A New Strategy for Providing Definitions In Task-Oriented Dialogues

Margaret H. SARNER
Department of Computer Science
University of Delaware
Newark, Delaware 19716 U.S.A.

Sandra CARBERRY
Department of Computer Science
University of Delaware
Newark, Delaware 19716 U.S.A.

Abstract

Definitions may be made up of one or more components, which correspond to strategic predicates. The selection of which components to use in giving a definition in a task-oriented dialogue depends heavily on the needs of the user. The selection strategy we present involves weighting possible strategic predicates and the propositions used to fill them at multiple points throughout an ongoing dialogue and at the actual time of giving the definition. Weighting will be influenced by a model of the user's domain knowledge, task-related plans and goals, and receptivity to the different kinds of information that could be presented. An utterance can then be produced that incorporates the most important information while adhering to common rhetorical practices.

1 Introduction

In the course of ongoing task-oriented expert-consultation dialogues, many occasions arise in which the expert must provide a definition. In this paper we will present a new strategy for a computer expert to use in giving definitions in a way that is most helpful to the individual user.

The strategy relies on a dynamically inferred model of the user's domain knowledge, task-related plans and goals, and receptivity to different kinds of information. It constructs a definition by weighting both the strategic predicates that might comprise a definition and the propositions that might be used to fill the strategic predicates. These weights are used to order what might be said according to its anticipated usefulness to the user. Rules can then be used to produce an utterance that incorporates the most important information while adhering to common rhetorical practices. This strategy reflects our overall hypothesis that beliefs about the appropriate content of a definition should guide selection of a rhetorical strategy, instead of the choice of a rhetorical strategy determining content.

Section 2 describes situations in task-oriented dialogues in which definitions are called for. Section 3 identifies three characteristics that differentiate definitions provided by experts during task-oriented dialogues from those provided in response to isolated requests for definitions, and argues that the choice of a rhetorical strategy should be made on the basis of being able to include in the definition those features deemed most important. Section 4 proposes a Principle of Usefulness as a guideline for selecting information to include in definitions. Section 5 discusses strategic predicates. Section 6 presents an overview of our strategy for weighting predicates and propositions and then ranking what might be said according to its usefulness to the user.

2 Definition Situations

In its simplest form, a definition-giving dialogue consists of an information-seeker asking "*What is an X?*" and an information-provider saying "*An X is a ....*" In actual practice, however, there are many ways a definition can be requested and many ways the request can be responded to by the information-provider. In order to identify the characteristics of definition-giving dialogues, we have analyzed transcripts of novice-expert dialogues from a variety of domains, including student/advisor dialogues, recipe-providing dialogues, taxpayer/tax-agent dialogues, and radio talk shows in which callers sought expert advice on investments and real estate.[1] This section describes definition-giving situations identified in this study.

An expert may give a definition either in response to a user's request or spontaneously. Occasions for providing definitions arise most obviously when the user asks a question of the form "*What is ...?*" or "*What is the significance of ...?*" The question doesn't have to be explicit, however, as illustrated in the exchange below, which is an excerpt from a money-management talk show transcript:

E: "I'd like to see you put that into two different Southern utilities."

U: "Southern utilities?"

As shown in [Carberry 1985], such elliptical fragments are often intended to elicit clarification and explanation of the repeated term.

In addition to giving definitions in response to a request by the user, the expert may provide a definition as part of correcting a user misconception [McCoy 1986], or may generate definitions spontaneously. There are several reasons an expert may give spontaneous definitions. He may provide additional definitional information to justify use of a concept. He may think it likely that the user doesn't know about the entity being introduced. The expert may want to ensure that he and the user are working with the same definition. The statement below is an example of a spontaneous definition from a recipe-giving dialogue:

E: "You use a spring-form pan — the kind that allows you to separate the bottom and the sides once you have prepared your dish."

3 Definitions in Task-Oriented Dialogues

McKeown [McKeown 1985] studied definitions in the context of requests for information about the objects modeled by a database system. She claimed that humans have mutually known conventions for organizing information and providing definitions, and that a natural language system should make use of these strategies in producing explanations. Given a definition request, her TEXT system selected a rhetorical strategy based on the information available. The rhetorical strategy was then responsible for selecting the information to be incorporated into the definition. TEXT treated requests for definitions as isolated queries, whereas we are interested in definitions generated in the course of ongoing task-oriented dialogues.

Our analysis of transcripts of naturally occurring interactions indicates that definitions generated in task-oriented dialogues differ significantly from those generated statically or as a result of isolated definition requests. The differences appear to be the result of several factors:

[1] These transcripts were provided by the Computer Science Departments of the University of Pennsylvania and the University of Delaware.

1. In task-oriented dialogues, the information-provider knows something about what the information-seeker is trying to accomplish, and will generate definitions that help the information-seeker achieve his goals. For example, the first response below would be an appropriate definition of baking soda if the information-seeker is baking a cake, whereas the second would be appropriate if he is trying to relieve indigestion.

   E: "Baking soda is an ingredient that, when heated, releases carbon dioxide, thereby causing the mixture to expand in size."

   E: "Baking soda is a substance that, when dissolved in water, produces a chemically basic solution that will counteract acidity."

2. Whereas static definitions or responses to one-shot requests for definitions must assume a generic model for the information-seeker, responses to definition requests during an ongoing dialogue can take into account acquired beliefs about the information-seeker's specific domain knowledge. For example, the information-provider might include an analogy to an entity that the information-seeker is already familiar with, as in the following definition of the course CS106:

   E: "CS106 is like CS105, except that it uses Fortran instead of Pascal and emphasizes scientific applications of computing."

3. Whereas static definitions and responses to one-shot requests for definitions must be generated all at once, dialogue allows the information-provider to produce what he thinks will be an acceptable definition and analyze the information-seeker's response to determine whether to elaborate on the definition. For example, in the following dialogue with a veterinarian about treating a cat with a hyperthyroid condition, the veterinarian (E) provides a definition that he believes will satisfy the information-seeker's needs, then must elaborate on it when the information-seeker's response reveals multiple goals: to improve the condition of the cat and to have medication that is easy to administer.

   E: "Tapazole is a drug that decreases the function of the thyroid."

   U: "How large are the pills?"

If a system carrying on a task-oriented dialogue is to be viewed by the information-seeker as *cooperative*, *intelligent*, and *natural*, it must take the above factors into account. Otherwise, it will not appear to be directed toward the user's goals (uncooperative), will not appear to make use of what the user already knows (unintelligent), and will not appear to take advantage of the fact that the interaction is ongoing, as opposed to one-shot (unnatural).

Our hypothesis is that, instead of using a rhetorical strategy to determine the content of a definition, the system should reason about the user's plans and goals and specific domain knowledge to decide the importance of incorporating individual propositions into the final definition. For this purpose a user model, preferably a dynamically constructed user model, is essential. The choice of a rhetorical strategy should be made on the basis of being able to include into the definition those features deemed most important. Thus beliefs about the appropriate content of the definition should guide selection of a rhetorical strategy, instead of the choice of a rhetorical strategy determining content.

McKeown, Wish, and Matthews [McKeown et al. 1985] addressed some of these issues in their work on an expert system that could provide explanations tailored to users. They described a method for using a model of the user's goals along with pre-built perspectives on the knowledge base to generate appropriate explanations. While they touched on some of the issues that concern us, they took a different approach from the one we are proposing. Their perspectives were built into the domain knowledge base, and their system did not make much use of information available from the system's model of the user's plans and goals. Also, they were concerned with answering *can* and *should* questions, whereas we are interested in definition explanations.

4 Appropriate Content of a Definition

Our analysis of naturally occurring consultation dialogues indicates that definitions can take many forms. They may be made up of one or more of a set of components, which correspond to rhetorical predicates described in [Grimes 1975, Williams 1893, McKeown 1985]. These predicates will be discussed further in Section 5.

Since we are studying cooperative dialogues in which the expert's goal is to help the information-seeker solve his problem, we hypothesize that the expert's overriding concern in selecting the information to include is that the response be as useful as possible to the individual user. Intuitively, to be truly useful to the user, the information must be something he doesn't already know but something relevant that he can understand. Our hypothesis, which appears to explain the definitions occurring in our dialogue transcripts, suggests the following Principle of Usefulness:

Principle of Usefulness

1. The response should be made at a high enough level that it is meaningful to the user.
   (a) Don't say something the user won't understand.
   (b) Don't give information that addresses more detailed aspects of the user's task-related plan than is appropriate for his current focus of attention.
2. The response should be made at a low enough level that it is helpful to the user.
   (a) Don't inform the user of something he already knows.
   (b) Don't give information that is unrelated to the user's goals and task-related plan, or is too general for his current focus of attention in the plan.

Grice [Grice 1975] stated that contributions should be as informative as required for the exchange, but not more informative than required. Paris [Paris 1988] suggested that an answer to a question should be both informative and understandable to the user, based on the user's level of knowledge about the domain of discourse. The Principle of Usefulness formalizes and extends these guidelines for definitions by selecting the appropriate level both in knowledge-related issues (1a, 2a) and in plans and goals (1b, 2b). This Principle will be used whenever a selection of appropriate level of information to fill a predicate is called for.

For example, consider a plant classification hierarchy.

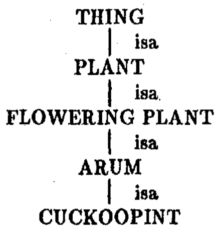

To describe a cuckoopint as an arum would have no meaning to an information-seeker who has never heard of an arum, while defining it as a thing is too general. The useful level of explanation for the information-seeker with no special knowledge of botany is defining a cuckoopint as a flowering plant. In task-oriented dialogues, additional care must be taken to avoid providing extra information that is unrelated to, or too detailed for, the user's current needs. Otherwise, the extra information may lead the user to erroneously assume that the system believes the distinguishing characteristics are important or that the system has mis-identified the aspect of his task on which the user is currently focused.

5 Strategic Predicates

The term *rhetorical predicate* has taken on several meanings in the literature of linguistics and computational linguistics. It has been used to describe relationships ranging from structural to conceptual in nature. Grimes [Grimes 1975] described rhetorical predicates that "relate the kinds of information communicated in discourse with each other." One of his predicates was the Attributive predicate which "adds qualities or color to another predicate as center." Hobbs [Hobbs 1979] chose to use the term *coherence relation* in preference to *rhetorical predicate* to place the emphasis on the coherence between sentential units. McKeown's description of rhetorical predicates [McKeown 1985] implied an association with sentential structure, but in practice the predicates she used, such as Constituency, dealt more with conceptual relationships.

We are using predicates to characterize the components of definitions in terms of relationships between conceptual units. Our predicates relate information about the entity being defined to the entity itself. This relationship is domain-independent and content-independent. For example, our Identification predicate is instantiated by finding in a generalization hierarchy an entity which is an ancestor of the entity being defined. This usage is close to McKeown's, but because of the apparent ambiguity of the term *rhetorical predicates*, we prefer to call the predicates *strategic predicates*, putting emphasis on the motivation of gaining an end (in this case, conveying useful information to the user) rather than on style.

From our study of definitions occurring in actual dialogues, we have identified fourteen distinct predicates that relate to definition content. Each predicate corresponds to a different type of information that can be put into a definition. Although we do not claim that the list is complete or unique, we do believe it is sufficient to generate appropriate definitions in an expert consultation system. Some of our predicates (Identification, Properties, Analogy, Components) are similar to McKeown's. Others (Effect, Prerequisites) are particular to a task-oriented environment.

Associated with each predicate are semantics that indicate how to instantiate it. For example, effect information will be found in the system's library of plans and goals, and property information will be found in the generalization hierarchy. In either case, the system must reason about the particular user's plans and goals in order to determine a proposition's relevance to what the user is trying to accomplish. When an occasion for a definition arises, a given predicate may be filled one or more times. The propositions thus produced are candidates for inclusion in the definition. Since our goal is to select the information that is most important to the user, we associate a measure of significance with each proposition. The significance metrics will be described in Section 6.

In the remainder of this section we will look at three types of definition components in some detail to illustrate how the user model influences selection.

5.1 Identification

Many naturally occurring definitions contain an Identification component. Identification consists of identifying the entity being described as a member of a generic class in a hierarchically structured knowledge base — for example,

E: "Amaretto is a liqueur."

The system's model of the user dictates what superclass from the generalization hierarchy to use in an identification. In order for identification to be helpful to the user, it is necessary that the user have knowledge of the parent category used in making the identification. This condition corresponds to the first part of the Principle of Usefulness. Knowledge of a parent category may not be sufficient, however, to cause that parent category to be given in the definition. If the system's beliefs indicate that the parent category is related to the user's plans and goals, then there is stronger reason to mention it. In the case in which the entity has several parents that the user has knowledge of, plans and goals should be used to select the one (or ones) most appropriate to mention. Suppose, for example, that a digital systems course is cross-listed as both a Computer Science and an Electrical Engineering course.

U: "What is Digital Systems?"

E: "It is a Computer Science course ..."

or

E: "It is an Electrical Engineering course ..."

The choice of answer depends on whether the user model indicates that the user is trying to satisfy Computer Science or Electrical Engineering requirements. A third alternative is

E: "It is both a Computer Science course and an Electrical Engineering course ..."

This response might be given if the model indicates both parent categories play a role in the user's plans and goals.

Following the Principle of Usefulness, the appropriate superclass is the lowest level parent category that would have meaning to the user and be relevant to what the system believes are the user's plans and goals.

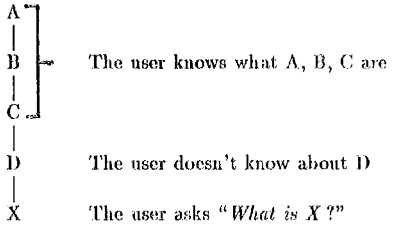
A
B
C
D
X
The user knows what A, B, C are
The user doesn't know about D
The user asks "*What is X* ?"

In the case illustrated above, the expert's Identification answer might be "*X is a C.*" The effect of answering "*X is a D*" would be to cause the user to ask "*What is a D*?" or give up without getting meaningful information. The answer "*X is a B*" would miss the distinguishing features shared by C and X but not by B. If these distinguishing features are not important to the user and would give the false impression that the system believes they are relevant to the user's task, however, a higher level than C should be selected.

5.2 Properties

A Properties response consists of naming characteristics of the entity. These are often expressed in descriptions given by humans as adjectival phrases attached to the Identification of the entity.

E: "A no-load fund is a mutual fund with no sales charge."

E: "Amaretti are crisp Italian almond-flavored macaroons."

In the TEXT system [McKeown 1985], attributes whose values distinguish one sub-type from another are marked in the knowledge base. In task-oriented dialogues, however, an entity's most important distinguishing attributes are not always static but instead may vary depending on the information-seeker's plans and goals. For example, the course Computer, Ethics and Society may have several distinguishing properties, including its content, its substantial writing component, its lack of programming projects, and its scheduling at night through continuing education. An information-seeker whose objective is to earn a BA degree at night while holding a full-time job would consider its scheduling property of interest in differentiating it from other computer science courses, whereas an electrical engineering major seeking a technical elective would probably consider its lack of programming projects of particular significance. Thus, although the properties of an entity are found in the generalization hierarchy, the system's beliefs about the user's plans and goals should play a major role in determining which properties of the entity are most appropriate to include in a definition.

5.3 Operation

An Operation response consists of a description of how something works. Paris [Paris 1988] has demonstrated that explanations given novices in a domain often take the form of process traces. An Operation definition may take the form of process information or steps in implementation. The difference between the two is that the process information is essentially a chain of cause-and-effect occurrences, while the steps in implementation are sequential, but not necessarily causal, as is shown in the example:

U: "Can you tell me what the money market is?"

E: " A money market fund is a group of people getting together — put their money together in a pool and it is invested by professional investors."

As with the Properties predicate, the system's beliefs about the user's plans and goals must be taken into consideration. The expert might identify the need for an Operation explanation in a task-oriented dialogue when the entity being explained appears in a step in a plan the user must carry out to meet a goal. For example, if the user is a traveler asking the expert for help planning a car trip and the expert advises the user to follow a "Trip Tik," the expert should explain how a Trip Tik works if the model of the user indicates lack of familiarity with it. The definitions of baking soda given earlier illustrate a case in which the appropriate Operation explanation depends on the use to which the entity will be put by the information-seeker.

6 Selecting Definition Content

Our strategy assumes a knowledge base consisting of a generalization hierarchy containing domain knowledge, a plan library, and a lexicon. The user model has three components:

1. a model of the user's domain knowledge in the form of markings on the knowledge base showing the pieces with which the user is familiar [Kass 1987],
2. a model of the user's underlying task-related plan and current focus of attention in the plan, given by a context tree [Carberry 1988],
3. a model of how receptive the user is to various kinds of information, given by weightings on strategic predicates.

The first two components will be dynamically updated during the dialogue as shown in [Kass 1987] and [Carberry 1988]. The third component will also be updated dynamically in response to the user's receptivity to types of definitions and his own usage of strategic predicates.

6.1 Weighting Predicates

When a definition occasion arises, a local predicate receptivity model is created. Starting with a copy of the current global weights representing the user's general receptivity to the kinds of information represented by the strategic predicates, as inferred from the preceding dialogue, further adjustments may be made to reflect the appropriateness of the predicates in the particular situation.

The question itself and the level of local domain expertise may cause further weighting of predicates. For example, if the user asks "*What is X?*" where X is an object, the Identification predicate would be more heavily weighted. If X is an action, the Operation predicate would be more heavily weighted. The level of local domain expertise can be ascertained when a definition is requested by looking at the parts of the plan library and generalization hierarchy that contain references to the entity in question. If they are heavily marked with things the user knows, the user can be considered to have a high level of expertise; otherwise, the user will be considered to be a novice. The weights for predicates that have been determined to be appropriate for expert and novice users will then be increased [Paris 1988].

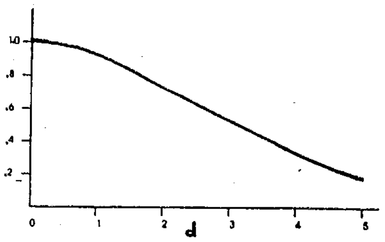

Figure 1: Graph of Relevance Formula

6.2 Weighting Propositions

After predicate weighting has been determined, predicates are filled with information from the knowledge base (generalization hierarchy, lexicon, plans and goals) relevant to the concept being defined. The semantics of each individual predicate dictate where to find the information to fill the predicate. For instance, the Identification and Properties predicates are filled with information found in the generalization hierarchy, and Necessity propositions are drawn from the plans of the user. Some predicates may produce several propositions. For example, an entity may have several properties. For others there might not be any corresponding propositions available.

Selection of propositions depends on both the weights of the possible predicates and a measure of significance of the information that could be used to fill them. Significance reflects where the proposition fits into the system's model of the user's goals and possible plans for accomplishing them (relevance) and what information in the generalization hierarchy has been marked as known by the user (familiarity).

The system's beliefs about the user's underlying task-related plan, as dynamically inferred from the preceding dialogue, are represented in a tree structure called a *context model* [Carberry 1988]. Each node in this tree represents a goal that the user has investigated achieving. Except for the root, each goal in the context model is a descendant of a higher-level goal whose associated plan, found in the system's plan library, contains the lower-level goal. One node in the tree is marked as the current focus of attention and indicates that aspect of the task on which the user's attention is currently centered. The context model may be expanded to arbitrarily many levels of detail by repeatedly replacing non-primitive subgoals with associated plans which themselves contain constituent subgoals.

If pursuing a subgoal in a plan represents a significant shift in focus, it is marked in the plan library as *introducing a new focus domain*. Within the context model, a focus domain of subgoals that are at approximately the same level of focus is generated by expanding the plan associated with a subgoal that introduces the focus domain. As long as this plan is expanded by substituting plans for just those subgoals that do not introduce another new focus domain, the subgoals appearing in the expanded plan are part of the same focus domain.

Our estimate of relevance is based on distance of the part of the context model in which the definition information is found from the current focus of attention in the context model. This distance is measured as the number of shifts in focus domains. If the plan is at the focus of attention, the information derived from it is of very high relevance. If it is in the immediately surrounding focus domain (one shift), the information is still of high relevance. As the number of focus domain shifts increases, the relevance of information in the plans begins to fall off, but as long as a plan has been activated the information found in it is of some relevance. This situation in which relevance remains high close to the focus of attention, but drops off more rapidly as the distance increases, is modeled by an inverse exponential function, as shown in Figure 1. The equation

$$r = e^{-\left(\frac{d}{4}\right)^2},$$

where r is the relevance rating and d is the number of shifts from the current focus of attention, captures the desired features.

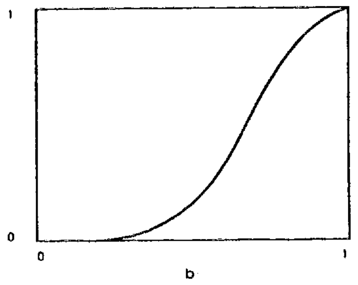

Figure 2: Graph of Familiarity Formula

Currently, our relevance metric treats all shifts among focus domains equally. It may be the case, however, that information in a higher-level plan *h* that led to the current focus of attention is more appropriate to include in a definition than is information extracted from a subplan *s* appearing in an expansion of the current focused plan, even if the two plans, *h* and *s*, represent the same number of shifts from the current focus of attention in the context model. The current focused plan is part of an expansion of *h*, so we know that the user is concerned with accomplishing *h*; therefore, information relevant to *h* may be more significant to the user than information relevant to details of carrying out the current focused plan. This is an issue that we plan to investigate further.

Our measure of familiarity is based on the knowledge the expert believes the user has about the objects, properties, or concepts that could be used in a definition. We are assuming a variant of the user modeling system described by Kass [Kass&Finin 1987], modified so that each node in the knowledge base is marked with a belief factor, ranging in value from 0 to 1, giving the system's level of belief that the user is familiar with the entity. Because of the importance of giving a definition in terms of something the person receiving the definition will understand, an entity known to have meaning to the user (belief factor = 1) should be treated as potentially useful to include, even if it is not germane to the hypothesized goals. If it is not believed strongly that the person is familiar with the entity, however, it is less useful to tie the definition to that entity. Note that since the dialogues under consideration are ongoing, as opposed to one-shot, a definition can include items that the system believes the user is probably familiar with, and the system can wait for the user's response to decide whether the definition was successful. The heuristic described here is modeled by the function shown in Figure 2. The formula

$$f = \frac{e^{6b(2-b)} - 1}{e^6 - 1},$$

where f is the familiarity rating and b is the belief factor, exhibits an appropriate amount of curvature to reflect the rapid drop-off in usefulness as the belief factor decreases.

The last step in computing a measure of significance for a piece of information is to form a weighted combination of the relevance rating and the familiarity rating. Since our primary goal is to provide information that will help the user accomplish a task, our formula for combining the two measures weights significance twice as heavily as familiarity. Our significance metric, then, is

$$S = \frac{2r + f}{3},$$

where S is significance, r is the relevance rating, and f is the familiarity rating.

The following example from a hypothetical travel domain illustrates how propositions are weighted according to significance. The dialogue pertains to planning a trip abroad.

U: "I need to have enough money with me to pay for anything I buy."

E: "You can carry as much as you like in travelers checks."

U: "Travelers checks?"

The first statement causes the have-money plan to be in focus. The have-money plan has subgoals

```
have-convertible-funds ((_agent: person)
                        (_amount1: funds))

have_currency ((_agent: person)
               (_country: country)
               (_amount2: funds)).
```

Suppose that the user's elliptical fragment is interpreted as a request for a definition. Figure 3 shows part of the context model. As a result of the expert's preceding response, the focus of attention is now on the have-convertible-funds plan. Suppose further that the other plans shown are in a focus domain at a distance of 1 from the focus of attention.

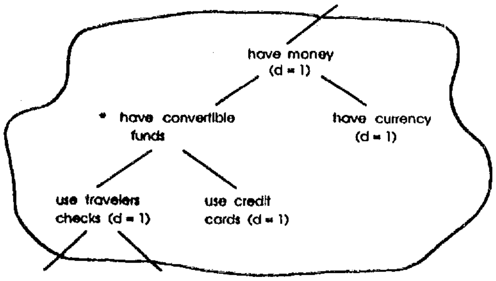

Figure 3: A Portion of the Context Model

The Operation predicate produces the candidate proposition formed from components of the use-travelers-checks subplan (not shown) equivalent to the statement

> "You can buy travelers checks at a bank here and cash them in the currency of the country."

The information comes from the body of the use-travelers-checks subplan, which is at distance d=1 from the focus of attention. Assuming that the expert believes that the user is familiar with the concepts of buying, banks, currency, and cashing things in, we have

$$\begin{aligned} r &= e^{-\left(\frac{d}{4}\right)^2} = e^{-\left(\frac{1}{4}\right)^2} = .939 \\ f &= \frac{e^{6b(2-b)} - 1}{e^6 - 1} = \frac{e^{6(1)} - 1}{e^6 - 1} = 1 \\ S &= \frac{2r + f}{3} = .959 \end{aligned}$$

The Analogy predicate is filled by a reference to a sibling with similar properties, equivalent to

"Travelers checks are like personal checks."

Suppose the belief factor for personal checks is .9 — that is, the expert believes it very likely but is not absolutely certain that the user knows about personal checks. Suppose further that the properties of travelers checks that are similar to those of personal checks appear in plans at a distance of two shifts of focus domain from the focus of attention. In this case we compute

$$\begin{aligned} r &= e^{-\left(\frac{d}{4}\right)^2} = e^{-\left(\frac{2}{4}\right)^2} = .779 \\ f &= \frac{e^{6b(2-b)}-1}{e^6-1} = \frac{e^{5.4(1.1)}-1}{e^6-1} = .942 \\ S &= \frac{2r+f}{3} = .833 \end{aligned}$$

The fact that the first definition component has higher computed significance than the second does not necessarily mean that it will be preferred, however. Recall that weights of candidate propositions must reflect both significance of the information and predicate receptivity.

Once weights have been assigned to the candidate propositions, they are then ranked according to weight and put into categories. There are four categories:

Must Say
Say if Convenient
Say if Needed for Coherence
Do Not Say

The higher weight categories receive the higher-weighted propositions; the lower-weighted propositions go into the lower weight categories. Some categories may be empty.

When all category assignments have been made, the resulting four groups of propositions are passed to an answer generator. Construction of this answer generator is a future project. The generator will take the classes of propositions, find a way to say all of the Must Say propositions and as many as possible of the Say if Convenient propositions, using Say if Needed for Coherence propositions whenever they help the construction of the response. We propose to do this task using rules of combination developed to produce an utterance that adheres to common rhetorical practices that people appear to follow.

7 A Comparison

Our strategy will produce different responses than would current definition systems. For example, consider a request for a definition of amaretti. McKeown's TEXT system would identify the entity and include all based database and distinguishing database attributes, and would produce a definition resembling

> "Amaretti are macaroons. They are made from apricot kernels, have almond flavor, are of Italian origin, and have crisp texture. The most popular brand is made by Lazzaroni and Company."

Our definition module would attempt to pick information appropriate to the individual user. If the user is selecting food items to sell in an international bazaar, it would say

> "Amaretti are Italian macaroons. The most popular brand is made by Lazzaroni and Company."

If the user is making Amaretti Amaretto Chocolate Cheesecake, for which amaretti are an ingredient, however, it would say

> "Amaretti are crisp almond-flavored macaroons."

8 Future Work

Our continuing research will work out additional details of our strategy for providing definitions in task-oriented dialogues. We need to investigate a strategy for dynamically weighting strategic predicates according to the user's perceived receptivity to different kinds of information, and putting this weighting together with our measure of significance for propositions. An answer generator that combines propositions, giving emphasis to including those propositions deemed most important to say, must be designed. This task includes ranking the candidate propositions by weight and combining the most heavily weighted ones in a way that will produce a coherent utterance. Finally, the system must be implemented to test and demonstrate the utility of our definition strategy.

9 Summary

We claim that determining the most important things to say for the individual user is the most significant task in providing definitions in task-oriented dialogues. In this paper we present a new strategy for generating definitions, using a weighting strategy that draws on a dynamically inferred model of the user's domain knowledge, task-related plans, and receptivity to different kinds of information. This strategy reflects our over-all hypothesis that beliefs about the appropriate content of a definition should guide selection of a rhetorical strategy, instead of the choice of a rhetorical strategy determining content. This approach will produce a system that exhibits cooperative, intelligent behavior by providing definitions tailored to the needs of the individual user.

References

Carberry, Sandra. 1985. A Pragmatics Based Approach to Understanding Intersentential Ellipsis. In: *Proceedings of the 23rd Annual Meeting of the Association for Computation Linguistics*, 188–197.

Carberry, Sandra. 1988. Modeling the User's Plans and Goals. *Computational Linguistics Journal*, To Appear.

Grice, H. Paul. 1975. Logic and Conversation. In: P. Cole and J. L. Morgan, Eds., *Syntax and Semantics III: Speech Acts*, Academic Press, N.Y.: 41–58.

Grimes, J. E. 1975. *The Thread of Discourse.* Mouton.

Hobbs, Jerry R. 1979. Coherence and Coreference. *Cognitive Science*, 3:67–90.

Kass, Robert. 1987. *Implicit Acquisition of User Models in Cooperative Advisory Systems.* Technical Report MS-CIS-87-05, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.

Kass, Robert and Finin, Tim. 1987. Rules for the Implicit Acquisition of Knowledge About the User. *Proceedings of the Sixth National Conference on Artificial Intelligence*, 295–300.

McCoy, Kathleen F. 1986. The ROMPER System: Responding to Object-Related Misconceptions Using Perspective. *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, 97–105.

McKeown, Kathleen R. 1985. *Text Generation.* Cambridge University Press.

McKeown, K., Wish, M., and Matthews, K. 1985. Tailoring Explanations for the User. In: *Proceedings of the 1985 Conference*, Int'l Joint Conference on Artificial Intelligence, Los Angeles CA.

Paris, Cecile L. 1988. Tailoring Object Descriptions to a User's Level of Expertise. *Computational Linguistics Journal.*

Williams, W. 1893. *Composition and Rhetoric.* Heath and Company.