Editing Speech Acts: A Practical Approach to Human-Machine Dialogue

Paul Piwek, Roger Evans & Richard Power

ITRI, University of Brighton

To appear in: *Proceedings of Amstelogue'99, Amsterdam*

Abstract

This paper introduces a new approach to dialogue. This approach is based on the idea that dialogue can be seen as a form of two-person knowledge base editing. In a dialogue, it is the dialogue history which is being edited. We propose some constraints on editing dialogue histories which distinguish dialogue from knowledge base editing in general. We claim that our proposal yields both interesting conceptual insights into the nature of dialogue, and provides a starting point for building practically feasible dialogue systems. A first prototype of such a dialogue system has been developed. We provide a description of this system.

1 Introduction

In this paper we describe a new approach to dialogue. We demonstrate that this approach provides both a fresh view on the nature of dialogue and helps us to build practically feasible dialogue systems. Central to our approach is the notion of a shared dialogue history. The idea that utterances both give rise to updates of this history and are dependent on this history (e.g., for the interpretation of ambiguous expressions, pronouns, etc.) is well established (see, e.g., Clark, 1996).

Our main contribution is to take the idea that the utterances in a dialogue depend on and affect a shared history quite literally. A dialogue is modelled as a two agents editing a knowledge base (i.e., adding, cutting, copying and pasting knowledge objects) which represents the dialogue history. For this purpose, we use the WYSIWYM knowledge editing technology (see Power *et al.*, 1998). We show that a dialogue can be seen as a special type of two-person knowledge base editing. We provide a number of constraints on two-person knowledge base editing which distinguish a dialogue from other forms of knowledge editing, thus obtaining a new view on the nature of dialogue.

The approach that we describe has been developed in the context of the CLIME project.[1] The overall aim of this project is a generic web-based system for accessing legal information. We describe the application that is being developed within the CLIME project to illustrate the suitability of our dialogue model for web-based human-computer interaction.

We proceed as follows. In section 2, we describe WYSIWYM style knowledge base editing. Next, in section 3, we introduce our model of dialogue as knowledge base editing. In section 4, an application, which is an information system for maritime law, is introduced. We identify the sort of situations in which systems like the one we developed for the maritime domain could be useful. Section 5 contains a discussion of the relation between our approach to dialogue and the approach in current commercially available dialogue systems. Furthermore, at the end of this section we list avenues for further research.

2 WYSIWYM Knowledge Base Editing

Generally speaking, there are two ways of editing a knowledge base: by means of a command language and by means of direct manipulation. Let us first consider the use of command languages. A command language requires an interpreter which updates the knowledge base on the basis of the user's typed or spoken input. Ideally, the user should be able to edit the knowledge base by means of the language which s/he is most accustomed to, i.e., unrestricted natural language. Unfortunately, for the foreseeable future unrestricted text is not feasible for practical applications, because natural language understanding from free text is not yet sufficiently reliable. Therefore, so-called controlled languages have been introduced (see, e.g., Fuchs and Schwitter, 1996; van der Eijck, 1998). But there is a downside to the use of controlled languages: a user will have to learn to write in the controlled language, which can involve a substantial amount of effort.

Alternatively, there are direct manipulation interfaces to knowledge bases. Usually, the knowledge base consists of a network, which is graphically presented to the user (see, e.g., Paley, 1996). The user can now directly cut objects from the network and insert objects into the network. Thus each action by the user has a direct semantic interpretation. Hence, no parsing and interpreting is required. Again, there is also a downside. Network representations of knowledge can become very complicated and therefore difficult to understand by the user.

In Power *et al.* (1998), a new solution, called WYSIWYM (What You See is What You Meant), to the problem of knowledge editing has been proposed. WYSIWYM aims at the best of both worlds: the practical feasibility of direct manipulation and the ease of use of natural language text. The basic idea underlying WYSIWYM can be presented by means of a simple diagram, see Figure 1.

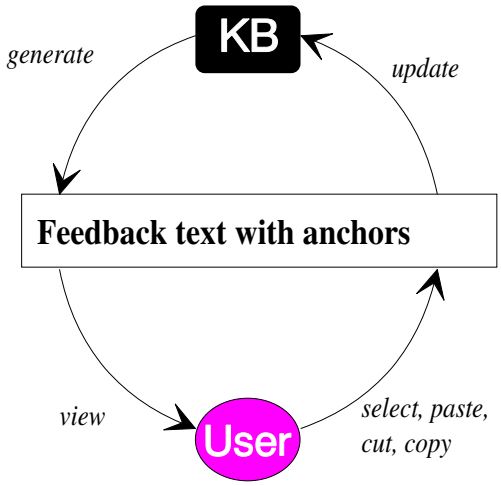

Figure 1: The editing cycle

Figure 1. represents the editing cycle. Given a Knowledge Base (KB), the system generates a description of the knowledge base in the form of a 'feedback text' containing 'anchors' representing places where the knowledge base can be extended. Each anchor is associated with pop-up menus, which present the possible extensions of the KB at that point. On the basis of the extension that the user selects, the knowledge base is updated and a new feedback text is generated from the new contents of the KB. Additionally, spans of feedback text representing an object in the KB can also be selected by means of the mouse. Cut and copy operations are available which allow the user to cut or copy the underlying knowledge base object into a buffer. Subsequently, such an object can be pasted into a location where the KB is still incomplete. After a cut or paste action, a new feedback text is generated which represents the updated KB.

Let us consider a very simple example which allows us to briefly illustrate the essential features of WYSIWYM editing. As usual, the knowledge base consists of two parts: a T(erminological)-box and an A(ssertion)-box. In the T-box, we specify the set of available concepts and their attributes:[2]

(1) a. `domain > [like, person].`
    b. `person > [man, woman, girl, boy].`
    c. `like intro [attr1:person, attr2:person].`

We start with the concept `domain`, which has two subconcepts: `like` and `person` (1.a). Subsequently, we introduce the subconcepts of the concept `person` in (1.b). Finally, in (1.c), it is stated that the concept `like` has two attributes. A concept following a colon indicates which objects are legitimate values of the attribute.

The A-box contains the actual knowledge that is to be edited. Formally, this knowledge has the structure of a Directed Acyclic Graph (DAG). The nodes in the graph stand for the instances of concepts, i.e., objects, and the directed arcs of the graph represent attributes. The

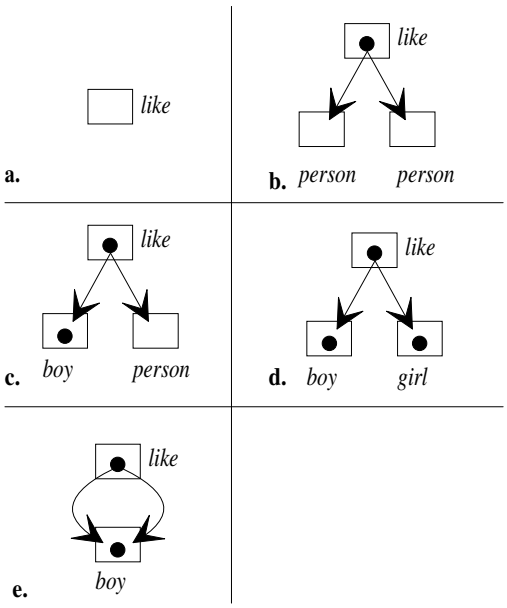

Figure 2: Editing a DAG

basic editing operation on a DAG is that of adding a new object, of a specified type, as the value of an attribute of an existing object. However, when we start knowledge editing, the A-box may be empty. To start the editing process, we need a predefined root attribute in the A-box, which is not part of the T-box proper, but whose value has to be an instance of one of the concepts that have been introduced in the T-box. For instance, let us assume that value of the root attribute has to be an instance of the concept `like`. This means that the state of the KB before we start editing can be represented with the graph in Figure (2.a). On the basis of this KB a feedback text is generated:

(2) **Some relation of liking.**

The entire span of text is in boldface. This indicates that the text is an anchor. By clicking on it, a menu appears which shows the possible alternatives for expanding the KB. In this case, the menu has only one option: `new object`. When the user selects this option a new object of the concept `like` is introduced into the KB (see fig. 2.b). On the basis of the new KB a fresh feedback text is generated:

(3) **Some person** likes **some person**.

The new feedback text has two anchors. When the user selects the first anchor, a menu with the option `new object` appears again. But now, when user user clicks on `new object`, another menu appears, which allows the user to select the concept to which the object belongs:

man
woman
boy
girl

Let us assume that the user selects `boy`. The new KB (fig. 2.c) gives rise to the following feedback text:

(4) A boy likes **some person**.

Along the same lines, the second anchor can be expanded. This leads to a fully completed knowledge base (fig. 2.d) and the following feedback text:

(5) A boy likes a girl.

Alternatively, instead of inserting a new object into the incomplete network (fig. 2.c), the user could also have chosen to copy an object that was already available, and paste it into the part of the network that was not yet developed. For instance, the user could have selected the span of text 'a boy'. A menu would have appeared with the options `cut` and `copy`. If the user had selected `copy`, the underlying object would have been stored in a buffer. Subsequently, the user could have selected the span of text which represents the incomplete part of the KB, i.e., '**some person**'. A menu would have appeared with the option `paste`. By performing the paste action, the network in fig. (2.e). would have been constructed, and the following feedback text would have been generated:

(6) A boy likes himself.

In this case, a reflexive pronoun is generated for the second attribute which takes the boy in question as a value. Note that if `copy` and `paste` had simply operated on the graphemic level of the sentence (i.e., the character sequence) instead of the underlying semantics, the result would have been 'A boy likes a boy'. For a more detailed account of coreference and WYSIWYM editing we refer to van Deemter and Power (1998).

3 Dialogue and Knowledge Base Editing

In a dialogue, two persons produce utterances and thus build up a shared history. This history is grounded in the physical events that constitute the utterances. On a higher level of analysis, the history consists of the words that were used and the syntactic structure of the phrases in which the words occur. Let us call this the linguistic level. Finally, on the highest level of analysis, the history consists of a sequence of speech acts. The speech acts each have a speech act type or force and a semantic content (Searle, 1969).

In ordinary dialogue the interlocutors edit the dialogue history by producing new utterances which extend the dialogue history. They also edit the dialogue history at its higher levels of analysis: by producing new utterances they also produce new speech acts. In this paper, we want to examine what it would be like if the interlocutors could directly edit the dialogue history at its highest level of interpretation: the speech act level.

WYSIWYM allows us carry out this experiment. For that purpose, we assume that the KB represent the dialogue history at the speech act level. Unrestricted two-person editing of this KB would, however, be nothing like ordinary dialogue. We need three simple constraints on the editing process to arrive at something which is similar to ordinary dialogue. We will state these constraints and then discuss their implementation in a WYSIWYM environment.

Turn taking In a dialogue, people normally speak one at a time. In order to mimic this phenomenon, we need to prevent the interlocutors from simultaneously editing the dialogue history. Let us state this in terms of two simple constraints:

(C1) Interlocutors are prevented from editing the dialogue history simultaneously.

(C2) After an interlocutor has finished editing a speech act, s/he passes the turn to the other interlocutor.

For the moment, these constraint will do, although it is a simplification: in real dialogues overlapping speech does occur.

Immutability of history Once something has been said, the fact that it has been said is immutable.[3] In terms of knowledge editing this means that only the most recent speech act can be changed in the editing process.

(C3) Interlocutors are prevented from editing speech acts other than the most recent one.

Without this constraint in place we run into some serious difficulties. Suppose that it were allowed to edit speech acts further back in history. Note that some of the speech acts that occur after time some time *T* will depend on the speech acts before time *T*. This means that if a speech act before time *T* is edited, this also has repercussions for the ensuing speech acts. These repercussions can, however, not be overseen by one participant alone, since the ensuing speech acts were produced by both participants. For instance, if participant *A* were to change a question that preceded an answer produced by *B*, then this change would also affect the answer. Hence, the editing operation on the question would have no clearly delineated effect with respect to the entire dialogue history.

Note that although it is not possible to edit a dialogue history somewhere in the middle, this does not prevent us from going back to some point *T* in the history, and then pursuing a alternative history to the one that actually took place after *T*. This simply involves making a copy of the history from its starting point until *T*, and editing that copy. In section 4,we shall see that this can be useful feature for practical dialogue systems.

As a consequence of the constraints (C1), (C2) and (C3), interlocutors are prevented from editing one and the same speech act (whether it be simultaneously or at different times). Note again that this is a simplification: sometimes one interlocutor does complete another interlocutor's utterance.

Accessibility of history Although in ordinary dialogue, we cannot change previous utterances, we can make use of the content that they introduced, by means of, for instance, anaphora:

A: The pump is not working.
B: You have to replace *it*.

and presuppositions:

A: The pump is out of order.
B: Why *is it out of order*?[4]

In terms of knowledge editing this means that we are allowed to copy items from the preceding history and paste them into the speech act that is currently under construction. In this respect, knowledge editing seems to require no further constraint. Note, however, that as soon as the semantic representations associated with speech acts become sufficiently expressive further constraints are needed. We have in mind the constraints on accessibility of discourse referents (which correspond to objects in the knowledge base) as stated in Discourse Representation Theory (Kamp & Reyle, 1993). It is beyond the scope of this paper to address this issue, but see Kibble *et al.* (1999) and Power (1999).

Implementing Dialogue as Knowledge Editing Let us now describe how we have implemented these ideas by means of a WYSIWYM system. The idea is that the interlocutors are a user and a software agent (e.g., an expert system). The user and the software agent talk with each other via a *dialogue manager* which maintains the dialogue history. The dialogue history is represented as a DAG. In this DAG, there will be nodes representing speech acts. At the beginning of a dialogue we start with a DAG whose speech acts are not yet specified. In other words, the dialogue history still has to be filled in by the interlocutors. The interlocutors can, one at a time, expand a speech act, thus establishing step by step a dialogue history.

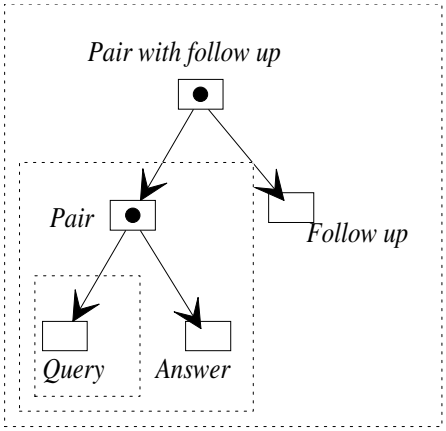

Figure 3: Incomplete DAG representing a dialogue history

Consider the DAG in figure 3. This DAG represents part of a dialogue history which still has to be filled in. There are locations representing three speech acts: a query, an answer and a follow up query to the answer (in speech act terms, we have two directives and one representative). The dotted boxes indicate in what order the interlocutors are allowed to edit the incomplete dialogue history. We start at the innermost box. Technically, this means that the root attribute is set to the location in the innermost box.

Let us assume that the user takes the first turn. The effect of having the `query` location as the value of the root is that the WYSIWYM generator will only generate a feedback text for this node. The feedback text is:

User: **Some enquiry**.

The user can now start editing this query, until s/he is satisfied with the result, e.g.:

User: I have a tanker with an opening in the upper deck. The opening is elliptical in shape (400x600mm). Is this acceptable?

Now that the user has finished editing the query, s/he can hand over the turn to the other interlocutor. This means that the user interface will have a button for this purpose. The software agent then constructs a representation of its answer (of course, it doesn't use the WYSIWYM interface, since the agent has been build in such a way that it can only construct well-formed representations). The software agent submits this answer to the dialogue manager, who slots it into the appropriate location of the dialogue history (the `answer` location).

Now, it is the user's turn again. At this point the root attribute is set to the `pair with follow up` location. Note that here the root is not set to an incomplete location, but rather to an object. The effect of this is that the generator will only generated the feedback text for the knowledge objects below this object. Thus a feedback text is generated which represents the full query and its answer, and an editable span of text for the follow-up query:

User:	I have a tanker with an opening in the upper deck. The opening is elliptical in shape (400x600mm). Is this acceptable?
System:	Yes it is acceptable without reinforcements.
User:	**Some follow up question**.

At this point, the user can edit the `follow up` location and copy objects which are situated under the `pair` node into the follow up location. In other words, although the content of the preceding speech acts can no longer be edited, it is available for constructing the new speech act. For instance, the user might expand the follow-up location into a why question:

User:	I have a tanker with an opening in the upper deck. The opening is elliptical in shape (400x600mm). Is this acceptable?
System:	Yes it is acceptable without reinforcements.
User:	Why is **something** the case?

At this point, the user can copy the proposition 'it is acceptable without reinforcement' into the location '**something**'. Thus, we end up with the following feedback text:

User:	I have a tanker with an opening in the upper deck. The opening is elliptical in shape (400x600mm). Is this acceptable?
System:	Yes it is acceptable without reinforcements.
User:	Why is it acceptable without reinforcements?

For the purpose of this example, we have identified the incomplete location of the dialogue history with very specific speech acts (query, answer, follow-up). We could also have chosen to use much broader types of speech acts as locations. For instance, we could have simply identified a location with a speech act of any type, and then let the user, by means of the WYSIWYM menu, choose whether it should be a question, conclusion, promise, declaration, etc. For a practical application, where the sequence of speech act types can be predicted reliably in advance, such freedom would, however, only have a negative effect on the ease of use of the system. Furthermore, if we were to use more a more general taxonomy, we would be faced with the problem of selecting an appropriate one. Traum (1999) points out that one needs to consider the task at hand when selecting a speech act taxonomy. For our purposes, a taxonomy which closely follows the intuitions of ordinary language users is most desirable.

There is one limitation to the system as we have described it here. The range of possible speech act *types* that can be constructed at a given point in the dialogue does not depend on the preceding dialogue history (note, however, that the *content* of a speech act can depend on what has been said before, since the user is allowed to copy material from the preceding dialogue). This range is determined by the networks that the T-Box allows us to build. This T-box is static during the dialogue. We might, however, want to allow the user to only ask a follow-up question of a particular type, after an answer of a particular type. This can be achieved by giving the dialogue manager a more active role. In addition to simply moving the root attribute to the next incomplete speech act location, the dialogue manager might pre-edit this speech act on the basis of the information it has concerning the preceding dialogue, or simply hide certain, no longer relevant, ways to expand the location from the user.

For instance, suppose the system has just replied 'No' to the question 'Are there any pumps on this ship?'. In that case, a where-question ('Where are the pumps located') is no longer a sensible follow-up. Alternatively, if the answer had been 'Yes', a where-question would have been a natural follow up. Here it is useful if the system can dynamically decide whether to allow the user to construct a where-question or not. This idea has, however, not yet been implemented in our first prototype.

4 An Application

The application that is described in this section is being developed in the context of the CLIME project whose aim is to develop software to support access to legal and regulatory information. The concrete application that is under development is in the area of maritime law.

Let us first describe the context of use of the application. The potential end users of the application are the surveyors of a classification society. The business of this society is the production and application of rules, in particular, in the maritime domain. In order for a ship to be insured, the owner of the ship has to make sure that the ship has been certified by at least one such classification society. Certification is carried out by the surveyors, who check whether a ship fulfils all the rules that are laid down by the classification society.

Currently, the rules of the classification society are available both in paper form and on CD-Rom (with simple information retrieval functionality). Our project partner at the university of Amsterdam has developed a legal reasoning system which can apply (a formalisation of) the rules to descriptions of ships (see Winkels *et al.*, 1998). The system can, given a (partial) description of a ship, determine which rules apply to it, and whether the described situation is acceptable according to the rules or not.

Requirements On the basis of discussions with the end-users of the system (i.e., experts of the French classification society Bureau Veritas) we established a number of *requirements* to which we have geared the architecture of our application prototype:[5]

- The system aims to help the surveyors in their daily practice of certifying ships. This means that the system should be available on site, for instance, during an inspection at a ship yard. Therefore, a web-based architecture has been chosen for the system. The user down loads the user interface as a JAVA applet into his/her computer, whereas the system itself is running on a possibly more powerful server elsewhere (legal information serving is computationally expensive for real life domains).

- The system is used in an internationally operating company and should be adaptable to the language of the local users.
- The user can ask queries and follow-ups to his/her queries, even days later than the original query.
- Given the nature of the domain, i.e., legal/regulatory information, there is a preference for having access to information in its written form.

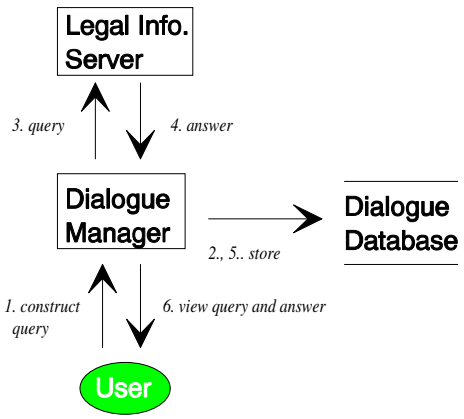

Figure 4: The system architecture

The Architecture A simplified representation of the system architecture is depicted in Figure 4. The legal information server and the dialogue database are located on a server. The user interface to the dialogue manager is running on the user's computer as a client, whereas the main functionality provided by the dialogue manager is running on the server. The idea is that 1. the user constructs a query using WYSIWYM. 2. This query is stored in the dialogue database. 3. The query is submitted to the legal information server. 4. The legal information server returns back an answer, which is integrated into the dialogue history (up till that point it only contained the user's query) which the dialogue manager retrieves from the dialogue database. 5. Subsequently, the dialogue manager stores the updated dialogue history in the dialogue database. 6. The user can view (the feedback text corresponding to) the new dialogue history (the user is notified by the system of any changes to the dialogue database), and if needed continue to edit it.

Database-oriented dialogue There are two reasons for maintaining a dialogue database. Firstly, it allows the user to quickly construct queries which deviate not too much from queries which s/he submitted earlier on. For that purpose, the user simply retrieves a query from the database which is sufficiently similar to the one s/he wants to submits and modifies this query before submitting it. The second reason for maintaining a dialogue database is that this makes the dialogue independent of a persistent internet connection between the user's computer and the server. It is now possible to construct a follow up query to a query which was posted hours, days or even weeks ago. Furthermore, the legal information server may sometimes take more than an hour to compute an answer. It would be rather awkward if the user needed to remain logged in to wait for such an answer. Fortunately, this is not necessary because his or her query is stored in the dialogue database and updated with the answer whenever the legal information

server has found one. The next time the user logs into the system s/he is notified of the change in the dialogue database (alternatively, the user may also be notified immediately by email) and can inspect the query and its answer.

Example of a session with the system After a user has logged in to the system, a main window appears. This window (see figure 5.) is divided into two panels: a panel with a directory structure and a view panel. The directory structure consists of folders and files in which dialogue histories are stored. By clicking on a file, the user can view the dialogue history which it contains in the view panel. Thus it is possible for the user to browse through the collection of dialogue histories which s/he has built up by querying the system.

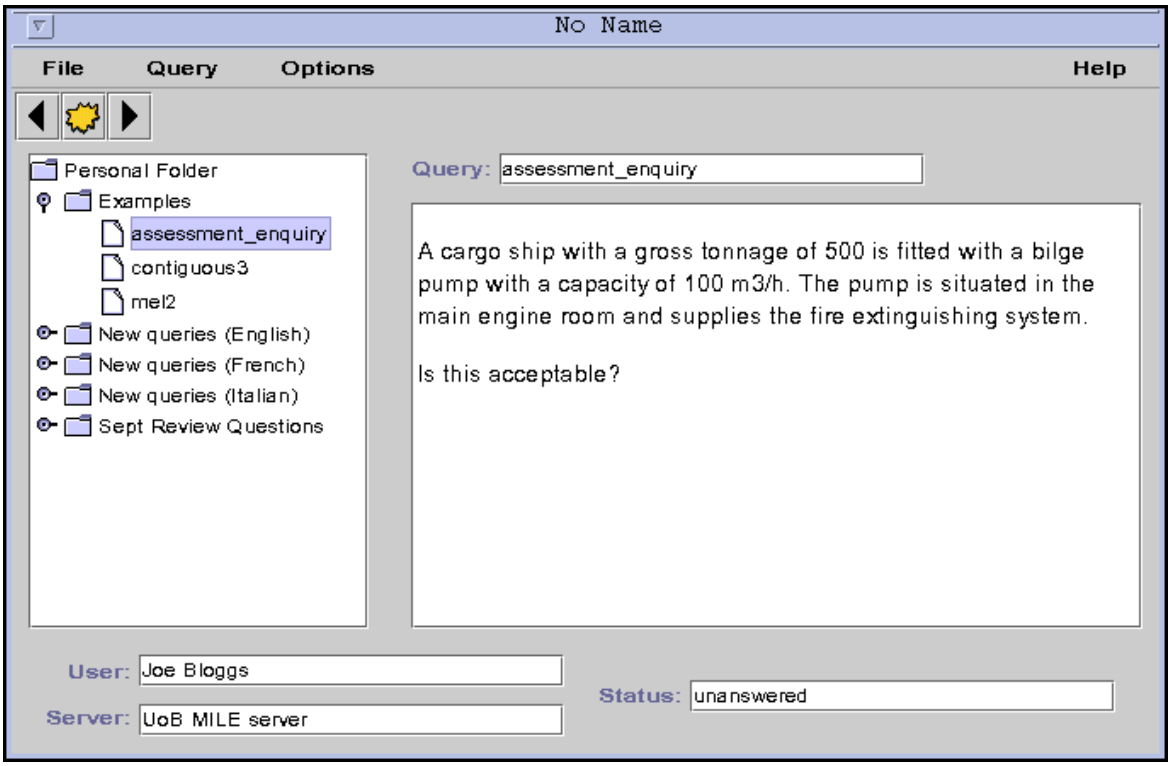

Figure 5: The main window

On the menu bar, there is a file, query and options menu. On the file menu, the user finds the usual operations for manipulating files and folders, such as 'delete', 'rename, open, close folder, etc. The query menu provides amongst other things the possibility to create new dialogue histories, edit existing dialogue histories and extend existing dialogue histories (i.e., construct follow-ups). For instance, when the user selects the option *edit*, a WYSIWYM editing window pops up. Within this window the user can then make alterations to the dialogue history which was depicted in the view panel (see figure 6.). When the user is satisfied with the changes which s/he made and has constructed a complete query (i.e., the last query in the dialogue history contains no locations which still need to be filled in), s/he can send the query to the legal information server by selecting submit on the file menu of the editing window.

Let us now have a closer look at the query in figure 6. The current system has a T-Box and a grammar which covers one chapter of the rules book of the classification society (there are 30 chapters in total). The example is modelled after the following query which was provided by experts from the classification society: 'I have a cargo ship with a gross tonnage of 500, that is fitted with a bilge pump with a capacity of 100m3/h in the main engine room. This

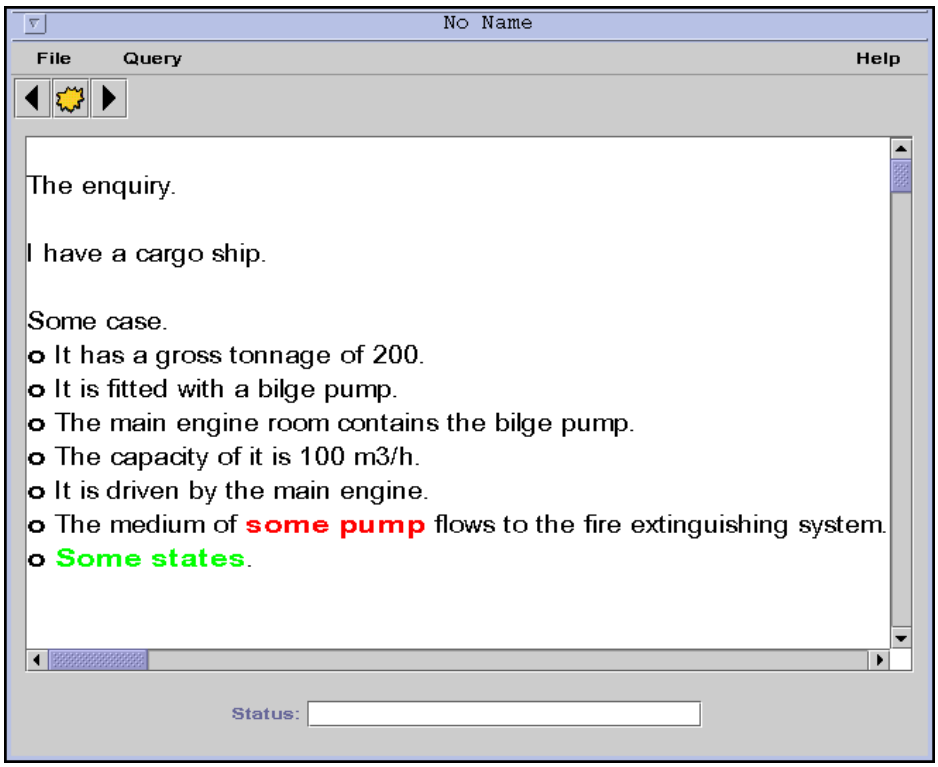

Figure 6: The dialogue history editing window (English)

pump is driven by the main engine and can provide water to the fire extinguishing system. Is this acceptable?'. Currently, a user can construct this query with the WYSIWYM system. The feedback text that is generated by the system is depicted in figure 6. (The feedback text is not yet fully complete. Note that the phrase 'some pump' is printed in a lighter grey. These are anchors which indicate in which respects the underlying content of the query is still incomplete).

On the basis of a formal representation of the answer to this query (the formal representation was handmade, because the legal information server and the dialogue manager have not yet been integrated) the text in Figure 7. is generated.

It is beyond the scope of this paper to go into the formal details concerning the representations of answers (but see Winkels *et al.*, 1998). Let us, however, briefly sketch how the surface structure of the answer is related to its underlying formal representation. The underlying representation makes reference to two rules which are applicable to the case that the user wants to examine. The first rule says that the firepumps of a ship are not allowed to be driven by the main engine. A second rule says that it is allowed that a firepump is driven by the main engine, if the gross tonnage of the ship is less than a 1000. Thus we have (where C = it is allowed):

$P \rightarrow$ not C
$(P \wedge Q) \rightarrow C$

The rule with the more restricted antecedent wins. We have used the following template to associate this type of answer with a natural language text (where $l(P)$ stands for the translation into natural language of the proposition P):

Although $l(P)$, $l(C)$, because $l(Q)$.

When the legal information system comes back with the answer, it is fitted into the representation of the dialogue history. The system then generates a text from the representation of the antecedent dialogue history which has been merged with the answer. It is essential that the antecedent history and the answer are merged in order to get cross speaker anaphora right. For instance, in figure 7., 'the bilgepump' is used to refer to an object which was introduced in the user's query. In the underlying representation, the bilgepump which is referred to in both the query and the answer is represented with one object (i.e., a discourse referent).

The answer.

Although the bilge pump is a firepump and it is driven by the main engine
(see 20–052–11), the case is allowed, because the cargo ship
has a gross tonnage of less than 1000 (see 20–052–31).

A follow up.

Can you tell me why the bilge pump is a firepump?

Figure 7: Answer and a follow-up

In Figure 7., we also see an example of a follow-up query. This follow up query has been constructed starting from the following feedback text: 'Can you tell me why **some answer state**?'. The user can fill the location represented by the span '**some answer state**', by selecting a part of the answer, copying the underlying knowledge object, and pasting it into the follow up. Thus a complete follow up question is constructed.

Finally, let us note that because the feedback texts which represent the dialogue history are generated from an underlying language independent representation, it is rather straightforward to change the language which is supported by the system. All that needs to be done, is to change the natural language generator which produces feedback texts from the formal representation of the dialogue history. The user can select a language on the options menu of the main window. For instance, if the user at this point selects French, the text in the query window is replaced with its French equivalent (see figure 8).[6]

5 Discussion and Conclusions

In summary, the model of dialogue as knowledge editing that we have introduced provides a fresh view on dialogue. In particular, by exploring the differences between dialogue and knowledge editing in general, we obtain a new perspective on the nature of dialogues. Furthermore, the model is promising for use in practical applications. It is based on the user friendly interaction provided by WYSIWYM editing. In combination with the use of a dialogue database, it has some features which make it particularly useful for human-machine dialogue over the internet. In the remainder of this section we compare our approach with alternative approaches to natural language interpretation in dialogue, and conclude with some remarks on directions for further research.

Natural language interpretation is a notoriously difficult task. Textbooks on human-computer interaction (e.g., Sutcliffe, 1995; Dix et al., 1998) often point out that it seems unlikely that a general natural language interface will be available for some time, if at all. This claim is based on the fact that natural language interpretation is heavily context-dependent. In particular, interpretation may require any kind of general background knowledge that is available

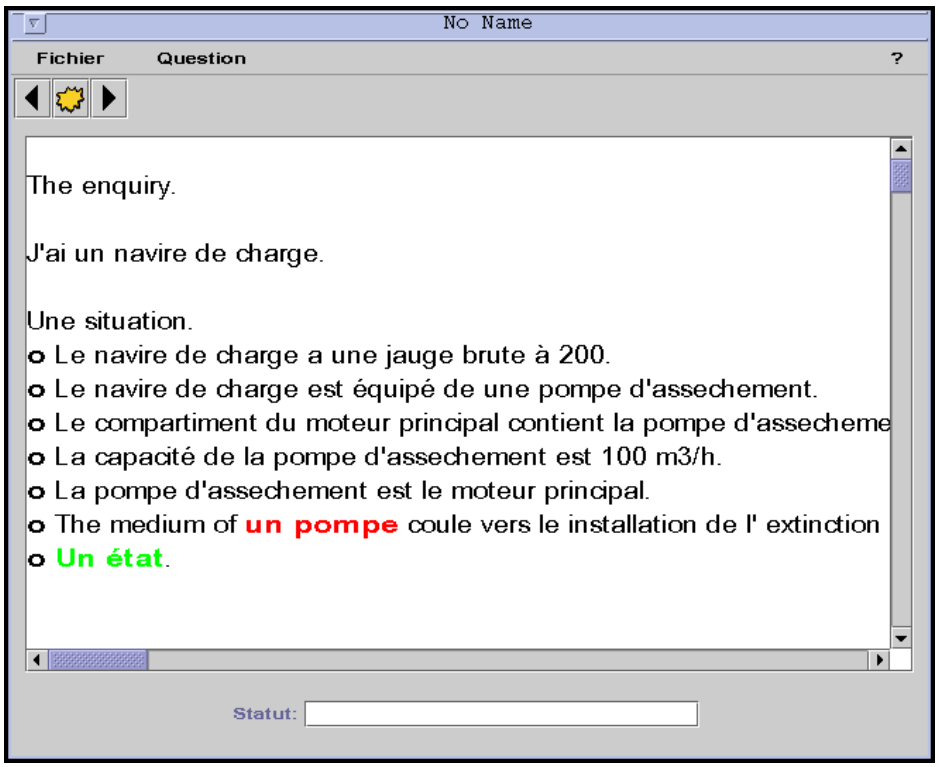

Figure 8: The dialogue history editing window (French)

to a human agent. Formalization and tractable implementation of such immense knowledge resources seems implausible for the foreseeable future (currently, the only large scale project aiming at that goal is the Cyc project, see Guha & Lenat, 1994).

However, natural language interpretation seems feasible for small scale domains. In recent times, commercially viable dialogue systems have been developed for such constrained domains. The Philips automatic train timetable information system (Aust *et al.*, 1995) is a case in point. However, the success of this and similar systems depends on the fact that there is a very limited number of simple informational items which the systems needs to elicit from the user in order to satisfy the user's request. For instance, in order to provide the user with a train connection, the Philips system needs the place of departure, the destination and the time at which the user wants to depart or arrive. The system can provide a connection when it has filled the slots for these parameters.

Suppose that the user tells the system "I'd like to go to Hamburg". In that case, the only part of the utterance that the system actually understands is "to Hamburg". It uses this information to fill in the slot for the destination. Thus, the system could not have told the difference between the aforementioned utterance and "I'd not like to go to Hamburg". Furthermore, the Philips system does not deal with anaphoric references such as pronouns and descriptions. These limitations make this type of approach unsuitable for the CLIME domain.

In CLIME, the information that the system needs to elicit from the user typically consists of multiple informational items which may be linked by co-reference and enter into logical (e.g., negation) and discourse relations (e.g., the relation between a legal case and a question about that particular case). Consider, the following query in the maritime domain which illustrates all three of the aforemention relations:

> I have an oil tanker. One of its bilge pumps is not functioning. Can I replace the pump with an ejector system?

The technology that is used in current state of the art applied dialogue systems is not suited for this type input. A legal analysis system will have to understand the meaning of the user's utterance on a much more detailed level than the aforementioned dialogue systems do in order to provide an answer to the user's query. In particular, the system has to recognize coreferential terms and the logical and rhetorical structure of the query. We already pointed out that for the foreseeable future unrestricted text seems to be not feasible for practical applications.

We hope to have made it plausible that our approach to dialogue as knowledge editing, and in particular, the use of the WYSIWYM technology, provides a promising alternative to the approaches taken in currently available commercial dialogue systems and research dialogue systems (the latter often rely on full parsing and interpretation of a user's utterances, but see Milward, 1999). Note that there is an interesting connection between WYSIWYM editing and the frame-driven or form filling approaches which are characteristic of commercially available dialogue systems (The frame-driven approach goes back to Bobrow *et al*. (1971). Its recent success seems to be partly due to the advent of high quality speech recognition technology). When a user constructs a DAG by means of WYSIWYM, the user can be seen as filling in slots. However, in doing so, the user builds up a DAG whose complexity goes beyond that of a simple frame structure. Firstly, the DAG may be build up using recursion (in other words, when we construct a DAG, we may fill in an indefinite number of slots, whereas frames will always have a fixed number of slots), and secondly we can copy parts of the dag and insert them into other locations in the DAG which have not yet been filled in (we make essential use of this feature to model coreference).[7]

Whereas building up a structure using recursion and copy and paste can be done by means of direct manipulation, these features seem to make it difficult to see how DAG's could be successfully used in spoken natural language dialogue systems, without having to fall back to full parsing and interpretation of the user's utterances.

Finally, let us indicate some issues which we would like to address in the future:

- Currently, a demonstrator of the dialogue manager exists. It has not yet been integrated with the full application. However, the first trials with end-users of the system have been very encouraging. Three end users were asked to formulate queries using the WYSIWYM technology and had no difficulties in doing so. Within the coming months, a formal evaluation of the system has been planned.

- We would like to further explore the possibilities of making the range of utterances that the user can construct at a given point in the dialogue more dependent on the antecedent dialogue history.

- The issue of sub dialogues has not been addressed in this paper. In our application domain these can, however, not be avoided (e.g., sometimes the legal information system needs to obtain more background information from the user before it can answer a question).

- Finally, we have been working on extending the expressive power of the formal language which is used to represent dialogue histories. Power (1999) shows how we can extend it to include Discourse Representation Structures (Kamp & Reyle, 1999), and recently a prototype which can deal with reference to plural objects has been devised.

Acknowledgements

We would like to thank Kees van Deemter for commenting on an earlier version of this paper. We would also like to thank the audiences at Amstelogue '99, the ITRI CL Meeting (University of Brighton) and the IMS Colloquium (University of Stuttgart) for useful comments and discussions. The front-end JAVA interface which is depicted in section 4 was implemented by Neil Tipper with help from Chris Douce.

PAUL PIWEK, ROGER EVANS AND RICHARD POWER
Information Technology Research Institute (ITRI)
University of Brighton, Lewes Road,
Brighton BN2 4GJ
UK
e-mail: firstname.lastname@itri.brighton.ac.uk

Notes

1. CLIME stands for Computerized Legal Information Management and Explanation. It is funded by the EC Esprit Programme under project number EP 25.414. The participants in CLIME are British Maritime Technology Ltd., Bureau Veritas, TXT Ingegneria, The University of Amsterdam and the University of Brighton.
2. We use the ProFit notation. See Erbach (1995).
3. Note that this does not mean that a participant can not contradict something that s/he said earlier on, but merely that the fact of having said it can not be 'undone'.
4. The question 'Why P?' is said to presuppose that P.
5. This list contains only a selection of the requirements. The full list of requirements can be found in Piwek *et al.* (1999:6).
6. Currently, our lexicon for French is not yet complete. Whereever the generator cannot find the appropriate French word, it defaults to English.
7. It would be fairly straightforward to devise, for instance, a WYSIWYM system for the train information domain. In the T-box we would need a concept `journey` with attributes whose values are a place of departure, a destination and departure time. On the basis of an object belonging to the concept journey, the system could then generate a feedback text saying 'I want to travel from **some place** to **some other place** at **some time**'. The user would subsequently have to expand the anchors which are in boldface. The final text might then be: 'I want to travel from Brighton to Gatwick on the second of August at 10:00 am.'

References

Aust, H., M. Oerder, F. Seide and V. Steinbiss (1995), 'The Philips automatics train timetable information system', Speech Communication, 17, 249–262.

Bobrow, D., R. Kaplan, M. Kay, D. Norman, H. Thompson and T. Winograd (1971), 'GUS, A Frame-Driven Dialog System', Artificial Intelligence, 8:2, 155–173.

Clark, H. (1996), *Using Language*, Cambridge University Press, Cambridge.

Dix, A., J. Finlay, G. Abowd, R. Beale (1998), *Human-Computer Interaction*, Prentice-Hall, Hertfordshire.

Erbach, G. (1995) 'Prolog with features, inheritance and templates', *Proceedings of EACL-95*, Dublin. 1995, 180–197.

Fuchs, N. and R. Schwitter (1996), 'Attempto Controlled Language (ACE)', *First International Workshop on Controlled Language Applications*, Leuven, Belgium.

Guha, R. and D. Lenat (1994), 'Cyc: A Mid-Term Report', AI Magazine, 11(3), 32-59.

Kibble, R., R. Power and K. van Deemter (1999), 'Editing logically complex discourse meanings', *Proceedings of the 3rd International Workshop on Computational Semantics*, 147–162.

Milward, D. (1999), 'Towards a Robust Semantics for Dialogue using Flat Structures', *Preproceedings of Amstelogue '99*.

Paley, S. (1996) 'Generic Knowledge-Base Editor User Manual', Technical Report, SRI International, California.

Piwek, P., N. Tipper & R. Evans (1999), 'Specification of the Natural Language Software Modules', CLIME Deliverable D.1.4, University of Brighton.

Power, R., D. Scott and R. Evans (1998), 'What You See Is What You Meant: direct knowledge editing with natural language feedback', *Proceedings of ECAI-98*, Brighton, UK, 1998, 180–197.

Power, R. (1999), 'Controlling logical scope in text generation' *Proceedings of the European Workshop on Natural Language Generation*, Toulouse, France.

Searle, J. (1969), *Speech Acts: An Essay in the Philosophy of Language*, Cambridge University Press, Cambridge.

Sutcliffe, A.G. (1995), *Human-Computer Interface Design*, Second Edition, Macmillan, London.

Traum, D. (1999), 'Extended Abstract: 20 Questions on Dialogue Act Taxonomies', *Preproceedings of Amstelogue '99*.

van der Eijck, P. (1998), 'Controlled Languages in Technical Documentation', *Selected Papers from the Eight CLIN meeting*, Rodopi, Amsterdam, 187–203.

van Deemter, K. and R. Power (1998), 'Coreference in knowledge editing', *Proceedings of the COLING-ACL workshop on Computational Treatment of Nominals*, Montreal, Canada, 1998, 56–60.

Winkels, R., A. Boer, J. Breuker & D. Bosscher (1998), 'Assessment Based Legal Information Serving and Cooperative Dialogue in CLIME'*Proceedings of JURIX-98*, GNI, Nijmegen, Netherlands 131–146.