

An Integrated Discourse Model for a Case-Based Reasoning Dialogue System

Karolina Eliasson

Linköping University

Department of Computer and Information Science

SE - 581 83 Linköping

karel@ida.liu.se

Abstract

We present a discourse model integrated with a case-based reasoning dialogue system which learns from experience. The discourse model is capable of solving references, manage sub dialogues and respect turn takings in a dialogue in natural language. The framework is flexible enough not to disturb the learning functions, but allows dynamic changes to a large extent. The system is tested in a traffic surveillance domain together with a simulated UAV and is found to be robust and reliable.

1 Introduction

For a dialogue in natural language to run smoothly, the participants have to know the history of it. If a computer dialogue system will be able to work properly in such a natural dialogue with a human user it has to maintain a discourse model of the dialogue so far to be able to interpret the utterances of the user in the right context. The discourse model helps the system to interpret references to utterances earlier in the dialogue. The system also need to know if an utterance shall be interpreted in the earlier discourse or if it is a start of a new dialogue with a new discourse.

In this paper, we will describe a discourse model which is integrated in a case-based reasoning (CBR) system used for dialogue with a robot. Case-based reasoning is a form of machine learning where the system stores problems and their corresponding solutions in a case base. When a new target case enters the system, it searches the case base for similar cases. When the most similar case is found, its corresponding solution is adapted to the new target case and the new solution is returned. The new target case and its solution are then stored in the case base for future use. See for example [Aamodt, 1994] for an overview. Phrase matching is necessary both in CBR and in discourse modeling, in the latter to allocate incoming new phrases to the correct dialogue thread, which makes CBR and discourses modeling a suitable combination.

One long term goal of this research is to construct a CBR-system which can gain new knowledge about the world, the current domain, and the knowledge about the dialogue structure. This can be done by using CBR-techniques such as adaption of the solution but also by asking the operator for

guidance and by using interactive mixed-initiative case-base planning. These goals leave us with a demand for a knowledge intensive CBR system with the ability to have a natural conversation.

We have chosen to work on the discourse model presented in [Pfleger *et al.*, 2003] for the SmartKom project. Our structure of the discourse model as described in section 3 is highly inspired by their model. Our contribution to their work is mainly the integration of the model with CBR which is described in section 4 and 5. Section 6 works through an example to get a feeling for the discourse model and section 7 shows some result. Section 8 gives some pointers to related work and in section 9 some conclusions and ideas of future work are presented.

2 Dialogue System

CEDERIC, Case-base Enabled Dialogue Extension for Robotic Interaction Control, is a dialogue system designed for dialogue with a physical robot, in particular the WITAS autonomous unmanned aerial vehicle (UAV). The WITAS project focuses on the development of an airborne computer system that is able to make rational decisions about the continued operation of the aircraft, based on various sources of knowledge including pre-stored geographical knowledge, knowledge obtained from vision sensors, and knowledge communicated to it by data link [Doherty *et al.*, 2000]. The UAV used in the project is a Yamaha RMAX helicopter which an operator can control by high level voice commands or by written commands. The operator can ask the UAV to perform different tasks and answer questions.

CEDERIC consists of a *case base*, *domain knowledge*, a *discourse module* and a *case-base manager* as shown in Figure 1. The domain knowledge contains an ontology of the world as the robot knows it, a categorization of the world items, and a grammar. The purpose is twofold. It serves as a world representation which gives CEDERIC knowledge about which buildings there are in the known world, what kind of buildings they are, where they are places, and their attributes such as color and material. It also gives CEDERIC fundamental knowledge about which items that can be called buildings in the dialogue and which can not and provides CEDERIC with a grammar so that the system can interpret natural language.

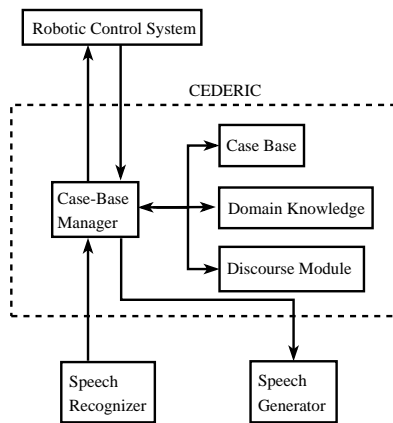


Figure 1: Architecture of CEDERIC.

The operator can choose to use either speech or text for the input to the dialogue system. The speech recognizer used is the off-the-shelf product Nuance and the speech generator used is one of the off-the-shelf products Festival or Bright-speech. When learning a new word using speech recognition, one can choose between having a considerably bigger grammar for the speech recognizer than the dialogue manager and only consider learning in the dialogue manager, or provide the new word in text form in the learning phase and then compile it into the speech recognition grammar if that can be done at runtime. We have chosen the second approach where the unknown words are provided in text and the learning phase extends the grammar.

When a new sentence arrives from the operator CEDERIC looks for cases similar to the new target case. The solution to it is either an utterance in return to the user or a request to the robotic control system. The robot acts upon the request and produces a response that is caught by CEDERIC who searches its case base and returns a message to the user. The system can manage simple cases of dialogue such as a command from the user that directly produces an answer even without a discourse model, but to be able to handle a more natural and sophisticated dialogue such as references to earlier objects and clarifying questions (where?, what?, which?, why?), a discourse model is necessary. This paper is particularly focused on the discourse model implemented in CEDERIC and how it can be used in a case-based system. Other parts of CEDERIC are not deeply explained.

The following dialogue problems are addressed in the paper:

Anaphora references. The discourse model should be able to solve references to objects which have occurred in an earlier stage of the dialogue.

Sub dialogues. It should also be able to recognize if an utterance is a sub dialogue to the present dialogue and hence should be interpreted within the limits of the current discourse or if it is the start of a new dialogue. It should also recognize a dialogue as completed which makes the old discourse no longer applicable. It should be possible to return to older non-completed dialogues which is not

presently in focus.

Turn taking. The discourse model should be able to figure out if it is a good moment to mention e.g. an observed event or if that utterance should wait for a better occasion when it does not disturb the present dialogue.

3 Discourse Model Design

The discourse model is built up of four different objects, which is linked to one another in a hierarchical manner which constitutes the meaning of the dialogue.

The linguistic objects. These objects are furthest down in the chain of objects and thus most specific on the word level. They contain information of how the nouns in the dialogue were uttered. They could for example have been references by the word *it* or by a noun and a determinant.

The discourse objects. These objects contain the different nouns together with their attributes mentioned in the dialogue. A discourse object can also be composite. An enumeration of several objects can be seen as a discourse object representing the enumeration as such and this object contains the enumerated objects as its children. This gives CEDERIC the opportunity to understand references referring to the order of the enumerations, e.g. *the first one*. The discourse objects have a link to the corresponding linguistic object.

The dialogue objects. These objects groups the sentences and their information together which have the same direct goal. The sentence *fly to the hospital* gives for example, when it is executed, a dialogue object which groups the sentences *fly to the hospital, ok and I am at the hospital now together*. If any sub dialogues come up, they will be saved in a new dialogue object with their direct goal to clarify some matter in the dialogue. Dialogue objects contain information about the topic of the dialogue, which discourse objects that were created due to the utterances, and which future utterances this dialogue object expects to consider the dialogue or the sub dialogue completed. These expectations on future dialogue are saved in a modified *initiative-response (IR) unit* [Ahrenberg *et al.*, 1991]. IR-units in our context can, unlike the original IR-units described by Ahrenberg, contain more than two subelements because they shall also be able to represent the response from the robot when the system sends a request. The *fly to the hospital* example above shows such an example.

The global focus space. The different objects in the dialogue layer which belongs to the same dialogue, including sub dialogues, are grouped together in a top object called the global focus space. It contains information about the main topic of the dialogue, if it is ok to interrupt the dialogue and which dialogue objects that belongs to it. Each global focus space also keeps track of the discourse object last mentioned, to be able to resolve references

such as *it*. This is known as the *local focus stack*. The last mentioned discourse object is said to be in focus.

To keep track of the current dialogue in focus, CEDERIC saves the different global focus spaces in a stack called the *global focus stack*. The global focus space on top of the stack is said to be the one in focus. If every IR-unit belonging to a global focus space is closed, that is, has received all its subelements, the global focus space is marked as closed and removed from the stack. Several dialogues can be open and ongoing at the same time and are thus members of the stack but only one dialogue can be in focus at the same time.

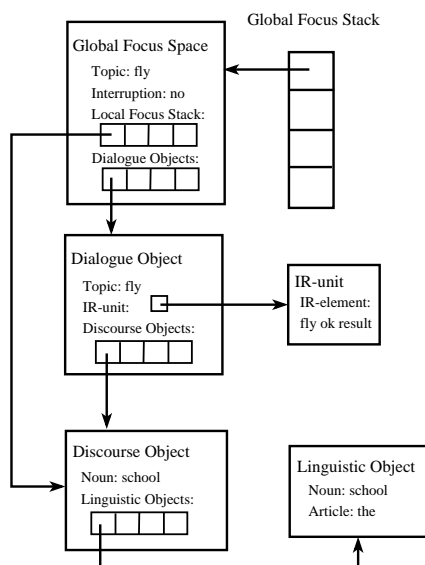


Figure 2: An example of a discourse model.

Figure 2 shows an example of how the discourse model looks like when the utterance *Fly to the school* has been executed.

4 Discourse Information in the Cases

When a new utterance enters the system, it is not only the utterance itself, but also the dialogue discourse, that tells the system how it should be interpreted. The simple answer *yes* to a question is an illustrative example of this. Without knowing the question, the answer carries no information at all. Therefore, to match a case in the case base, not only the utterance by itself but also the discourse needs to match. When a matching case is found, the system knows which information the new utterance carried and the discourse has to be updated accordingly to reflect this new information.

A case in our approach is divided into five different parts:

The problem. The problem is a description of the utterance. It contains the words and their classification according to the grammar in the domain knowledge.

The discourse information. This part describes how the global focus space in focus and its discourse object in focus should look like. It makes sure that utterances

such as answers to questions are executed with the correct case.

The update according to problem. Depending on the problem, the discourse model has to be updated with the new information. This information is stored in this part.

The solution. This part contains the reaction to the problem. It can be a request to the robot to perform an action or an answer in natural language to the operator.

The update according to solution. When the solution has been executed, the discourse model has to be updated to reflect it.

If a new dialogue is started, a new global focus space with one or more dialogue objects with corresponding IR-units, one or more discourse objects, and one or more linguistic objects are created. This newly created global focus space is put on top of the global focus stack and the local focus stack of the new global focus space is populated with the new discourse objects. Possible old open global focus spaces on the global focus stack are left in the stack as they are and are still reachable although not in focus.

If the new problem is an expected continuation of an ongoing dialogue, the case returns the newly satisfied elements of the IR-unit and CEDERIC updates the above IR-units accordingly. In case all elements in the IR-unit have been satisfied, the IR-unit is closed and CEDERIC checks if the global focus space of that IR-unit only consists of closed IR-units. In that case the whole global focus space is marked as closed.

In case CEDERIC needs to ask a clarifying question to a given problem to be able to unambiguously interpret the meaning of the operators utterance, a new dialogue is created. The new dialogue object is created in the same global focus space that matched the case, because the new dialogue is only a sub dialogue to the main one. A new IR-unit is created and possible discourse and linguistic objects are created as well. If a new discourse object is created, it is put on top of the local focus stack.

If the solution to the case is a request to the robot, the discourse model notices it and starts to expect a response from the robot.

5 Case Matching

When a new utterance from the operator or a message from the robot enters the system, it starts by classifying the included words according to the grammar. Then the case base is searched for cases with similar utterances. The current discourse in focus is matched with the discourse information saved in the case, hence a match implies that the utterance can be evaluated in the current discourse in focus.

If no case matches the new problem and the discourse currently in focus, one of the following scenarios has happened:

- The operator or the robot returns to an older open discourse.
- The operator or the robot changed topic and started a new dialogue.
- CEDERIC did not understand the new utterance either because the utterance as such is not represented in the

Operator: Fly to the school.
 CEDERIC: I have two schools to choose between.
 Which one do you mean?
 Operator: Which can I choose between?
 CEDERIC: You can choose between the one on
 Harborroad and the one on Mainstreet.
 Operator: The one on Harborroad.
 CEDERIC: Ok.
 CEDERIC: I am at the school now.

Figure 3: A dialogue example.

case base or it is totally out of context and no suitable open discourse is found.

The operator is free to change subject of the dialogue at any time by starting a new dialogue or return to an old open one, i.e. the operator can at any time take the turn. If no matching case is found using the present discourse in focus and the utterance originates from the operator, CEDERIC investigates if it is possible to match the utterance with a discourse from an earlier open dialogue. These dialogues are stored in the global focus stack. If a match is found using an old global focus space the solution to the case will be evaluated in the discourse represented by that old global focus space and it will be put in focus. It will also be updated so it correlates with the result of the evaluation of the solution.

If no matching case is found this way either, the operator did not return to an earlier dialogue and the utterance is again matched against the cases in the case base. This time with no discourse in focus at all. If matching, this indicates that the operator has started a new dialogue. If a match is found, the adapted solution is evaluated with an empty discourse and a new global focus space is put on the global focus stack. If there is no match with an empty discourse either, CEDERIC matches with a default error case.

If it, on the other hand, is a message from the robot that does not match with the present discourse in focus, CEDERIC has to take the turn taking into consideration. A report of a result of a performed command shall for example not be mentioned right away if the operator waits for an answer to a question. CEDERIC decides what to do by investigating the global focus space currently in focus and checks if it is ok to interrupt in the present discourse or not. If it is ok to interrupt, the same algorithm as the one for an utterance from the operator is performed, but if it is not, the message is put in a queue and is evaluated as soon as it is ok to interrupt or the present dialogue is closed.

6 An Example

We will present an example of a dialogue with a couple of sub dialogues and show the state of CEDERIC in the different situations to give a feeling for how the different parts fit together. The dialogue is given in Figure 3.

The operator starts by asking the UAV to fly to the school. All earlier dialogues are closed and there is no current dialogue in focus. Given the domain knowledge, CEDERIC recognizes two schools which the operator could have meant.

With this knowledge, the actual utterance and the empty discourse, the case base is gone through in the search for a matching case. Such a case is found and the update according to problem part of the case is evaluated. This part updates the discourse model by creating a new global focus space and populating it with a dialogue object describing the utterance. The dialogue object includes a discourse object which contains the noun `school`. The linguistic object created contains the noun and the determinant used. Then the solution part of the case is executed which produces a clarifying question. Now the discourse model needs to be updated again to reflect the solution utterance. This time a new dialogue object related to the clarifying question is created and linked to the global focus space already created. It is linked to the same discourse object because it is the same school they are referring to. But the word used to refer to the school is not the same. The solution utterance used the word `one` and this information is saved in a new linguistic object.

The two dialogue objects are also provided with one IR-unit each. The IR-unit of the first dialogue object is a list of the elements `fly`, `ok` and `result` and the next expected element is `ok`. The other IR-unit is a list of the elements `mean` and `specify` where the next expected element is `specify`.

The operator answers by asking a counter question, more exactly by asking which (schools) he or she can choose between. CEDERIC matches the utterance together with the current discourse and finds a suitable case. This case states that the two known schools should be enumerated. The discourse is updated and there is now one additional dialogue object corresponding to the new sub dialogue started by the operator. This dialogue object contains a discourse object representing a sequence of objects. This object has in turn references to two other discourse objects representing the two schools. The sequence object is put on top of the local stack. A new IR-unit with `choose` and `chooserresponse` as elements is created and coupled to the new dialogue object. Since this IR-unit already has got all the elements due to CEDERIC answering the question, it is marked as completed directly.

The operator selects one of the schools by specifying the unique attribute of that particular school. CEDERIC is now matching the utterance together with the discourse information stating that the discourse will be a discourse object consisting of a sequence with several other objects. A matching case is found which creates an acknowledgement in form of an `ok` to the operator and a request to the robot to fly to the school on Harborroad. The discourse is updated partly by putting the chosen discourse object on top of the local focus stack and partly by the case returning the IR-elements `specify` and `ok` making CEDERIC update the concerned IR-units and set the interruption status to `yes`. The only IR-unit which is not yet completed is the one with elements `fly`, `ok` and `result`. When the UAV arrives at the school, it reports the event to CEDERIC. CEDERIC matches this response together with the current discourse and finds a matching case. This case reports the event to the operator and notifies the discourse that the IR-element `result` has been satisfied. The discourse sets the last IR-unit as completed and the whole global focus space is then marked as closed and popped from the global focus stack.

Operator: Fly to the school.
 CEDERIC: I have two schools to choose between.
 Which one do you mean?
 Operator: Take off.
 CEDERIC: Ok.
 Operator: Which can I choose between.

CEDERIC gets a message from the robot saying that the action take off has been successfully completed

CEDERIC: You can choose between the one on Harborroad and the one on Mainstreet.
 Operator: Fly to the hospital.
 CEDERIC: Ok.
 CEDERIC: I have taken off now.
 Operator: What is your altitude?
 CEDERIC: It is 20 meters.
 CEDERIC: I am at the hospital now.

Figure 4: An example of turn taking and dialogue topic changes.

7 Result

CEDERIC has been tested connected to a simulated UAV situated in a simulated environment. The simulated UAV can perform high level actions such as flying to a building with a certain unique identity, flying in a certain direction, take off, land, ascend and descend. It reports the result of an action and also report if it observes any buildings as it flies. It is able to answer status questions such as the current altitude, velocity and heading etc.

A number of dialogues have been implemented where both references, dialogue topic changes, and turn taking have been carefully tested. Figure 4 shows an example where the operator switches dialogue topic back and forth. CEDERIC is able to keep the different discourses in mind and use the correct discourse model for every new utterance. It is also an example of a case where CEDERIC gets a message from the robot, but due to the current dialogue, the message is restrained until a suiting moment in the dialogue appears. The operator's utterances are never restrained and he or she is free to lead the dialogue as he or she wishes.

More complex examples where the operator teaches CEDERIC new information have also been tested. An example of such a dialogue is provided in Figure 5, where the operator teaches CEDERIC a new word and a domain item, in this case a new building. The result of this dialogue is, besides the movement of the robot, an update of the grammar and the domain knowledge. As seen in the end of the dialogue, CEDERIC can successfully fly to the church after the explaining dialogue.

The tests have proven the discourse model to be robust with ability to cope with various types of operator input as partly shown above. Due to the flexibility of the information stored in the discourse model it is well suited for CBR systems and various learning strategies.

Operator: Fly to the church.
 CEDERIC: I do not know what the noun church means.
 Is it a building?
 Operator: Yes.
 CEDERIC: Ok, where is the church you want me to fly to?
 Operator: East of the hospital.
 CEDERIC: Ok, I will start by flying to the hospital then.
 CEDERIC: I am at the hospital now.
 Operator: Fly east and look for a white building.
 CEDERIC: Ok, I'll stop when I see such a building.
 CEDERIC: I see the church now and hover over it.
 Operator: Fly to the school on Mainstreet.
 CEDERIC: Ok.
 CEDERIC: I am at the school now.
 Operator: Fly to the church.
 CEDERIC: Ok.
 CEDERIC: I am at the church now.

Figure 5: An example of a teaching situation between the operator and CEDERIC.

8 Related Work

Dialogue in CBR is commonly used in conversational CBR where the system asks the user questions which guides the search for the most similar case in the case base [Shimazu, 2002; Aha *et al.*, 2001; McSherry, 2003; Carenini *et al.*, 2003; Bridge, 2002; Göker and Thompson., 2000]. This approach is commonly used when the user is querying a system with products or in help desk scenarios using a graphical user interface. Conversational CBR has lately gained more interest in other areas as well, such as conversational CBR with natural language. As the systems built on conversational CBR get more complex, we believe the need for a discourse model grows. Some work has already been done in that direction, e.g. Brantings discourse model for conversational CBR [Branting *et al.*, 2004]. The dialogue in conversational CBR is however characteristically a simple question-answer dialogue where the questions and their corresponding answers are saved in the cases as a description of the item. This separates the discourse model from the CBR procedure and limits the use of CBR for dialogue purposes.

Because our CBR-system for dialogue with a robot is not a pure conversational CBR system, but has with respect to its use of dialogue more in common with non-learning dialogue systems such as [Allen *et al.*, 2001; Rosset and Lamel, 1999; Seneff, 2002], we have integrated a discourse model built on the traditional principles with CBR.

Within the WITAS project, several dialogue systems with various capabilities have been developed. The first WITAS Dialogue System [Lemon *et al.*, 2001] was a system for multi-threaded robot dialogue using spoken I/O. The DOSAR-1 system [Sandewall *et al.*, 2003] was a new implementation using another architecture and a logical base. This system has been extended into the current OPAS system. Our work takes a rather different approach to discourse modelling, compared to these predecessors, as we are in-

tegrating CBR techniques, but it reuses major parts or the OPAS implementation for other aspects of the system. For additional information, please refer to the WITAS web site at <http://www.ida.liu.se/ext/witas/>.

9 Conclusion and Future Work

We present a discourse model which is integrated with a CBR-system for communication with a robot. We have shown how the cases updates the discourse model which gives scope for learning of new dialogues and dialogue structures within the loose framework the discourse model defines. This way, we can control which cases matches the new problem not just by comparing the problem statements but also by comparing the discourse, which gives us the opportunity to solve problems such as references, sub dialogues and turn taking in a learning system.

Our implementation has been tested connected to a simulated UAV operating in a simulated environment. The resulting system is robust and allows the operator to take the initiative in the dialogue at any time without losing track of the discourse. It has also proven easy to work with and new cases can easily be automatically generated from new target case problem, the adapted discourse description, the adapted solution and the adapted discourse update. In fact, the adapted discourse description is generated per se because it is the same discourse as the one currently in focus.

The integrated discourse model is an aid for our primary goal to design a dialogue system not only capable of learning in a restricted area but to be able to handle a large amount of utterances and advanced dialogue both from the operator and from the robot. The advanced dialogue features provides a platform for further research regarding giving the operator the opportunity to explain new domain and dialogue knowledge to the system and the ability for the system to ask for confirmation to a solution. Another interesting area is mixed initiative case-based planning with user interaction where the system is able to reuse parts of plans from several cases in the case base to solve a new planning problem. Using the dialogue features, the operator can guide the search and teach the system new information during a planning action.

Acknowledgement

This research work was funded by the Knut and Alice Wallenberg Foundation, Sweden.

References

- [Aamodt, 1994] Agnar Aamodt. Case-based reasoning; foundational issues, methodological variations, and system approaches. *AI Communications*, 7(1):39–59, 1994.
- [Aha et al., 2001] David W. Aha, Leonard A. Breslow, and Hector Munoz-Avila. Conversational case-based reasoning. *Applied Intelligence*, 14(1):9–32, 2001.
- [Ahrenberg et al., 1991] Lars Ahrenberg, Arne Jönsson, and Nils Dahlbäck. Discourse representation and discourse management for a natural language dialogue system. Technical report, Institutionen för Datavetenskap, Universitetet och Tekniska Högskolan Linköping, 1991.
- [Allen et al., 2001] James Allen, George Ferguson, and Amanda Stent. An architecture for more realistic conversational systems. In *IUI '01: Proceedings of the 6th international conference on Intelligent user interfaces*, pages 1–8. ACM Press, 2001.
- [Branting et al., 2004] Karl Branting, James Lester, and Bradford Mott. Dialogue management for conversational case-based reasoning. In *Proceedings of the Seventh European Conference on Case-Based Reasoning*, 2004.
- [Bridge, 2002] David Bridge. Towards conversational recommender systems: a dialogue grammar approach. In *In Proceedings of the EWCBR-02 Workshop on Mixed-Initiative Case-Based Reasoning*, 2002.
- [Carenini et al., 2003] Giuseppe Carenini, Jocelyn Smith, and David Poole. Towards more conversational and truly collaborative recommender systems. In *Proceedings of the 7th International Conference on Intelligent User Interfaces*, 2003.
- [Doherty et al., 2000] Patrick Doherty, Gösta Granlund, Krzysztof Kuchinski, Erik Sandewall, Klas Nordberg, Erik Skarman, and Johan Wiklund. The witas unmanned aerial vehicle project. In *Proceedings of the 12th European Conference on Artificial Intelligence*, 2000.
- [Göker and Thompson., 2000] Mehmet Göker and Cynthia Thompson. The adaptive place advisor: A conversational recommendation system. In *Proceedings of the 8th German Workshop on Case Based Reasoning*, 2000.
- [Lemon et al., 2001] Oliver Lemon, Anne Bracy, Alexander Gruenstein, and Stanley Peters. The WITAS multi-modal dialogue system. In *Proceedings of EuroSpeech*, 2001.
- [McSherry, 2003] David McSherry. Explanation of retrieval mismatches in recommender system dialogues. In *Proceedings of the ICCBR-03 Workshop on Mixed-Initiative Case-Based Reasoning*, 2003.
- [Pfleger et al., 2003] Norbert Pfleger, Jan Alexandersson, and Tilman Becker. A robust and generic discourse model for multimodal dialogue. In *Workshop Notes of the IJCAI-03 Workshop on Knowledge and Reasoning in Practical Dialogue Systems*, 2003.
- [Rosset and Lamel, 1999] Sophie Rosset and Samir Benacef Lori Lamel. Design strategies for spoken language dialog systems. In *Proceedings of EuroSpeech*, 1999.
- [Sandewall et al., 2003] Erik Sandewall, Patrick Doherty, Oliver Lemon, and Stanley Peters. Words at the right time: Real-time dialogues with the witas unmanned aerial vehicle. In *Proceedings of the 26th Annual German Conference in AI*, 2003.
- [Seneff, 2002] Stephanie Seneff. Response planning and generation in the MERCURY flight reservation system. *Computer Speech and Language*, 16(1):283–312, 2002.
- [Shimazu, 2002] Hideo Shimazu. ExpertClerk: A conversational case-based reasoning tool for developing salesclerk agents in e-commerce webshops. *Artif. Intell. Rev.*, 18(3-4):223–244, 2002.