

USING PLAUSIBLE INFERENCE RULES IN DESCRIPTION PLANNING

Alison Cawsey*

Computer Laboratory, University of Cambridge
New Museum Site, Pembroke St, Cambridge, England.

ABSTRACT

Current approaches to generating multi-sentence text fail to consider what the user may *infer* from the different statements in a description. This paper presents a system which contains an explicit model of the inferences that people may make from different statement types, and uses this model, together with assumptions about the user's prior knowledge, to pick the most appropriate sequence of utterances for achieving a given communicative goal.

INTRODUCTION

Examples, analogies and class identification are used in many explanations and descriptions. Yet current text generation techniques all fail to tackle the problem of *when* an example, analogy or class is appropriate, *what* example, analogy or class is best, and exactly what the user may infer from a given example, analogy or class. McKeown, for example, in her *identification* schema (given in figure 1) includes the 'rhetorical predicates' *identification* (as an instance of some class), *analogy*, *particular-illustration* and *attributive* (McKeown, 1985). From each of these, different information could be inferred by the user. In a human explanation they might be used to efficiently convey a great deal of information about the object, or to reinforce some information about an object so it may be better recalled. Yet in McKeown's schema based approach the only mechanism for selecting between these different explanation options is the

*This work was carried out while the author was at the department of Artificial Intelligence, University of Edinburgh, funded by a post doctoral fellowship from the Science and Engineering Research Council. Thanks to Ehud Reiter, Paul Brna and to the anonymous reviewers for helpful comments.

Identification (class & attribute/function)
{Analogy/Constituence/Attributive/Renaming/
Amplification}*
Particular-Illustration/Evidence+
{Amplification/Analogy/Attributive}
{Particular-Illustration/Evidence}

Note: '{}' indicates optionality, '/' alternatives, '+' that item may appear 1-n times, '*' 0-n times.

Figure 1: McKeown's *identification* schema [McKeown 85]

initial pool of knowledge available to be conveyed, and *focus* rules, which just enforce some local coherence on the discourse. A particular example or analogy could perhaps be selected using the functions interfacing the rhetorical predicates to the domain knowledge base, but this is not discussed in the theory.

More recently, Moore has included examples, analogies etc. in her text planner (Moore, 1990). She includes planning operators to *describe-by-superclass*, *describe-by-abstraction*, *describe-by-example*, *describe-by-analogy* and *describe-by-parts-and-use*. Two of these are illustrated in figure 2. But again there are no principled ways of selecting which strategy to use (beyond, for example, possibly selecting an analogy if the analogous concept is known), and the *effect* of each strategy is the same - that the relevant concept is 'known'. In reality, of course, the detailed effects of the different strategies on the hearer's knowledge will be very different, and will depend on their prior knowl-

```

(define-text-plan-operator
 :NAME describe-by-example
 :EFFECT (BEL ?hearer (CONCEPT ?concept))
 :CONSTRAINTS (AND (ISA ?concept OBJECT)
                  (IMMEDIATE-SUBCLASS
                   ?example ?concept))
 :NUCLEUS ((FORALL ?example
             (ELABORATE-CONCEPT-EXAMPLE
              ?concept ?example)))
 :SATELLITES nil)

(define-text-plan-operator
 :NAME describe-by-analogy
 :EFFECT (BEL ?hearer (CONCEPT ?concept))
 :CONSTRAINTS
  (AND (ISA ?concept OBJECT)
        (ANALOGOUS-CONCEPT
         ?analogy-concept ?concept)
        (BEL ?hearer (CONCEPT
                    ?analogy-concept)))
 :NUCLEUS (INFORM ?speaker ?hearer
           (SIMILAR ?concept
                    ?analogy-concept))
 :SATELLITES ((CONTRAST ?concept
                    ?analogy-concept)))

```

Figure 2: Moore's example and analogy text planning operators

edge. Failing to take this into account results in possible incoherent dialogues which don't address the speaker's real communicative goals.

The rest of this paper will present an approach to the problem of selecting between different statement types in a description, based on a set of inference rules for *guessing* what the hearer could infer given a particular statement. These guesses are used to guide the choice of examples, analogies, class identification and attributes given particular goals, and influence how the user model is updated after these kinds of statements are used. The paper first describes the overall framework for explanation generation. This is followed by a brief discussion of the inference rules and knowledge representation used, and a number of examples where the system is used to generate leading descriptions of bicycles. The approach is intended to be complementary to existing approaches which emphasise the coherence of the text, and could reasonable be

combined with these.

OUTLINE OF EXPLANATION 'PLANNER'

The system described below¹ aims to show how plausible inference rules may be used to guide explanation planning given different communicative goals. The basic approach is to find some set of possible utterances, and select the one which – assuming that the user makes certain plausible inferences – contributes most to the stated communicative goal. This process is repeated until some terminating condition is met, such as the communicative goal being satisfied.

This explanation 'planning' strategy is a kind of heuristic search, using a modified best-first search strategy. The search space consists of the space of all possible utterance sequences, and the heuristic scoring function assesses how far each utterance would contribute to the communicative goal. Because this gives a potentially very large search space, only certain utterances are considered at each point. Currently these are constrained to be those which appear to make *some* contribution to the communicative goal – for example, the system might consider describing an object as an instance of some class if that class had some attributes which contributed to the target state. These possible utterances are then scored by using the plausible inference rules to predict what might reasonably be inferred by the user from this statement, given his current knowledge, and comparing that with the communicative goal.

For example, if the communicative goal is for the user to have a positive impression of the object, and the system knows of some feature which the user believes is desirable in an object, then the system may select utterances which allow the user to plausibly infer this feature given their current assumed knowledge about this and other objects.

The search space is defined by the range of possible utterance types. Currently the following types (and associated plausible inference procedures) are allowed, where there may be many possible statements about a given object of each type:

¹Referred to from now on as the GIBBER system – Generating Inference-Based Biased Explanatory Responses.

- Identification, as an instance (or sub-class) of some class.
- Similarity, given some related object with many shared attributes².
- Examples, of instances or sub-classes.
- Attributes of that object.

The selection of possible utterances, and their scoring (given the probable inferences which might be made) depends on the communicative goal set. In the current system, given some object to describe, two different types of communicative goal may be set. The system may either be given an explicit set of attribute values which should be inferrable from the generated description, or it can be given a 'property' that the inferrable attributes should have. This property can be, for example, that the user believes the attribute value to be a desirable one, where an 'evaluation form' similar to Jameson's (1983) is used to rate different values. Where a set of attribute values are given these can be either specific values, or value ranges.

This approach uses a set of rules which may be used to propose a possible move/statement (given the target/communicative goal), a set of rules which may be used to guess what would be inferred or learned from that statement, given the assumed current state of the user's knowledge, and a scoring function which assesses how far the 'guessed at' inferences would contribute to the target. Statements are generated one at a time, with currently³ the only relation between the utterances being enforced by the common overall communicative goal and by the fact that the statements are selected to incrementally update the user's model of the object described.

Using plausible inference rules in this way is undoubtedly error-prone, as assumptions about the user may be wrong and not all hearers will make the expected inferences. However, it is certainly better than ignoring these inferences entirely. So long as the user can ask follow-up questions in an explanatory dialogue (e.g., Cawsey, 1989; Moore, 1990) any such errors are not crucial.

²Note that full analogies, where a complex mapping is required between two conceptually distinct objects, are currently not possible in the system.

³Adding further coherences relations and global strategies may be the subject of further work.

INFERENCE RULES AND KNOWLEDGE REPRESENTATION

For this approach to text planning to be effective, the rules used for guessing what the reader might infer should correspond as far as possible to human plausible inference rules. There are a relatively small number of AI systems which attempt to model human plausible inferences (compared with those attempting to model efficient learning strategies in artificial situations). Zuckerman (1990) uses some simple plausible inference rules in her explanation system, in order to attempt to block incorrect plausible inferences, while a more comprehensive model of human plausible reasoning is provided by Collins and Michalski (1989). This latter theory is concerned with how people make plausible inferences given *generalisation*, *specialisation*, *similarity* and *dissimilarity* relations between objects, using a large number of certainty parameters to influence the inferences. The theory assumes a representation of human memory based on *dynamic hierarchies*, where, for example, given the statement colour(eyes(John))=blue then colour, eyes, John and blue would all be objects in some hierarchy. The theory is used to account for the plausible inferences made when people guess the answer to questions given uncertain knowledge.

The GIBBER system uses inference rules somewhat differently to Collins' and Michalski's. Whereas they are concerned with the competing inferences which may be made from existing knowledge to answer a single question, the GIBBER system is concerned with mutually supporting inferences from multiple given relationships in order to build up a picture of an object. So, although the basic knowledge representation and relationship types (apart from dissimilarity) are borrowed from their work, the actual inference rules used are slightly different.

It should be possible to use the inference rules to incrementally update a representation of what is currently known about an attribute, where generalisation, similarity and specialisation relationships may all contribute to the final 'conclusion'. In order to allow such incremental updates, the representation used in Mitchell's version space learning algorithm is adopted (1977), where each attribute has a pointer to the most specific value that attribute could take, and to the most gen-

eral value, given current evidence. Positive examples (or *generalisation* relationships) are used to generalise the specific value (as in Mitchell's algorithm)⁴ while class identification (specialisation) is used to update the general value using the inherited attributes. Similarity transforms are done by first finding a common context for the transform (a common parent object), and then transferring those attributes which belong to that context which are not ruled out by current evidence. Explicit statement of attribute values fix the attribute value, but further evidence may be used to increase the certainty of any value.

The system also allows for other kinds of domain specific inference rules to be defined - for example, if a user has just been told that a bike has derailleur gears, a rule may be used to show that the user could probably guess that the bike had between 5 and 21 gears. The different kinds of inference rules are used to incrementally update the representation of the user's assumed knowledge of the object and the scoring function, discussed in the previous section, will compare that assumed knowledge of the object with the target.

The knowledge representation is based on a frame hierarchy describing the objects in the domain, where the slot values may point to other objects, also in some hierarchy. In figure 4 a small section of a knowledge base of different kinds of bicycle is illustrated, along with some simple hierarchies of attribute values. In the GIBBER system separate hierarchies are defined for the system's and for the user's assumed knowledge, where the latter is initialised from a user stereotype and updated following each query and explanation.

Of course, the knowledge representation and inference rules described in this section are by no means definitive - there is no implied claim that people really use these rules rather than others in learning from descriptions. They simply provide a starting point for exploring how explanation generation may take into account possible learning and inference rules, and thus better select statements in a description given knowledge of the domain and of the user's knowledge.

⁴Note that Collins' and Michalski's theory does not appear to allow multiple examples to be used by generalising the inferred values.




Figure 3: Partial Bicycle Hierarchies

EXAMPLE DESCRIPTIONS

This section will give two examples of how descriptions of bicycles may be generated using this approach. We will assume that the system's knowledge includes the hierarchy given in figure 4, and (for simplification) the user's knowledge includes all the items except the 'Cascade', but includes the fact that Alison's bike has shimano indexed gears. The first example will show how the system will select utterances to economically convey information given some target attribute values, while the second will show how biased descriptions may be generated given a specification of the desired property of inferrable attributes.

Suppose the user requests a description of the Cascade and that the communicative goal set by the system (by some other process) is to convey the following attributes:

```

type_of(saddle) = anatomic
type_of(tires) = knobby
weight = 31lb
number_of(gears) = 18
type_of(gears) = shimano_index

```

There are many possible statements which could be made about the Cascade. The user knows Alison's bike, so this example could be mentioned. It could be described as an instance of a mountain bike, or just as a bicycle; a comparison could be made with the Trek-800; or any one of the bikes attributes could be mentioned. In this case if it is identified as an instance of a mountain bike the system guesses that the user could infer the first two attributes, which gives the highest score given the target⁵. A comparison with the Trek-800 also gives two possible inferrable attributes, (though one incorrect value, which is currently allowed), and this is the next choice. Finally the system informs the user of the number of gears, blocking the incorrect inference in the previous utterance. The resulting short description is the following⁶:

"The Cascade is a kind of mountain bike.
It is a bit like the Trek-800.
It has 18 gears."

If the scoring function is changed so that it is biased further towards highly certain inferences, rather than efficient presentation of information, then given the same communicative goal the description may end up as an explicit list of all the attributes of the bike, or in a less extreme case, a class identification and three explicit attributes. This scoring function therefore allows for further variation in descriptions, given a communicative goal, and different scoring functions should be used depending on the type of description required.

Suppose now that the same bike is to be described, but the communicative goal is that the user has a positive impression of the Cascade. If the user regards it to be good for a bike to be black with 21 shimano index gears then the following description will be generated.

⁵The scoring function compares the plausibly inferred information with the target, preferring more certain inferences, and inferences bring the knowledge of the object closer to the target (given the attribute value hierarchy). For example, an inference that the bike had 18-21 gears, or an uncertain inference that it had 18, would be given a lower score than a certain inference that it had 18 gears. The total score is the sum of the scores of each possibly inferred value.

⁶Of course this description would be more coherent if a higher level *compare-contrast* relation was used to generate the last two inferences, with resulting text: "It is a bit like the Trek-800 but has 18 gears.". Allowing these higher level strategies within an inference-based approach is the subject of further work.

"The Cascade is a bit like the Trek-800.
Alison's bike is a Cascade.
The Cascade has Shimano Index Gears."

Here the system evaluates each statement by comparing the plausible *inferences* against an evaluation form (Jameson, 1983). The evaluation form describes how far different attribute values are appreciated by different classes of user. Instead of comparing inferred values with some target attribute values the scoring function will score each against the evaluation form. For example, the first utterance (comparison with the Trek-800) is selected because the attributes which might be plausibly inferred from this statement by this user are rated highly on the evaluation form for that class of user. In this case the system assumes that this type of user will prefer a bike with a large number of indexed gears. Of course, one of the plausible inferences which can be made will be incorrect (the fact the Cascade has 21 gears). The system is not required to block such false inferences if they contribute to its goals (though the ethics of generating such leading descriptions might be doubted!).

It should be clear from this that the descriptions generated by the system are very sensitive to the assumptions about the user's prior knowledge, and the inference rules and the scoring function used, as well as to the communicative goal set. There is much possibility for error (and further research required) in each of these. However, the approach still seems to provide the potential for generating improved descriptions, and provides a new principled way of making choices in a description which is absent in schema-based (and RST-based) approaches. It gives a simple example of how, given a model of how people update their beliefs, utterances may be strategically generated to change those beliefs.

CONCLUSION

This paper has discussed how, by anticipating the user's inferences, better explanations may be generated and assumptions about the user's knowledge updated in a more principled way. Although there are problems with the approach - particularly the difficulty of reliably predicting the user's inferences - it seems to provide a more principled way of selecting certain utterance types than existing multi-sentence text generation systems. Other question

answering systems have attempted to simulate the user's inferences in order to block false inferences (Joshi et al., 1984; Zuckerman, 1990), and particular inferences have been considered in lexical choice (Reiter, 1990) and in generating narrative summaries (Cook et al., 1984). However, it has not been used previously as a general technique for selecting between different options in a description.

Considering what is implicitly conveyed in different types of description may also begin to explain some of the empirically derived results used in other systems. For example, the GIBBER system generally chooses to begin a description with class identification or with a comparison, as most information may be inferred from these (compared with mentioning specific attributes). This may be one of the principles influencing the organisation of the discourse strategies developed by McKeown (1985). The general approach would also suggest that experts might prefer structural descriptions to process descriptions (Paris, 1988) because they can already *infer* the process description from the structural, the former therefore conveying more implicit information.

By looking at possible plausible inferences when planning descriptions we attempt give a better solution to the problem of determining what to say given a particular communicative goal. The approach has potential for generating more memorable descriptions, where different types of statement are used to re-inforce some information, as well showing us how to economically convey a great deal of information, where some of this information may be implicit. It does not provide a solution to the problem of determining how to structure this communicative content (considered in much other research), though we may find that by considering further how people incrementally learn from descriptions we may obtain better structured text.

The prototype system has been fully implemented, but much further research is needed. The inference rules, user modelling and scoring functions need to be further developed, and other influences on text structure (such as focus and higher level rhetorical relations) incorporated into the overall model.

REFERENCES

- Cawsey, Alison (1989), *Generating Explanatory Discourse: A Plan-Based, Interactive Approach*, Unpublished PhD thesis, Department of Artificial Intelligence, University of Edinburgh.
- Collins, Allan & Michalski, Ryszard (1989) The logic of plausible reasoning: A core theory. *Cognitive Science*, 14:1-49.
- Cook, Malcolm, E., Lehnert, Wendy, G. and McDonald, David, D. (1984) Conveying Implicit Content in Narrative Summaries. In *Proceedings of COLING-84*, pages 5-7.
- Jameson, Anthony (1983), Impression monitoring in evaluation-oriented dialogue. In *Proceedings of the 8th International Conference on Artificial Intelligence*, pages 616-620.
- Joshi, Aravind, Webber, Bonnie and Weiscedel, Ralph, M. (1984) Living up to expectations: computing expert responses. In *Proceedings of the 7th National Conference on Artificial Intelligence*, pages 169-175.
- McKeown, Kathleen, R. (1985), *Text Generation : Using discourse strategies and focus constraints to generate natural language text*. Cambridge University Press.
- Mitchell, Tom, M. (1977), Version spaces: A candidate elimination approach to rule learning. In *Proceedings of 5th International Conference on Artificial Intelligence*, pages 305-310.
- Moore, Johanna, D. (1990), *A Reactive Approach to Explanation in Expert and Advice-Giving Systems*. PhD thesis, Information Sciences Institute, University of Southern California (published as ISI-SR-90-251).
- Paris, Cecile. (1988), Tailoring Object Descriptions to a User's Level of Expertise. In *Computational Linguistics (Special Issue on User Modelling)*, vol 14.
- Reiter, Ehud (1990), Generating descriptions that exploit a user's domain knowledge. In R. Dale, C. Mellish, and M. Zock, editors, *Current Research in Natural Language Generation*, Academic Press.
- Zuckerman, Ingrid (1990), A Predictive Approach for the Generation of Rhetorical Devices. In *Computational Intelligence*, vol 6, issue 1.