In "Cybernetics and Systems '94", R. Trappl editor, 1994, pages 1879–1886.

TOBI
An Ontological Based Interpreter for Temporal Presuppositions and Counterfactuals

STEFANO MIZZARO

Dipartimento di Matematica e Informatica, Università degli Studi di Udine, V. Zanon 6, Udine, I-33100, Italy
e-mail: mizzaro@udmi5400.cineca.it

ABSTRACT

This paper describes TOBI, a system that interacts with the user in natural language and can treat temporal relations and counterfactual utterances. This system presents three peculiar characteristics. First, it tackles the classical problem of temporal presuppositions from a novel and general point of view; second, the system draws a clear distinction between two types of knowledge namely ontology and content; third, the system is not based on a classical deductive system, but it uses the more primitive and flexible notion of model-based evaluation.

1. INTRODUCTION

In this paper I describe TOBI (acronym for "Temporal presuppositions and counterfactuals: an Ontological Based Interpreter"), a system that interacts with the user in natural language and treats temporal relations and counterfactual utterances.

TOBI presents three peculiar characteristics. First, it tackles the classical problem of temporal presuppositions (see Gazdar[4]) from a novel and general point of view; second, knowledge and inference in TOBI are divided into two distinct components, namely ontology and content; third, TOBI is not a classical deductive system that implements some nonmonotonic logic and uses a TMS, but is based on the more primitive and flexible notions of models and model-based evaluation.

The paper is structured in the following way. In Section 2, I present the linguistic phenomena handled by TOBI, namely temporal presuppositions and counterfactuals. In Section 3, the distinction between ontology and content is illustrated, with particular emphasis on the ontology and content of time. In Section 4 I describe *recursive models*, the data structures used in TOBI for modelling natural language utterances. Section 5 presents the system. Finally, Section 6 summarizes the work done so far and proposes some future extensions.

2. THE LINGUISTIC PHENOMENA

The linguistic phenomena handled by TOBI are situated on the boundary between semantics and pragmatics. This section informally describes such phenomena through a restricted corpus of examples.

To understand the meaning of an utterance, it is important to analyse the relations between utterances. Following Gazdar, an utterance *implies* another utterance if the latter is a consequence of the former. For example, utterance (1)

"Mary met John before she left" (1)

implies utterances (2) and (3):

"Mary met John" (2)
"Mary left". (3)

Here I give no formal definition of implication. A particular case of implication between utterances is *entailment*: utterance (1) entails (2). However, entailment is not the only type of implication, as utterance (3) proves: the relation between (1) and (3) is not an entailment, as showed by the fact that the following utterance is consistent

"Mary met John before she left and he persuaded her to stay at home". (4)

If we admit that (3) is entailed by (1), then (3) is also entailed by (4). But (4) entails

"Mary did not leave"

which contradicts (3). Using Gazdar's terminology, (3) is a *(temporal) presupposition* of (1). A presupposition is a form of implication weaker than entailment: the second part of (4) *cancels* the presupposition, so we do not have a contradictory utterance.

It is important to remark that although the event 'Mary left' has not happened, it is used in (4) in order to date the event 'Mary met John'. Moreover, from a logical point of view it seems more correct to say

"Mary met John before she *did not* leave and he persuaded her to stay at home"

instead of (3), but no human would do so. In other words, the problem is in *nonmonotonicity*: utterance (1) implies (3) only *by default* and the second part of (4) deletes the default. Therefore, a system handling such phenomena must be nonmonotonic. The most widely used instrument for this purpose is nonmonotonic logic; however, the system described in this paper is not based on this instrument, as will be shown later.

Temporal presuppositions, together with other kinds of presuppositions and implications, have been studied by Gazdar in the work referred to above. His treatment is not entirely satisfactory: there is no deep explanation of why 'before' introduces a presupposition, while, as utterance

"Mary met John after she left"

shows, 'after' introduces an entailment. In fact, Gazdar does not consider the *ontology of time* in his work, while, I think, the facts that time is ordered and the future unknown and partially unpredictable must be taken into account when dealing with utterances containing 'before' and 'after'. This point is investigated in the next section. Furthermore, Gazdar does not explicitly take into account relationships between events, which are necessary for example to explain the utterance

"Mary left before meeting John", (5)

in which the leaving event prevents the meeting.

A related phenomenon treated by TOBI is that of *counterfactuals*. In fact, (4) implies:

"If Mary had not met John, she would have left", (6) that is used for referring to a *non-real world* (the world in which Mary did not leave).

3. TIME: ONTOLOGY AND CONTENT

In TOBI, the treatment of the linguistic phenomena illustrated in the previous section is based on the dichotomy between *ontology* and *content*. Informally speaking, ontology is the component of knowledge that has a general logical status; on the contrary, content is the component of knowledge that is highly situation dependent.

As an example of this dichotomy, consider the case of (*subjective*) *time*. The ontology of time is its *ordering* and the fact that while the past is in a sense *closed*, the future is *open*. On the contrary, the *metric* of time is a content characteristic, in that the subjective evaluation of the duration of a time interval may vary depending on the situation.

It is possible to split an inference process into two parts: an ontological part and a content one. In the former, only *ontological inferences* (i.e., inferences based only on ontology, not on content) will take place; in the latter, only *content inferences* will occur. In the same way, it is possible to speak of *ontological knowledge* and *content knowledge*. This division is rather approximate: there seems to be a deeper link between ontology and content than what is illustrated here; anyway it is interesting to study how far it is possible to push this dichotomy.

I think that the phenomenon of temporal presuppositions can be explained in the following way: an event in the future cannot be certain, because of the partial unpredictability of the future.* This is why, using Gazdar's terminology, 'before' introduces a temporal presupposition, while 'after' does not.

In implementing TOBI, I have assumed that ontology can be handled using classical symbolic methods; there are reasons, however, to believe that this might not be true for content (see for example Airenti and Colombetti[1]). Therefore, in the present version of TOBI, content inferences are replaced by an interface to an external user. This interface is activated upon request of a master module, which implements ontological inferences.

4. RECURSIVE MODELS

TOBI is a simple natural language comprehension system, able to understand a text (sequence of utterances) and to provide correct answers to questions regarding the text. I have considered only *polar* questions, i.e. questions admitting as answers only 'yes' (`true`), 'no' (`false`) or 'I don't know' (`unknown`).

TOBI deals with text representation by building a model of the utterances, and it answers questions by evaluating them in the model. To understand how TOBI works, three notions must be examined: the notion of *recursive model* (*RM*), the operation of building an RM from a text, and the operation of evaluating an utterance in an RM. To build a model of a text, a function that *integrates* a previous model with the information of a new utterance is needed; this function will be named `int` (for 'integrate'). The function

* It is important to point out that 'future' refers to the *point of reference*, not to the *point of speech* (see Reichenbach[6]). In utterance (1), both the events ('met' and 'left') happened in the past, but the second is in the future of the point of reference.

that evaluates an utterance in a model will be named `eval`.

Roughly speaking, an RM is constituted by *instances* of classes of an encyclopedia and *relations* among those instances. Therefore, the system must contain an *encyclopedia*, that is a taxonomy of *categories* and *concepts*. The encyclopedia is the knowledge base of the system, and is needed to let the system know that Mary and John are persons, hence living beings, and so on; that the meeting of Mary and John is an event, etc.

Using the operation of *instantiation* it is possible to create a *token* for each individual mentioned in the utterance. Referring to utterance (1), there will be tokens for 'Mary', 'John', 'met' and 'left' (the last two being instances of the class `event`). Every token has an associated identifier; I shall use uppercase letters for instances of objects (`M` for 'Mary', `J` for 'John'), and lower case letters for events (`m` for 'met', `l` for 'left', etc.). As usual, tokens inherit *slots* from their parent concepts, so `M` is the value of the slot `agent` of `m` and `J` is the value of slot `theme` of `m`. Moreover, between tokens `m` and `l` there is a temporal *relation* to indicate that the meeting took place before the leaving.

Tokens, slots and relations, are not sufficient to obtain a complete model, since by using only these components, one would obtain the same RM for the utterance

"Mary did not meet John before she left"

and this is clearly a problem.

To deal with event occurrence and object existence, other elements are introduced in the RM: *spaces*, *attachments* and *signs*. The first notion is that of *space*. It is needed because not only an object exists, or an event takes place; it is more correct to say that an object exists (or an event takes place) *in a world*. Consider utterance (4): Mary did not leave in the *real world*, but it is correct to say that Mary left in the "counterfactual world" (see utterance (6)) in which she did not meet John. Analogously, it is possible to say that Donald Duck does not exist in the real world, but he exists in Walt Disney’s world.

A space is a formal tool for representing alternative worlds. I indicate the real world with `[]`. It is possible to represent the object existence and the event verification *attaching* every token to the right world: the relation between token and world is named *attachment*. Finally, attachments are labelled with a *sign* in order to deal with non-existence and non-verification, both of which cause a *negative* sign, whereas a *positive* sign has the obvious meaning.

As illustrated in Section 2, the occurrence of an event may be certain (the meeting of (1)) or uncertain (the leaving of (1)); this can be dealt with using the *certain* and *uncertain* signs. In the RM of (1), the signs labelling the attachments of the tokens for 'met' and 'left' are both positive, but only the first is certain, while the second is uncertain.

The RM obtained from (1) is illustrated in Figure 1. Only a little portion of the encyclopedia is presented (the portion needed to build the RM of the utterance); *is-a* and *instance-of* relations are represented by labelled grey arcs, tokens are shown as circled letters, slots are illustrated by means of oriented arcs, relations, as usual in entity-relationship diagrams used in data base theory, are represented by arcs labelled with a rhombus (the symbol < stands for 'precedes temporally'), a dashed arc represents an attachment, a bold sign is certain and a plain text sign is uncertain. For the sake of simplicity, in the graphic representation the names of the slots are not illustrated.

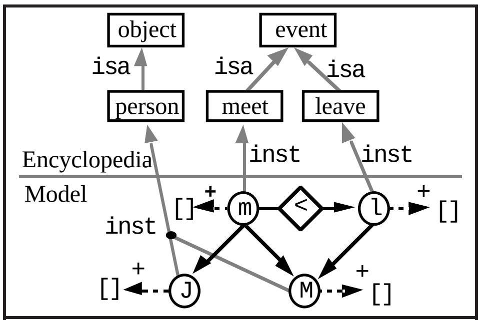

Figure 1. Graphic representation of the RM of (1).

The RM in Figure 1 is complete. Nevertheless, there is another element to add for dealing with the *causal links* relating the occurrence (or non-occurrence) of events. Examples can be found in utterances (4) and (6) (the occurrence of the meeting with John causes the occurrence of the event 'Mary stayed at home') and (5) (the occurrence of Mary's leaving causes the non-occurrence of the meeting with John).

The elements used in RMs to represent such causal relations are named *justifications*, and are represented by curved arcs. As signs, also justifications may be certain or uncertain. In order to understand the role of these new elements, consider Figure 2, in which the RM of (4) is represented. Here and in the following, for the sake of simplicity, I have omitted the representation of the encyclopedia (i.e. the classes and the `isa` and `inst` relations: the letters labelling the tokens should be sufficient for understanding which class each token is instance of. Furthermore, the token `p` is assumed to be an instance of the ad-hoc class `persuade to stay at home`.

The justification between the signs of `m` and `p` is uncertain (graphically represented by a thin curved line), whereas the one that links the signs of `p` and `l` is certain (thick curved line). The reason for this distinction is that the meeting implies persuading in a very weak sense (it is a precondition), while persuading (to stay at home) entails non-leaving.

Note furthermore that in Figure 2 `l`'s attachment is labelled with two signs: the positive one (uncertain) models the presupposition of the leaving and the negative one (certain) reflects the fact that the leaving actually did not take place. The last sign is the *preferred* sign (and it overcomes the uncertain one); graphically, this is represented putting it near to the end of the arc.

Justifications are needed not only by abstract completeness considerations, but also to deal with counterfactual utterances. This will be clear later, after the discussion about the evaluation of counterfactual utterances.

At this point, the structure of the RMs should be clear. Now, I present the way the functions `eval` and `int` work. Both such functions can be defined in the same way (by structural recursion on the *logical form* of an utterance, see below), therefore I describe only the way `int` builds the model of an utterance.

Roughly speaking, a raw RM is built on the ground of ontological considerations and it is then redefined using content knowledge. Consider for example the RM of utterance (5) represented graphically in Figure 3. The following steps take place during its creation:

- token `l`, from 'left', is created and it is attached to the space `[]` with a positive and certain sign. The sign is certain because of linguistic considerations: 'left' belongs to the main proposition;
- token `M` is created and it becomes the value of slot

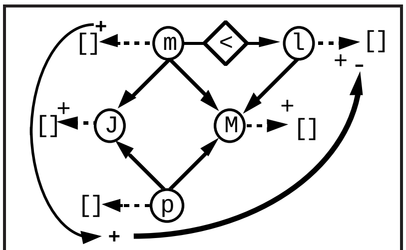

Figure 2. Graphic representation of the RM of (4). Pay attention to the role of justifications.

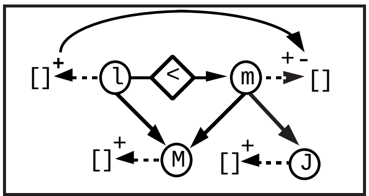

Figure 3. Graphic representation of the RM of (5).

agent of token l. Now, the building of the RM of the main proposition is terminated;

- token m, from the event of the secondary proposition, is created and attached to []. The sign of this attachment is still positive, but *uncertain* because the event is in a secondary proposition;
- the slot agent of token l assumes as value the token M, already present in the model; token J is instead created and it becomes the value of slot theme of token m;
- the temporal relation between the tokens l and m is created;
- all the above operations take place on the ground of ontological considerations. However, to complete the construction of the model, some content inferences are needed to create a negative certain sign (preferred to the positive uncertain one) on the attachment of m and the corresponding justification.

Thus, the division of ontological and content work is clear. Ontologically, tokens are created, slot values are filled, relations explicitly referred in the utterance are produced and attachments are created. On the ground of content considerations, justification arcs, representing the causal relations between events implicit in the utterance, are added and the same happens for new signs.

However, the separation between ontology and content is not so simple: the content part may create temporal relations and the ontology one may create justifications. This happens, for example, in the creation of the RM of (6). It is similar to the one represented in Figure 3, the only differences being the attachment of m (that is labelled by only one negative certain sign) and the justification (that is certain too). In this case, the temporal relation is created by the content module, in that the fact that the leaving takes place before the meeting is indubitably a content inference. Furthermore, the justification derives from ontological considerations, in that it appears explicitly in the word 'if' of the utterance.

As already specified, the discussion above regards exclusively the function int. Nevertheless, the function eval works in a similar way; instead of creating tokens, it verifies that the elements already exist in the model.

The function eval works in a particular way for the evaluation of counterfactual utterances. Such evaluation takes place in three steps: first, the antecedent and the consequent of the counterfactual utterance are evaluated in the current model; second, the current RM is modified accordingly to what said in the antecedent of the counterfactual, obtaining the *counterfactual model*; third, the consequent of the counterfactual is evaluated in the counterfactual model. Consider the evaluation of utterance (6) in the model obtained from (4) (the model in Figure 2). The evaluation takes place in the following way:

- the antecedent and the consequent of (6) are evaluated in the RM; both are false (and they must be false in order to evaluate the counterfactual utterance true);
- the counterfactual model is obtained modifying the original RM in such a way that the antecedent is evaluated false. The counterfactual model is illustrated in Figure 4. Observe that the token m is attached with a negative sign to [], in that the antecedent must be evaluated false. This, by means of the justification between the

signs of m and l (see the original model in Figure 2), leads to the removal of the positive sign on p's attachment and to label this token with an opposite (negative) one. The same happens for the token l; here the removal of the negative sign brings up the positive sign;

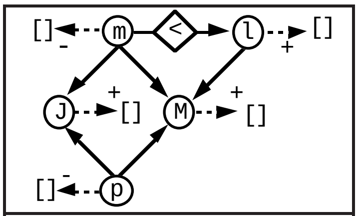

Figure 4. Counterfactual model for the evaluation of (6).

- the consequent of the counterfactual ("Mary left") is evaluated in the counterfactual model, obtaining true as result. The counterfactual utterance itself is then evaluated true.

5. TOBI'S IMPLEMENTATION

I now briefly present the system TOBI (for implementation details, see Mizzaro[5]).

TOBI dialogues with the user in natural language (English) and uses the RMs illustrated in the previous section as internal representations of utterances. The system is implemented in LPA Prolog 3.0 on a Macintosh II.

TOBI processes an utterance in three steps. The first step is *morphosyntactic analysis*: the input utterance is parsed into its syntactic structure. During this task, a lexicon and a DCG grammar (see Gazdar and Mellish[3]) are used as knowledge sources.

The syntactic structure is input to *semantic analysis*, that produces another representation of the initial utterance, namely its logical form. This representation is in a slot-filler notation, in which events and semantic roles are singled out. In this step, I use a semantic dictionary associating syntactic terms with the corresponding (semantic) concepts.

The last step is *interpretation*: here the logical form is used to build the RM of the utterance (or, more generally, to *integrate* the old model with the new information in the utterance). It is in this phase that TOBI's peculiarity comes to evidence. In most natural language systems, content knowledge is encapsulated in the encyclopedia, together ontological knowledge. In TOBI the two kinds of knowledge are separated; the encyclopedia contains only ontological knowledge, that can easily be dealt with in symbolic terms. The content module is presently an *oracle* that asks questions to the user.

To complete the description of TOBI, I present its modularization, sketched in Figure 5. Here is a list of TOBI's modules with a short description of their tasks:

- TOBI: module that loads and inits the system;
- SYNT: morphoSYNTactic analyzer that parses the utterance in input;
- SEM: SEMantic analyzer; it takes in input the syntactic structure produced by SYNT and produces as output the logical form of the utterance;
- ENC: ENCyclopedia; this is not a module, rather a knowledge base. It is a taxonomy of categories and concepts;
- FRAM: FRAme Manager; manager of encyclopedia and models. It implements the functions needed to work on classes (the encyclopedia) and instances (the models);

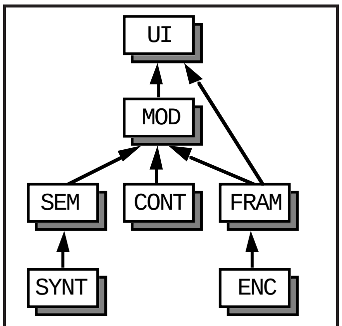

Figure 5. TOBI's architecture.

- `CONT`: this is the module devoted to exploit the CONTent inferences. It is not really implemented, because of its problematic computability; it is replaced by an oracle that asks the user the information needed;
- `MOD`: MODel builder; module that implements the functions `int` and `eval` using functions from SEM, FRAM and CONT;
- `UI`: User Interface; it accepts the utterance from the user (via keyboard) and answers his (her) questions. This interface is developed using the features of LPA Prolog for windows and menus management.

6. REMARKS AND FUTURE WORK

The main points discussed in this paper are: the distinction between ontology and content, the influence of the ontology of time on temporal presuppositions, and the implementation of the TOBI system based on the hypothesis that some aspects of the human mind (namely, ontology) can be simulated in a natural way by symbolic instruments, while other aspects (namely content) are better handled using other instruments.

TOBI is a prototype, and therefore it can be enhanced in various ways. To extend the set of cases it can deal with, an extension of the vocabulary is obviously needed. This, in conjunction with an improvement of the grammar, will allow for the treatment of utterances different from the ones presented here, but with some common characteristics. For example, counterfactual phenomena are very common in language, and do not need a specific syntactic construction: another common case is for instance the use of the verb 'to wish', as in "Mary really wishes she had left".

Another extension regards RMs. By now, they cannot handle utterances like "Mary left with George. This hurt John". In order to treat this kind of utterances, it will be necessary to introduce the concept of *situation* (see Barwise and Perry[2]) in RMs.

ACKNOWLEDGEMENTS

I am indebted to Marco Colombetti for valuable suggestions, and for long and stimulating discussions.

REFERENCES

1. G. Airenti and M. Colombetti, *Ontology of Mind, Subjective Ontology, and the Example of Temporal Presuppositions*, Politecnico di Milano, Dipartimento di Elettronica, Report 92-018 (1992).
2. J. Barwise and J. Perry, *Situations and attitudes*, MIT Press, Cambridge, MA (1983).
3. G. Gazdar, G., and C. S. Mellish, *Natural Language Processing in Prolog*, Addison Wesley, Wokingham, England (1989).
4. G. Gazdar, "Presuppositions and conversational implicature", in Cole P. ed., *Radical Pragmatics*, 188-198 (1979).
5. S. Mizzaro, *Elaborazione del linguaggio naturale: analisi e inferenze a livello pragmatico,* Degree Thesis, Università degli Studi di Udine, Udine, Italy (1992).
6. H. Reichenbach, *Elements of Symbolic Logic*, Macmillan, London (1947).