Delaware
Newark, De. 19716

Abstract¹

As a user interacts with a database or expert system, s/he may reveal a misconception about the objects modeled by the system. This paper discusses the ROMPER system for responding to such misconceptions in a domain independent and context sensitive fashion. ROMPER reasons about possible sources of the misconception. It operates on a model of the user and generates a cooperative response based on this reasoning. The process is made context sensitive by augmenting the user model with a new notion of object perspective which highlights certain aspects of the user model due to previous discourse.

1 Introduction

A study of transcripts of expert-user dialogues reveals that users often exhibit misconceptions about the objects modeled in a domain. This paper describes the ROMPER system (Responding to Object-Related Misconceptions using PERSpective) which is able to respond to certain classes of these misconceptions in a principled manner. In doing so the system sheds light not only on the process of correcting misconceptions, but also on issues in natural-language generation, user models, and modeling certain contextual effects by a "filtering" of the knowledge representation.

The ROMPER system functions as a part of a natural-language interface to a database or expert system. Input to ROMPER is a specification that a misconception has been detected. In this work a *misconception* is defined to be some discrepancy between what the system believes (i.e., what is contained in the system knowledge base) and what the user believes (as exhibited through the conversation). The system knowledge base includes an object taxonomy and knowledge about object attributes and their possible values.

Several factors may influence the structure and content of responses to queries that reveal misconceptions. These include the goals of the conversational participants. If the misconception is not important to these goals, the response may not address the misconception or may address it only minimally. ROMPER is concerned with correcting misconceptions that are important to the current goals of the conversational participants and is thus concerned with generating a maximal

response. This response is aimed at eliminating the discrepancy between what the user believes and what the system believes by bringing the user's knowledge into line with the system's. This means that the system must not only give the user the correct information, but must present it in such a way so as to have the user adopt that information. ROMPER has a user model available to aid in this task. The user model constitutes what the system believes the user believes about the domain. It contains the same kind of information as is contained in the system's knowledge base — an object taxonomy and information about objects' attributes and their values. The content of the user model, however, may be very different from the content of the system's knowledge base. For instance, it may contain less information than is contained in the system knowledge base, or it may contain some information that is inconsistent with the system knowledge base. The user model will not, however, contain more information than is contained in the system knowledge base since the system is assumed to be an expert in the domain.

In an attempt to respond to a misconception in a natural way, the system operates on the model of the user attempting to find certain structural configurations which might indicate support for the misconception. If one of the configurations is found, then a response is generated that refutes the found support. ROMPER is specifically concerned with responding to two kinds of misconceptions: those involving an object's classification (which I call misclassifications) and those involving an object's attributes (which I call misattributions). Certain structural configurations have been identified indicating possible support for both kinds of misconceptions. Each identified configuration has a response strategy associated with it which may be instantiated to respond to the misconception. The whole process is made context sensitive by a new notion of *object perspective* which acts to filter the user model, highlighting those aspects which are made important by previous dialogue, while suppressing others. The filtering gained by object perspective allows the same misconception by the same user to be responded to differently in different contextual situations.

Output from ROMPER is a formal specification of a response. This specification is then input to the MUMBLE system [McD80] which, using a dictionary and grammar supplied by Robin Karlin [Kar85], produces actual English text.

¹Much of this work was done while the author was at the University of Pennsylvania and was partially supported by the ARO grant DAA20-84-K-0061 and the NFS grant MCS81-07290.

2 Misconception Responses

The view of natural-language generation taken in this system is the same as that taken in [McK82]. The generation process is seen as consisting of two parts: (1) determining the content and structure of the response and producing a formal message specification, and (2) transforming that specification into actual English text. My work has concentrated on determining the content and structure of a response to a misconception. It attempts to automate the process of deciding what information to include in a response to a misconception by giving the system the ability to reason about certain classes of misconceptions and typical ways of correcting misconceptions in one of the identified classes. This should be contrasted with the *a priori* listing of misconceptions and responses found in most existing systems that handle misconceptions ([SC80], [BB78], and [Woo84]).²

The form of the responses generated by ROMPER derived from an analysis of transcripts of human conversational partners. These transcripts revealed that responses to statements containing misconceptions often include more than a simple denial of the wrong information. This is particularly true in circumstances where the misconception is about something important to the current goals of the participants. In addition to denying the information involved in the misconception, many misconception responses include both the corresponding correct information, and additional justification for the denial and correction given. The justification often involves refuting faulty reasoning that may have led the user to the misconception.

While it may seem that the kinds of faulty reasoning that the user may be using to arrive at a misconception are limitless, the transcript analysis revealed a surprisingly small number of misconception support relations that are refuted by the human experts. In addition, these few misconception support relations can be couched in terms of a knowledge base (KB) structure rather than its content. Thus a system reasoning on a model of the user might look for such relations in a domain independent fashion. If one is found, information refuting the misconception support might be included in the corrective response.

To see this, let us examine the number of ways that a human expert was found to correct one of the misconception types handled by ROMPER: *misclassifications*. The strategies used by the human experts to respond to a misclassification can be exemplified by the number of possible responses to the following misconception:

U. I thought whales were fish.

R1. No, they are mammals. You may have thought they were fish because they are fin-bearing and live in the water. However, they are mammals since, (while fish have gills) whales breathe through lungs and feed their young with

milk.

R2. No, they are mammals. You may have thought they were fish since they are like the fish, sharks, in that both are large aquatic creatures and both scare people. However, whales are mammals since, (while fish have gills) whales breathe through lungs and feed their young with milk.

R3. No, they are mammals.

Before analyzing each one of these in detail, let us first note their similarities. Each of the above strategies can be seen as consisting of three parts: (1) a denial of the incorrect classification, (2) a statement of the correct classification, and (3) an offering of justification for the denial/correction pair. This three part strategy is, in fact, typical of all of the responses found in the transcript analysis. Notice that the denial of the incorrect information and the offering of the corresponding correct information is the same in each of the sample responses given. What distinguishes one kind of response strategy from another is the kind of justification given in each case. Responses R1 and R2 offer two different kinds of justification while R3 offers no justification.

Given that examples of each of the above three kinds of responses were found in the transcripts, we must ask what causes one to be used in preference to another in a particular situation. One explanation is that different beliefs about what the user believes trigger the use of each strategy. Notice that each strategy can be seen as refuting a different kind of support for the misconception. My claim is that a speaker may choose a strategy depending on the support that he/she believes the user may be using to come up with the misconception. Let us take each strategy in turn, examine what beliefs might have led to the use of that strategy, and then investigate how this information might be used by a system to generate responses to misconceptions.

3 Using Response Strategies

The justification in R1 consists in the expert conceding properties that whales have in common with fish (fin-bearing and water-living), and overriding that conceded information with properties that distinguish whales from fish. The use of this strategy by an expert might be explained by the expert believing that the user believes that whales and fish are similar, and that that similarity may have led to the misclassification. An expert having these beliefs might very well find it reasonable to concede that the similarity between whales and fish does indeed exist, but then go on to show that that similarity is not enough to classify whales as fish. S/he may do this, as above, by offering properties that make whales mammals instead of fish.

Given that this analysis might explain a human's response to a misconception, we might have a computer system

²[Woo84] represents a departure from the canned response in that she is concerned with appropriately structuring a response to reflect a certain tutoring style.

adopt this strategy to respond to a misconception in a natural way. First, the information included in a response like R1 can be captured in a response schema as shown below. R1 can be seen as an instantiation of this schema where OBJECT is instantiated with whale, POSITED with fish, and REAL with mammal. The shared attributes are instantiated in the obvious way.

```
((deny (classification OBJECT POSITED))
 (state (classification OBJECT REAL))
 (concede (share-attributes OBJECT
          POSITED
          ATTRIBUTES1))
 (override (share-attributes -----
            POSITED
            ATTRIBUTES2))
 (override (share-attributes OBJECT
          REAL
          ATTRIBUTES3)))
```

The above schema is called the "like-super" schema because it is used by ROMPER when the user exhibits a misconception by wrongly classifying some OBJECT as a POSITED superordinate and when ROMPER determines that a probable reason for the misclassification is that the user believes that the OBJECT and the POSITED superordinate are similar to each other. The schema captures a response like R1 by specifying a denial of the incorrect classification, a statement of the correct classification, and then an offering of justification. The justification is in the form of conceding the similarity that may have led to the misclassification (e.g., the shared attributes), but overriding that conceded information with attributes that are not shared by the OBJECT and the POSITED superordinate but instead distinguish the two.

It should be pointed out that this schema encodes two kinds of information: a domain-independent specification of the *content* of each proposition included in the response (e.g., an object classification or shared attributes between objects), as well as information about the *rhetorical force* or *communicative role* played by each proposition (e.g., a denial or statement or conceded information). The content specification is derived from the transcript analysis. The rhetorical force is derived from both the transcript analysis and from work done by [McK82], [MT83], and [Man84] who have developed theories about the role that a proposition can play in a discourse. The goal in using such a schema is to have a specification of a response that may be filled in with information from the user model and that, when instantiated, contains enough rhetorical information to be turned into a cohesive English text by a tactical component. The schema above meets both of these requirements.

The justification included in R2 is also in the form of a concede/override pair. However, in the case of R2, rather than concede a similarity between whales and all fish, a similarity between whales and some subset of fish (i.e., the sharks) is conceded. The use of this response might be explained by the

expert believing that the user believes whales and sharks to be similar and salient at this point in the discourse. The expert might imagine the user to have reasoned: "I don't know how to classify whales, but I do know that they are similar to sharks and I know that sharks are fish. Perhaps whales are fish as well."

This analysis was used in developing ROMPER by associating a schema based on responses like R2 with a user model configuration showing a similarity between the misclassified object and some descendent of the posited superordinate. The schema is termed the "like-some-super" schema and is shown below:

```
((deny (classification OBJECT POSITED))
 (state (classification OBJECT REAL))
 (concede (similarity
          OBJECT
          DESCENDENT
          (share-attributes OBJECT
           DESCENDENT
           ATTRIBUTES1)))
 (override (share-attributes OBJECT
          REAL
          ATTRIBUTES2)))
```

Response R3 can be thought of as the degenerate strategy since it contains no justification for the denial/correction pair. ROMPER instantiates the schema corresponding to R3 when neither of the two above mentioned knowledge base configurations can be found in the user model.

So far this paper has concentrated on misclassifications. ROMPER also handles misconceptions involving an object's attributes. The transcript analysis revealed three correction strategies for misattributions as exemplified by the following responses:

- U. What is the interest rate on this stock?
- R4. Stock doesn't have an interest rate. Were you thinking of a bond?
- R5. Stock doesn't have an interest rate. Did you mean dividend?
- R6. Stock doesn't have an interest rate.

ROMPER employs three correction schema to handle misattributions; one for each of the response strategies shown. R4 can be seen as an instantiation of ROMPER's wrong-object schema. This schema offers an object which has the attribute involved in the misconception that the user may have either confused with the misconception object or made a bad analogy from. It is instantiated when an object is found that has the attribute involved in the misattribution and is similar to the misconception object.

R5 exemplifies the wrong-attribute schema which offers an attribute that the object involved in the misconception does

have. This response is used when there is reason to believe the user may have confused the attribute involved in the misconception with a similar attribute that the object does have. ROMPER uses the schema when the misconception object has an attribute that is similar to the attribute involved in the misconception.

As is the case with the misclassifications, there is a “degenerate” schema for misattributions. This schema contains no justification for the correction and is exemplified by R6.

In summary, a study of transcripts of humans responding to misconceptions reveals a great deal of regularity in the way misconceptions about objects are corrected. One can abstract a small number of response strategies for each of the various knowledge base features that might be involved in a misconception. Each of these strategies can be seen as refuting a different kind of support that the user may have for the information involved in the misconception. These strategies are captured as schemas in the ROMPER system and each schema is associated with a domain independent description of the kind of support it refutes. ROMPER, when faced with a misconception, operates on a model of the user looking for evidence for one of the identified kinds of support. If enough evidence is found, the response to the misconception is generated by instantiating the corresponding schema.

4 Effects of Context

The above section outlined a method for correcting misconceptions. While the method does seem to be appealing, at first glance it seems to have a major flaw. It does not seem to take into account the role that previous context plays in correcting misconceptions. The responses given by the human experts were very context dependent. In two different contexts a human expert might choose to correct the same misconception by the same user in two different ways. For example, in response to the misconception exhibited by “I thought whales were fish”, an expert might choose R1 in one context and R2 in another. How can this be explained if the process described above is used to respond?

I claim that the process of correcting misconceptions is context sensitive not because the process changes with context, but because what the process works on changes with context. In particular the piece of the user model that is analyzed in looking for possible sources of the misconception changes with context. Instead of doing the user model analysis on a flat representation containing everything that the user knows at equal levels of importance, the analysis is done on a model that has been highlighted by previous discourse. Previous discourse serves to highlight certain aspects of the user model while suppressing others. Different highlighting resulting from different previous discourse may cause the user model analysis to conclude that different support had been used for the misconception and therefore cause a different response strategy to be selected. *Object perspective* is a notion which can be used to model this

contextual effect.

5 Object Perspective

In this section I introduce a new notion of object perspective as an augmentation to a standard semantic network representation. Before introducing this notion let us first examine what we want this notion to account for.

The notion of object perspective has previously been discussed in the literature. It can be likened to the “point of view” one takes on an object in a particular discussion. From a particular point of view certain characteristics of the object seem more important than others. For instance, a particular building may be discussed from the point of view of being someones home on the one hand, and from the completely different point of view of being an architectural work on the other. The two different views of the same building cause different groups of attributes to be important. It is this highlighting of a whole group of attributes that must be explained. Notice that it could not be explained by a focusing mechanism which highlights attributes which have been mentioned in the preceding discourse because many of the highlighted attributes may not have been explicitly mentioned. What needs to be captured is the feeling that each view calls to mind a “precompiled” set of attributes that seem to be important while that view is in effect.

An attempt to explain this effect has been made by defining object perspective as viewing an object as a member of one superordinate when, in fact, it may have many superordinates ([Gro77], [BW77], and [TWF*82]). The highlighting is achieved through a limited inheritance mechanism. An object inherits only those attributes contributed by the one superordinate deemed “in perspective”. Thus, when a building is viewed as an architectural work, for example, it inherits only those attributes associated with the concept architectural-work in the generalization hierarchy. Any attributes that it might inherit from other superordinates (e.g., home) are ignored. While this notion is intuitively appealing, in practice it is problematic (see [McC85] for details) and is unable to handle some additional effects which intuitively should be handled by object perspective. Two of these effects will be discussed here.

During the course of a conversation it is usually the case that more than one object will be discussed. When this happens, usually the same kinds of things are discussed about the objects. In essence, a particular highlighting of attributes (or point of view) seems to be in force during the conversation. Yet, this highlighting is applied to different objects. What seems to be happening is that the conversational partners are viewing an entire group of objects from the same perspective. This cannot be accounted for by the previous definition of object perspective unless each of the objects under discussion can be said to have the same superordinate.

A second effect which is not accounted for by the above definition, yet seems to hinge on object perspective, has to do

with the heightened importance of some objects during a discourse. For instance, in the responses R1-R3 above, the correct classification of whale was given as mammal. It is the case, however, that whales are cetaceans and cetaceans are mammals. If the expert above thought that U. knew about cetaceans, why wasn't cetaceans given as the correct classification? Since there was no preceding discourse given in this case, some default context would have to be in force. Apparently, in this context cetacean did not seem important enough to mention. Yet in other contexts, one can imagine cetacean being given as the correct classification even though it had not yet been explicitly referred to in the preceding discourse. The importance of the object cetacean seems to have something to do with the current perspective from which objects are being viewed. The previous definitions of object perspective do not address this issue.

5.1 Perspective: Definition and Representation

I claim that all of the above criteria can be met by a simple notion of object perspective which has the following properties:

First, instead of tying perspective into the generalization hierarchy of objects as has been done in the past, the new notion of perspective will be independent of that hierarchy. "Perspectives" which can be taken on the objects in the domain will be defined and will sit in a structure which is orthogonal to the generalization hierarchy.

Second, the number of such perspectives that need be defined for the objects in a given domain of discourse is small and finite. Moreover, any given domain object may be viewed from any one of several perspectives defined for that domain. As it turns out, it will make more sense to view some of the objects in the domain through some perspectives and not others, but this is a feature of perspectives which will be taken advantage of later.

Third, each perspective comprises a set of attributes with associated salience values. It is these salience values that dictate which attributes are highlighted and which are suppressed.

Fourth, one such perspective is designated *active* at a particular point in the discourse.

This notion of object perspective works as follows. An object or group of objects is still said to be viewed through a perspective. In particular any object which is accessed by the system is viewed through the current *active* perspective. However, instead of dictating which attributes an object inherits, the active perspective affects the salience values of the attributes that an object possesses (either directly or inherited through the generalization hierarchy). The active perspective essentially acts as a filter on an object's attributes – raising the salience of and thus highlighting those attributes which have a high salience rating in the active perspective, and lowering the salience of and thus suppressing those attributes which are either given a low salience value or do not appear in the active perspective.

The importance of an object in a discourse is determined by the salience values given to the attributes it possesses. The idea is that the whole becomes highlighted by having its parts highlighted. Thus, during a discussion in which the active perspective highlights many attributes contributed by the object "cetacean" in our generalization hierarchy, cetacean will be seen as an important object. If, on the other hand, none of the attributes associated with cetacean are highlighted, then that object will be suppressed.

This notion of object importance realizes the intuitive notion that it makes "more sense" to view some objects through particular perspectives than others. It makes more sense to view an object through perspectives that highlight many of the object's attributes and thereby make the object more dominant. Notice that we can see a certain amount of symmetry here. The perspective determines the salience of an object's attributes and the object's importance; the object and its attributes determines how likely the object is to be viewed from a particular perspective.

5.2 Using Perspective

A model of a particular domain would include the usual object taxonomy containing all of the objects in the domain and all of the attributes those objects possess. So in our fish-mammal domain we would have sharks as a kind of fish with attributes like "scare-people" and "large-aquatic-creature". In addition, all of the attributes of fish would also be represented and sharks would inherit those attributes as well.

In addition to the object taxonomy, we must build a separate structure containing the perspectives that can be taken on the domain objects. One perspective we might imagine defining for the fish-mammal domain would be the "body-characteristics" perspective. In this perspective attributes like "fin-bearing", "have-gills", and "breathe-through-lungs" would be given high salience and thus highlighted. Other attributes would be suppressed by this perspective.

Another perspective that might be defined for the fish-mammal domain might be the "common-people's-perception" perspective. This perspective might highlight attributes like "large-aquatic-creatures" and "scare-people". Other attributes, like "have-gills" and "fin-bearing" might be suppressed by this perspective.

ROMPER uses the highlighting from object perspective in two ways. First, during the user model analysis it uses the information to check for user model configurations which might indicate particular kinds of support for a misconception. Section 2 introduced two user model configurations which were associated with response schemas. The like-super schema was associated with a user model configuration that indicated that the user believed the misclassified object was *like* the posited superordinate. The like-some-super schema was associated with a user model configuration that indicated that the user believed the misclassified object was *like* some descendent of the posited

superordinate. Notice that both of these user model configurations hinge on a similarity assessment between objects. The similarity metric used by ROMPER is one that is based on the objects' common and disjoint attributes which takes attribute salience into account [Tve77]. This metric will be discussed below. Since the similarity metric takes attribute salience into account and attribute salience is effected by object perspective, the active perspective can influence the selection of a misconception response schema.

Second, ROMPER uses the highlighting from object perspective to instantiate the selected response schema. It attempts to do this using only attributes deemed important by the current perspective.

5.3 Object Similarity

As was mentioned above, the object similarity metric used by ROMPER must be sensitive to context. To date, most AI systems that use object similarity use a metric that is based on distance in the generalization hierarchy. Such a metric is not context sensitive.

The ROMPER system uses a similarity metric based on work done in [Tve77] which allows contextual information to be taken into account. Tversky's metric, called a *contrast model*, is based on the common and disjoint features/properties of the objects involved. Suppose we have 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 . Tversky's measure can be expressed as:

$$s(a, b) = \theta f(A \cap B) - \alpha f(A - B) - \beta f(B - A)$$

for some θ, α , and $\beta \geq 0$.

In the above equation θ, α , and β are parameters which alter the importance of each piece of the equation. The function f maps over the features and yields a salience rating for each. In essence, the contrast model 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 particular feature involved (determined by the function f) and the importance of each piece of the equation (determined by θ, α , and β) may change with context.

In order to use the metric, we must come up with values for the functions in the equation. Tversky suggests that the θ, α , and β functions might be affected by the relative prominence of objects a and b in the discourse. If a is relatively more important, then function θ and α should be greater than β resulting in the attributes of the more prominent object having a greater influence over the similarity assessment. While I would conjecture that information about the focus of the discourse [Gro81], [Sid83], [GJW83] might give an indication of an object's prominence and would therefore be useful in setting the values of θ, α , and β , in this work I have assumed a value of 1 for the θ, α , and β and have concentrated on setting the f function.

In the ROMPER system the f function has been set using the salience values returned after the knowledge base has been filtered through object perspective. Using this setting of f the same two objects may be seen as very similar when the active perspective highlights attributes that the objects have in common and suppresses those that are disjoint between them. On the other hand, the same two objects may be seen as very different when the active perspective suppresses attributes that they have in common and highlights those that are disjoint between them.

This similarity metric is used by ROMPER in deciding which schema to use to respond to a particular misconception. Suppose that ROMPER must respond to the misconception "I thought a whale was a fish" when the active perspective is the "body-characteristics" perspective defined above. Recall that this perspective highlighted attributes like fin-bearing, have-gills, and breathe-through-lungs. Under this perspective, attributes common to whales and all fish are highlighted. Using a Tversky-like similarity metric this highlighting causes whales and fish to be seen as similar. ROMPER would thus respond using the like-super schema producing a response similar to R1.

If, on the other hand, the same misconception were encountered when the perspective was "common-people's-perception", the attributes that whales and all fish have in common would not be highlighted. Rather, attributes like scare-people and large-aquatic-creatures shared with just a subset of fish, the sharks, would be highlighted. Under these conditions, the similarity metric would return a low similarity rating for whales and all fish (and thus the "like-super" schema would not be applicable), but a high similarity rating for whales and sharks. Thus, the "like-some-super" schema would be used to produce a response similar to R2 above.

One can imagine how other perspectives might make neither the "like-super" nor the "like-some-super" schemas applicable, causing the "no-support" schema to be used.

5.4 Choosing the Active Perspective

In order for the notion of object perspective to be truly beneficial, there must be a mechanism for choosing the active perspective based on previous discourse. While this topic is still very much open to investigation, some preliminary research has revealed several factors that might influence the choice of active perspective.

Perhaps one of the most influential pieces of information useful in choosing a perspective is the user's current goal. In [MWM85] the user's goal completely determines which perspective is active. In their work each perspective which can be taken on the domain objects is indexed by potential goals. Thus, once the system has determined what the user's goal probably is, it has also determined what perspective the user has probably taken on the domain objects.

While it is true that the user's goal is a good source of information to use to determine the probable perspective,

other factors may also influence this choice. These include the attributes and objects mentioned so far in the dialogue. The mentioned attributes are obviously thought to be important and one would therefore expect them to be given a fairly high salience rating in the active perspective. Thus, the choice of active perspective can be narrowed down to those in which the mentioned attributes appear with high salience.

By the same token, the objects mentioned so far in the dialogue can also give a clue concerning the active perspective. One would expect that the active perspective would deem these objects important. Therefore the system might look for perspectives that give high salience ratings to many of the attributes associated with objects that have been mentioned in the discourse.

In this section I have identified several factors which influence the choice of active perspective. This choice, however, is a question which remains as an open research topic. Still unanswered are questions such as: When does a perspective change? How long is a perspective active? Is there a relationship between a discourse unit [GS85] and perspective? Is there any structure to the space of perspectives that would put constraints on moving from one active perspective to another? These questions must be taken up in future research on perspective.

5.5 An Example

In this section an example is given which indicates how the choice of perspective influences how a misconception may be corrected. Recall that in correcting a misattribution one of the correction schemas used by ROMPER called for a similar object to be offered as a possible object of confusion. A study of transcripts reveals, however, that this schema may be instantiated in different ways depending on the context. Consider the following dialogue:

U. I am interested in investing in some securities to use as savings instruments. I want something short-term and I don't have a lot of money to invest so the instrument must have small denominations. I am a bit concerned about the penalties for early withdrawal. What is the penalty on a T-bill?

S. Treasury Bills don't have a penalty. Were you thinking of a Money Market Certificate?

In this case the money market certificate was seen as being similar to the treasury bill and therefore included in the response. A different object might be used in a different context. Consider:

U. I am interested in investing in some securities. Safety is very important to me, so I would probably like to get something from the government. I am a bit concerned about the penalties for early withdrawal. What is the penalty on a T-bill?

S. Treasury Bills don't have a penalty. Were you thinking of a Treasury Bond?

The difference in these two responses can be explained by different perspectives being taken on the objects. Suppose that our knowledge base contains the following objects and attributes in the financial securities domain.

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

The following perspectives might be reasonable for the domain (here we are assuming salience values from low salience of 0 to high salience of 1):

Savings Instruments

Maturity - 1.0
denominations - 1.0
safety - 0.5

Issuing Company

issuer - 1.0
safety - 1.0
purchase-place - 0.5

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 company is and how stable an investment in the company is, become important. Other attributes of the securities are ignored (recall that attributes not mentioned in the perspective get assigned a low salience rating).

Consider how perspective might effect the misconception response. Given the discourse preceding the utterance containing the misconception in our first dialogue, it is reasonable to assume that the perspective of "Savings Instruments" is the active perspective at the time of the misconception utterance.³ A system attempting to respond to this misconception might proceed by attempting to instantiate the wrong object schema described above. Recall that this schema is applicable when there is a similar object which has the property involved in the misconception. The system might collect all objects which have the attribute in question and then test their similarity with the object involved in the misconception. In our knowledge base there are two objects which have the attribute involved in the misconception: Money Market Certificates and T-Bonds.

Suppose the attributes of these objects were assigned the salience values given by the Savings Instrument perspective. Applying the Tversky metric using the salience values attached by this perspective (and assuming a value of 1 for θ , α , and β) we get:

$$s(\text{T-Bill, MM-Cert}) = f(\text{maturity, denom}) - f(\text{safety}) \\ = 2 - .5 = 1.5 \implies \text{high similarity}$$

$$s(\text{T-Bill, T-Bond}) = f(\text{safety}) - f(\text{maturity, denom}) \\ = .5 - 2 = -1.5 \implies \text{low similarity}$$

With these calculations the system would choose the Money Market Certificate as the possible object of confusion and respond:

S. Treasury Bills don't have a penalty. Were you thinking of a Money Market Certificate?

Contrast the above calculations with calculations that might occur given a different active perspective. The discourse preceding the misconception utterance in the second example suggests the active perspective of "Issuing Company". Using the salience values attached by this perspective the similarity metric would produce the following calculations:

$$s(\text{T-Bill, MM Cert}) \\ = f() - f(\text{issuer, safety, purchase}) \\ = 0 - 2.5 = -2.5 \implies \text{low similarity}$$

$$s(\text{T-Bill, T-Bond}) \\ = f(\text{issuer, safety, purchase}) - f() \\ = 2.5 - 0 = 2.5 \implies \text{high similarity}$$

In this case a reasonable response by the system would be:

S. Treasury Bills don't have a penalty. Were you thinking of a Treasury Bond?

As the examples show, changes in the active perspective can account for the same misconception being responded to in two different ways.

³ROMPER does not calculate the active perspective. Instead, it is input to the system.

6 Conclusion

If we want our natural-language front-ends to database or expert systems to mimic human behavior, they must have the ability to handle misconceptions. This paper has described a methodology for handling object-related misconceptions and has illustrated this methodology on misconceptions involving object misclassifications.

The proposed method for responding to object-related misconceptions requires associating response schemas with certain structural configurations of the user model. The response schemas described in this paper were derived from a corpus of transcripts and were associated with user model configurations that would explain their use by a human expert in responding to a misconception.

A system might use the pairing of strategies to configurations upon encountering an object-related misconception by searching the user model for one of the identified configurations. If one was found, the associated schema could be instantiated to generate a corrective response.

The context-dependent nature of responses to misconceptions is accounted for not by having the *process* of correcting misconceptions change with context, but rather by having *what the process works on* change with context. A new notion of object perspective was introduced as an augmentation to a flat semantic network representation of the user. Object perspective provides a highlighting of the user model as a result of previous discourse. This resulting user model was shown sufficient for accounting for different responses being given to the same misconception in different situations.

7 Acknowledgements

I would like to thank my advisors, Aravind Joshi and Bonnie Webber, for their many helpful comments throughout the course of this work. Special thanks also go to Sandra Carberry and Martha Pollack for their comments on various drafts of this paper.

References

- [BB78] J.S. Brown and R.R. Burton. Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science*, 2(2):155-192, 1978.
- [BW77] D. G. Bobrow and T. Winograd. An overview of krl, a knowledge representation language. *Cognitive Science*, 1(1):3-46, January 1977.
- [GJW83] B. Grosz, A.K. Joshi, and S. Weinstein. Providing a unified account of definite noun phrases in discourse. In *Proceedings of the 21st Annual Meeting*, pages 44-50, Association for Computational Linguistics, Cambridge, Mass, June 1983.

- [Gro77] B. Grosz. *The Representation and Use of Focus in Dialogue Understanding*. Technical Report 151, SRI International, Menlo Park Ca., 1977.
- [Gro81] B. Grosz. Focusing and description in natural language dialogues. In B. Webber A. Joshi and I. Sag, editors, *Elements of Discourse Understanding*, pages 85–105, Cambridge University Press, Cambridge, England, 1981.
- [GS85] B. Grosz and C. Sidner. Discourse structure and the proper treatment of interruptions. In *Proceedings of the 1985 Joint Conference on Artificial Intelligence*, IJCAI85, Los Angeles, Ca., August 1985.
- [Kar85] Robin Karlin. *Romper Mumbles*. Technical Report, University of Pennsylvania, May 1985.
- [Man84] W. C. Mann. Discourse structures for text generation. In *Proceedings of Coling84*, pages 367–375, Association for Computational Linguistics, Stanford University, Ca., July 1984.
- [McC85] K.F. McCoy. *Correcting Object-Related Misconceptions*. PhD thesis, University of Pennsylvania, December 1985.
- [McD80] D. D. McDonald. *Natural Language Production as a Process of Decision Making Under Constraint*. PhD thesis, MIT, 1980.
- [McK82] K. McKeown. *Generating Natural Language Text in Response to Questions About Database Structure*. PhD thesis, University of Pennsylvania, May 1982.
- [MT83] W. C. Mann and S. A. Thompson. *Relational Propositions in Discourse*. Technical Report ISI/RR-83-115, ISI/USC, November 1983.
- [MWM85] K. McKeown, M. Wish, and K. Matthews. Tailoring explanations for the user. In *Proceedings of the 1985 Conference*, Int'l Joint Conference on Artificial Intelligence, Los Angeles CA, August 1985.
- [SC80] A.L. Stevens and A. Collins. Multiple conceptual models of a complex system. In Pat-Anthony Federico Richard E. Snow and William E. Montague, editors, *Aptitude, Learning, and Instruction*, pages 177–197, Erlbaum, Hillsdale, N.J., 1980.
- [Sid83] C. L. Sidner. Focusing in the comprehension of definite anaphora. In Michael Brady and Robert Berwick, editors, *Computational Models of Discourse*, pages 267–330, MIT Press, Cambridge, Ma, 1983.
- [Tve77] A. Tversky. Features of similarity. *Psychological Review*, 84:327–352, 1977.
- [TWF*82] F. Tou, M. Williams, R. Fikes, A. Henderson, and T. Malone. Rabbit: an intelligent database assistant. In *Proceedings of AAAI-82*, pages 314–317, AAAI, Carnegie-Mellon University, August 1982.
- [Woo84] Beverly P. Woolf. *Context Dependent Planning in a Machine Tutor*. PhD thesis, University of Massachusetts, May 1984.