

# UNDERSTANDING PRAGMATICALLY ILL-FORMED INPUT

M. Sandra Carberry  
Department of Computer Science  
University of Delaware  
Newark, Delaware 19711 USA

## ABSTRACT

An utterance may be syntactically and semantically well-formed yet violate the pragmatic rules of the world model. This paper presents a context-based strategy for constructing a cooperative but limited response to pragmatically ill-formed queries. Suggestion heuristics use a context model of the speaker's task inferred from the preceding dialogue to propose revisions to the speaker's ill-formed query. Selection heuristics then evaluate these suggestions based upon semantic and relevance criteria.

## I INTRODUCTION

An utterance may be syntactically and semantically well-formed yet violate the pragmatic rules of the world model. The system will therefore view it as "ill-formed" even if a native speaker finds it perfectly normal. This phenomenon has been termed "pragmatic overshoot" [Sondheimer and Weischedel, 1980] and may be divided into three classes:

- [1] User-specified relationships that do not exist in the world model.

EXAMPLE: "Which apartments are for sale?"

In a real estate model, single apartments are rented, not sold. However apartment buildings, condominiums, townhouses, and houses are for sale.

- [2] User-specified restrictions on the relationships which can never be satisfied, even with new entries.

EXAMPLE: "Which lower-level English courses have a maximum enrollment of at most 25 students?"

In a University world model, it may be the case that the maximum enrollments of

-----  
This material is based upon work supported by the National Science Foundation under grants IST-8009673 and IST-8311400

lower-level English courses are constrained to have values larger than 25 but that such constraints do not apply to the current enrollments of courses, the maximum enrollments of upper-level English courses, and the maximum enrollments of lower-level courses in other departments. The sample utterance is pragmatically ill-formed since world model constraints prohibit the restricted relations specified by the user.

- [3] User-specified relationships which result in a query that is irrelevant to the user's underlying task.

EXAMPLE: "What is Dr. Smith's home address?"

The home addresses of faculty at a university may be available. However if a student wants to obtain special permission to take a course, a query requesting the instructor's home address is inappropriate; the speaker should request the instructor's office address or phone. Although such utterances do not violate the underlying domain world model, they are a variation of pragmatic overshoot in that they violate the listener's model of the speaker's underlying task.

A cooperative participant uses the information exchanged during a dialogue and his knowledge of the domain to hypothesize the speaker's goals and plans for achieving those goals. This context model of goals and plans provides clues for interpreting utterances and formulating cooperative responses. When pragmatic overshoot occurs, a human listener can modify the speaker's ill-formed query to form a similar query X that is both meaningful and relevant. For example, the query

"What is the area of the special weapons magazine of the Alamo?"

erroneously presumes that storage locations have an AREA attribute in the REL database of ships [Thompson, 1980]; this is an instance of the first class of pragmatic overshoot. Depending upon the speaker's underlying task, a listener might infer that the speaker wants to know the REMAINING-CAPACITY, TOTAL-CAPACITY, or perhaps even the LOCATION (if "area" is interpreted as referring to "place") of the Alamo's Special Weapons Magazine. In each case, a cooperative participant uses the preceding dialogue and his knowledge of the

speaker to formulate a response that might provide the desired information.

This paper presents a method for handling this first class of pragmatic overshoot by formulating a modified query X that satisfies the speaker's needs. Future research may extend this technique to handle other pragmatic overshoot classes.

Our work on pragmatic overshoot processing is part of an on-going project to develop a robust natural language interface [Weischedel and Sondheimer, 1983]. Mays[1980], Webber and Mays[1983], and Ramshaw and Weischedel[1984] have suggested mechanisms for detecting the occurrence of pragmatic overshoot and identifying its causes. The main contribution of our work is a context-based strategy for constructing a cooperative but limited response to pragmatically ill-formed queries. This response satisfies the user's perceived needs, inferred both from the preceding dialogue and the ill-formed utterance. In particular,

- [1] A context model of the user's goals and plans provides expectations about utterances, expectations that may be used to model the user's goals. We use a context mechanism [Carberry, 1983] to build the speaker's underlying task-related plan as the dialogue progresses and differentiate between local and global contexts.
- [2] Only alternative queries which might represent the user's intent or at least satisfy his needs are considered. Our hypothesis is that the user's inferred plan, represented by the context model, suggests a substitution for the proposition causing the pragmatic overshoot.

## II KNOWLEDGE REPRESENTATION

Our system requires a representation for each of the following:

- [1] the set of domain-dependent plans and goals
- [2] the speaker's plan inferred from the preceding dialogue
- [3] the existing relationships among attributes and entity sets in the underlying world model
- [4] the semantic difference of attributes, relations, entity sets, and functions.

Plans are represented using an extended STRIPS [Fikes and Nilsson, 1971] formalism. A plan can contain subgoals and actions that have associated plans. We use a context tree [Carberry, 1983] to represent the speaker's inferred plan as constructed from the preceding dialogue. Nodes within this tree represent goals and actions which

the speaker has investigated; these nodes are descendants of parent nodes representing higher-level goals whose associated plans contain these lower-level actions. The context tree represents the global context or overall plan inferred for the speaker. The focused plan is a subtree of the context tree and represents the local context or particular aspect of the plan upon which the speaker's attention is currently focused. This focused plan produces the strongest expectations for future utterances.

An entity-relationship model states the possible primitive relationships among entity sets. Our world model includes a generalization hierarchy of entity sets, attributes, relations, and functions and also specifies the types of attributes and the domains of functions.

## III CONSTRUCTING THE CONTEXT MODEL

The plan construction component is described in [Carberry, 1983]. It hypothesizes and tracks the changing task-level goals of a speaker during the course of a dialogue. Our approach is to infer a lower-level task-related goal from the speaker's explicitly communicated goal, relate it to potential higher-level plans, and build the complete plan context as the dialogue progresses. The context mechanism distinguishes local and global contexts and uses these to predict new speaker goals from the current utterance.

## IV PRAGMATIC OVERSHOOT PROCESSING

Once pragmatic overshoot has been detected, the system formulates a revised query QR requesting the information needed by the user. Our hypothesis is that the user's inferred plan, represented by the context model, suggests a substitution for the proposition that caused the pragmatic overshoot. The system then selects from amongst these suggestions using the criteria of relevance to the current dialogue, semantic difference from the proposition in the user's query, and the type of revision operation applied to this proposition.

### A. Suggestion Mechanism

The suggestion mechanism examines the current context model and possible expansions of its constituent goals and actions, proposing substitutions for the proposition causing the pragmatic overshoot. This erroneous proposition represents either a non-existent attribute or entity set relationship or a function applied to an inappropriate set of attribute values.

The suggestion mechanism applies two classes of rules. The first class proposes a simple sub-

stitution for an attribute, entity set, relation, or function appearing in the erroneous proposition. The second class proposes a conjunction of propositions representing an expanded relationship path as a substitution for the user-specified proposition. These two classes of rules may be used together to propose both an expanded relationship path and an attribute or entity set substitution.

### 1. Simple-Substitution Rules

Suppose a student wants to pursue an independent study project; such projects can be directed by full-time or part-time faculty but not by faculty who are "extension" or "on sabbatical". The student might erroneously enter the query

"What is the classification of Dr. Smith?"

Only students have classification attributes (such as Arts&Science-1985, Engineering-1987); faculty have attributes such as rank, status, age, and title. Pursuing an independent study project under the direction of Dr. Smith requires that Dr. Smith's status be "full-time" or "part-time". If the listener knows the student wants to pursue independent study, then he might infer that the student needs the value of this status attribute and answer the revised query

"What is the status of Dr. Smith?"

The suggestion mechanism contains five simple substitution rules for handling such erroneous queries. One such rule proposes a substitution for the user-specified attribute in the erroneous proposition. Intuitively, a listener anticipates that the speaker will need to know each entity and attribute value in the speaker's plan inferred from the domain and the preceding dialogue. Suppose this inferred plan contains an attribute ATT1 for a member of ENTITY-SET1, namely ATT1(ENTITY-SET1, attribute-value), and that the speaker erroneously requests the value of attribute ATTU for a member ent1 of ENTITY-SET1. Then a cooperative listener might infer that the value of ATT1 for entity ent1 will satisfy the speaker's needs, especially if attributes ATT1 and ATTU are closely related.

The substitution mechanism searches the user's inferred plan and its possible expansions for propositions whose arguments unify with the arguments in the erroneous proposition causing the pragmatic overshoot. The above rule then suggests substituting the attribute from the plan's proposition for the attribute specified in the user's query. This substitution produces a query relevant to the current dialogue and may capture the speaker's intent or at least satisfy his needs.

### 2. Expanded Path Rules

Suppose a student wants to contact Dr. Smith to discuss the appropriate background for a new

seminar course. Then the student might enter the query

"What is Dr. Smith's phone number?"

Phone numbers are associated with homes, offices, and departmental offices. Course discussions with professors may be handled in person or by phone; contacting a professor by phone requires that the student dial the phone number of Dr. Smith's office. Thus the listener might infer that the student needs the phone number of the office occupied by Dr. Smith.

The second class of rules handles such "missing logical joins". (This is somewhat related to the philosophical concept of "deferred ostension" [Quine, 1969].) These rules apply when the entity sets are not directly related by the user-specified relation RLU --- but there is a path R in the entity relationship model between the entity sets. We call this path expansion since by finding the missing joins between entity sets, we are constructing an expanded relational path.

Suppose the inferred plan for the speaker includes a sequence of relations

R1(ENTITY-SET1, ENTITY-SETA)  
R2(ENTITY-SETA, ENTITY-SETB)  
R3(ENTITY-SETB, ENTITY-SET2);

then the listener anticipates that the speaker will need to know those members of ENTITY-SET1 that are related by the composition of relations R1, R2, R3 to a member of ENTITY-SET2. If the speaker erroneously requests those members of ENTITY-SET1 that are related by R2 (or alternatively R1 or R3) to members of ENTITY-SET2, then perhaps the speaker really meant the expanded path R1#R2#R3. The path expansion rules suggest substituting this expanded path for the user-specified relation.

We employ a user model to constrain path expansion. This model represents the speaker's beliefs about membership in entity sets. If pragmatic overshoot occurs because the speaker misused a relation

R(ENTITY-SET1, ENTITY-SET2)

by specifying an argument that is not a member of the correct entity set for the relation, then path expansion is permitted only if the user model indicates that the speaker may believe the erroneous argument is not a member of that entity set.

EXAMPLE: "Which bed is Dr. Brown assigned?"

Suppose beds are assigned to patients in a hospital model. If Dr. Brown is a doctor and doctors cannot simultaneously be patients, then path expansion is permitted if our user model indicates that the speaker may recognize that Dr. Brown is not a patient. In this case, our expanded path expression may retrieve the beds assigned to patients of Dr. Brown, if this is suggested by the inferred task-related plan.

To limit the components of path expressions to those relations which can be meaningfully combined in a given context, we make a strong assumption: that the relations comprising the relevant expansion appear on a single path within the context tree representing the speaker's inferred plan. For example, suppose the speaker's inferred plan is to take CS105. Expansion of this plan will contain the two actions

```
Learn-From-Teacher-In-Class(SPEAKER,
                             section,faculty)
such that Teach(faculty,section)
```

```
Obtain-Necessary-Extra-Help(SPEAKER,
                              section,teaching-assistant)
such that Assists(teaching-assistant,section)
```

The associated plans for these two actions specify respectively that the speaker attend class at the time the section meets and that the speaker meet with the section's teaching assistant at the time of his office hours. Now consider the utterance

"When are teaching assistants available?"

A direct relationship between teaching assistants and time does not exist. The constraint that all components of a path expression appear on a single path in the inferred task-related plan prohibits composing Assists(teaching-assistant,section) and Meet-Time(section,time) to suggest a reply consisting of the times that the CS105 sections meet.

#### B. Selection Mechanism

The substitution and path expansion rules propose substitutions for the erroneous proposition that caused the pragmatic overshoot. Three criteria are used to select from the proposed substitutions the revised query, if any, that is most likely to satisfy the speaker's intent in making the utterance.

First, the relevance of the revised query to the speaker's plans and goals is measured by three factors:

- [1] A revised query that interrogates an aspect of the current focused plan is most relevant to the current dialogue.
- [2] The set of higher level plans whose expansions led to the current focused plan form a stack of increasingly more general, and therefore less immediately relevant, active plans to which the user may return. A revised query which interrogates an aspect of an active plan closer to the top of this stack is more expected than a query which reverts back to a more general active plan.
- [3] Within a given active plan, a revised query that investigates the single-level expansion of an action is more expected, and therefore more relevant, than a revised query that investigates details at a much deeper level of expansion.

Second, we can classify the substitution  $T \rightarrow V$  which produced the revised query into four categories, each of which represents a more significant, and therefore less preferable, alteration of the user's query (Figure 1). Category 1 contains expanded relational paths  $R1 * R2 * \dots * Rn$  such that the user-specified attribute or relation appears in the path expression. For example, the expanded path

Treats(Dr. Brown,patient) \* Is- Assigned(patient,room)

is a Category 1 substitution for the user-specified proposition

Is- Assigned(Dr. Brown,room)

CATEGORY	SUBSTITUTION TERM T	SUBSTITUTION VARIABLE V
1	Expanded relational path including the user-specified attribute or relation	User-specified attribute or relation
2	Attribute, relation, entity set, or function semantically similar to that specified by the user	User-specified attribute, relation, entity set, or function
3	Expanded relational path, including an attribute or relation semantically similar to that specified by the user	User-specified attribute or relation
4	Double substitution: entity set and relation semantically similar to a user-specified entity set and relation	User-specified entity set and relation

Figure 1. Classification of Query Revision Operations

contained in the semantic representation of the query

"Which bed is Dr. Brown assigned?"

Category 2 contains simple substitutions that are semantically similar to the attribute, relation, entity set, or function specified by the speaker. An example of Category 2 is the previously discussed substitution of attribute "status" for the user specified attribute "classification" in the query

"What is the classification of Dr. Smith?"

Categories 3 and 4 contain substitutions that are formed by either a Category 1 path expansion followed by a Category 2 substitution or by two Category 2 substitutions.

Third, the semantic difference between the revised query and the original query is measured in two ways. First, if the revised query is an expanded path, we count the number of relations comprising that path; shorter paths are more desirable than longer ones. Second, if the revised query contains an attribute, relation, function, or entity set substitution, we use a generalization hierarchy to semantically compare substitutions with the items for which they are substituted. Our difference measure is the distance from the item for which the substitution is being made to the closest common ancestor of it and the substituted item; small difference measures are preferred. In particular, each attribute, relation, function, and entity set ATTRFENT is assigned to a primitive semantic class:

PRIM-CLASS(ATTRFENT, CLASSA)

Each semantic class is assigned at most one immediate superclass of which it is a proper subset:

SUPER(CLASSA, CLASSB)

We define function  $f$  such that

$f(\text{ATTRFENT}, i+1) = \text{CLASS}$

if PRIM-CLASS(ATTRFENT, CLASSa1)  
and SUPER(CLASSa1, CLASSa2)  
and SUPER(CLASSa2, CLASSa3)  
and ...  
and SUPER(CLASSa1, CLASS)

If a revised query proposes substituting ATTRFENT<sub>new</sub> for ATTRFENT<sub>old</sub>, then

semantic#difference(ATTRFENT<sub>new</sub>, ATTRFENT<sub>old</sub>)

=NIL if there does not exist  $j, k$  such that  
 $f(\text{ATTRFENT}_{\text{new}}, j) = f(\text{ATTRFENT}_{\text{old}}, k)$   
=min  $k$  such that there exists  $j$  such that  
 $f(\text{ATTRFENT}_{\text{new}}, j) = f(\text{ATTRFENT}_{\text{old}}, k)$   
otherwise

An initial set is constructed consisting of those suggested revised queries that interrogate an aspect of the current focused plan in the context model. These revised queries are particularly relevant to the current local context of the

dialogue. Members of this set whose difference measure is small and whose revision operation consists of a path expansion or simple substitution are considered and the most relevant of these are selected by measuring the depth within the focused plan of the component that suggested each revised query. If none of these revised queries meets a predetermined acceptance level, the same selection criteria are applied to a newly constructed set of revised queries suggested by a higher level active plan whose expansion led to the current focused plan, and a less stringent set of selection criteria are applied to the original revised query set. (The revised queries in this new set are not immediately relevant to the current local dialogue context but are relevant to the global context.) As we consider revised queries suggested by higher level plans in the stack of active plans representing the global context, the acceptance level for previously considered queries is decreased. Thus revised queries which were not rated highly enough to terminate processing when first suggested may eventually be accepted after less relevant aspects of the dialogue have been investigated. This relaxation and query set expansion is repeated until either an acceptable revised query is produced or all potential revised queries have been considered.

## V EXAMPLES

Several examples are provided to illustrate the suggestion and selection strategies.

### [1] Relation or Entity Set Substitution

"Which apartments are for sale?"

In a real-estate model, single apartments are rented, not sold. However apartment buildings, condominiums, townhouses, and houses are for sale. Thus the speaker's utterance contains the erroneous proposition

For-Sale(apartment)

where apartment is a member of entity set APARTMENT.

If the preceding dialogue indicates that the speaker is seeking temporary living arrangements, then expansion of the context model representing the speaker's inferred plan will contain the possible action

Rent(SPEAKER, apartment)  
such that For-Rent(apartment)

The substitution rules propose substituting relation For-Rent from this plan in place of relation For-Sale in the speaker's utterance.

On the other hand, if the preceding dialogue indicates that the speaker represents a real estate investment trust interested in expanding its holdings, an

expansion of the context model representing the speaker's inferred plan will contain the possible action

Purchase(SPEAKER,apartment-building)

where apartment-building is a member of entity set APARTMENT-BUILDING. Purchasing an apartment building necessitates that the building be for sale or that one convince the owner to sell it. Thus one expansion of this Purchase plan includes the precondition

For-Sale(apartment-building)

The substitution rules propose substituting entity set APARTMENT-BUILDING from this plan for the entity set APARTMENT in the speaker's utterance.

## [2] Function Substitution

"What is the average rank of CS faculty?"

The function AVERAGE cannot be applied to non-numeric elements such as "professor". The speaker's utterance contains the erroneous proposition

AVERAGE(rank,fn-value)  
such that Department-Of(faculty,CS)  
and Rank(faculty,rank)

If the preceding dialogue indicates that the speaker is evaluating the CS department, then an expansion of the context model representing the speaker's inferred plan will contain the possible action

Evaluate-Faculty(SPEAKER,CS)

The plan for Evaluate-Faculty contains the action

Evaluate(SPEAKER,ave-rank)  
such that ORDERED-AVE(rank,ave-rank)  
and Department-Of(faculty,CS)  
and Rank(faculty,rank)

If a domain D of non-numeric elements has an explicit ordering, then we can associate with each of the n domain elements an index number between 0 and n-1 specifying its position in the sorted domain. The function ORDERED-AVE appearing in the speaker's plan operates upon non-numeric elements of such domains by calculating the average of the index numbers associated with each element instead of attempting to calculate the average of the elements themselves. The substitution rules propose substituting the function ORDERED-AVE from the speaker's inferred plan for the function AVERAGE in the speaker's utterance. ORDERED-AVE and AVERAGE are semantically similar functions so the difference measure for the resultant revised query will be small.

## [3] Expanded Relational Path

"When does Mitchel meet?"

A university model does not contain a relation MEET between FACULTY and TIMES. However, faculty teach courses, present seminars, chair committees, etc., and courses, seminars, and committees meet at scheduled times. The speaker's utterance contains the erroneous proposition

Meet-Time(Dr. Mitchel,time)

If the preceding dialogue indicates that the speaker is considering taking CS105, then an expansion of the context model representing the speaker's inferred plan will contain the action

Earn-Credit-In-Section(SPEAKER,section)  
such that Is-Section-Of(section,CS105)

Expansion of the plan for Earn-Credit-In-Section contains the action

Learn-From-Teacher-In-Class(SPEAKER,  
section,faculty)  
such that Teach(faculty,section)

and the plan for this action contains the action

Attend-Class(SPEAKER,place,time)  
such that Meet-Place(section,place)  
and Meet-Time(section,time)

The two relations Teach(Dr. Mitchel,section) and Meet-Time(section,time) appear on the same path in the context model. Therefore the path expansion heuristics suggest the expanded relational path

Teach(Dr. Mitchel,section)\*Meet-Time(section,time)

as a substitution for the relation

Meet-Time(Dr. Mitchel,time)

in the user's utterance. Only one arc is added to produce the expanded relational path and it contains the user-specified relation Meet-Time, so the difference measure for this revised query is small.

## VI RELATED WORK

Erik Mays[1980] discusses the recognition of pragmatic overshoot and proposes a response containing a list of those entity sets that are related by the user-specified relation and a list of those relations that connect the user-specified entity sets. However he does not use a model of whether these possibilities are applicable to the user's underlying task. In a large database, such responses will be too lengthy and include too many irrelevant alternatives.

Kaplan[1979], Chang[1978], and Sowa[1976] have investigated the problem of missing joins between entity sets. Kaplan proposes using the shortest relational path connecting the entity sets; Chang proposes an algorithm based on minimal spanning trees, using an a priori weighting of the arcs; Sowa uses a conceptual graph (semantic net) for constructing the expanded relation. None of these present a model of whether the proposed path is relevant to the speaker's intentions.

## VII LIMITATIONS AND FUTURE WORK

Pragmatic overshoot processing has been implemented for a domain consisting of a subset of the courses, requirements, and policies for students at a University. Our system assumes that the relations comprising a meaningful and relevant path expansion will appear on a single path within the context tree representing the speaker's inferred plan. This restricts such expansions to those communicated via the speaker's underlying inferred task-related plan. However this plan may fail to capture some associations, such as between a person's Social Security Number and his name. This problem of producing precisely the set of path expansions that are meaningful and relevant must be investigated further. Other areas for future work include:

- [1] Extensions to handle relationships among more than two entity sets
- [2] Extensions to the other classes of pragmatic overshoot mentioned in the introduction.
- [3] Extensions to detect and respond to queries which exceed the knowledge represented in the underlying world model. We are currently assuming that the system can provide the information needed by the speaker.

## VIII CONCLUSIONS

The main contribution of our work is a context-based strategy for constructing a cooperative but limited response to pragmatically ill-formed queries. This response satisfies the speaker's perceived needs, inferred both from the preceding dialogue and the ill-formed utterance. Our hypothesis is that the speaker's inferred task-related plan, represented by the context model, suggests a substitution for the proposition causing the pragmatic overshoot and that such suggestions then must be evaluated on the basis of relevance and semantic criteria.

## ACKNOWLEDGMENTS

I would like to thank Ralph Weischedel for his encouragement and direction in this research

and for his suggestions on the style and content of this paper and Lance Ramshaw for many helpful discussions.

## REFERENCES

1. Carberry, S., "Tracking User Goals in an Information-Seeking Environment", Proc. Nat. Conf. on Artificial Intelligence, Washington, D. C., 1983
2. Chang, C. L., "Finding Missing Joins for Incomplete Queries in Relational Data Bases" IBM Res. Lab., RJ2145, San Jose, Ca., 1978
3. Fikes, R. E. and N. J. Nilsson, "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving", Artificial Intelligence 2, 1971
4. Kaplan, S. J., "Cooperative Responses from a Portable Natural Language Data Base Query System", Ph.D. Diss., Univ. of Pennsylvania, 1979
5. Mays, E., "Failures in Natural Language Query Systems: Applications to Data Base Query Systems", Proc. Nat. Conf. on Artificial Intelligence, Stanford, 1980
6. Quine, W. V., "Ontological Relativity" in Ontological Relativity and Other Essays, Columbia University Press, New York 1969
7. Ramshaw, L. A. and R. M. Weischedel, "Problem Localization Strategies for Pragmatic Processing in Natural Language Front Ends", Proc. of 9th Int. Conf. on Computational Linguistics, 1984
8. Sondheimer, N. K. and R. M. Weischedel, "A Rule-Based Approach to Ill-Formed Input", Proc. 8th Int. Conf. on Computational Linguistics, 1980
9. Sowa, J. F., "Conceptual Graphs for a Data Base Interface", IBM Journal of Research and Development, July 1976
10. Thompson, B. H., "Linguistic Analysis of Natural Language Communication with Computers", Proc. 8th Int. Conf. on Computational Linguistics, 1980
11. Webber, B. L. and E. Mays, "Varieties of User Misconceptions: Detection and Correction", Proc. 8th Int. Joint Conf. on Artificial Intelligence, Karlsruhe, West Germany, August 1983
12. Weischedel, R. M. and N. K. Sondheimer, "Meta-Rules as a Basis for Processing Ill-Formed Input", (to appear in American Journal of Computational Linguistics, Vol. 9, #3, 1983)