

Understanding narrative is like observing agents

Guido Boella, Rossana Damiano and Leonardo Lesmo

Dipartimento di Informatica - Università di Torino

Cso Svizzera 185 10149 Torino, ITALY

email: {guido,rossana,lesmo}@di.unito.it

Abstract

In this paper, we suggest that AI techniques, especially the ones developed in the field of agent modeling and intelligent interfaces, can be exploited both to build systems for narrative understanding and to allow agents to describe their own behavior in a narrative style.

In particular, we will show how it is possible to exploit a model of dialog interpretation for building a system that understands stories and produces a representation of the characters' plans and intentions.

Introduction

(Bruner, 1991) claims that the interpretation of intentional behavior of people works by structuring it into narratives, that is, by finding the relations among the different actions a person does and by inferring what he thinks or feels.

We believe that, while Bruner is right in establishing a relationship between the interpretation of intentional behavior and narrative, the most basic activity is the interpretation of intentional behavior. In fact, it is standard practice in everyday life to try and infer what are the intentions of other people by observing their actions (for instance, think of driving in the traffic). Moreover, in many cases, intention recognition is needed to understand the coherence links underlying complex action sequences. As far as narrative is concerned, humans have to do the same: the only difference is that they have as their starting point not direct observations of actions but the description or representation of actions provided by the narrator. The narrator, who must select the relevant information and decide what aspects of the behavior of a character are worth describing, exploits this ability of humans to reconstruct the intentions of a character from the known facts.

In this paper, we suggest that AI techniques, especially the ones developed in the field of agent modeling and intelligent interfaces, can be exploited both to build systems for narrative understanding and to allow agents to describe their own behavior in a narrative style.

In particular, we will show how it is possible to exploit a model of dialog interpretation for building a system that understands stories and produces a representation of the characters' plans and intentions.

First of all, the notion of belief-desire-and-intention (BDI) agents seems to be relevant to dialog interpretation. The aim of BDI models is to build autonomous systems able to mimic human practical reasoning (Bratman et al., 1988): given a set of beliefs, goals and plans, an autonomous agent should pursue his objectives; in case this requires interaction with other agents, the agent builds a model of his partners in terms of their beliefs, goals and of the plans they are carrying on.

Moreover, reconstructing other agents' intentions is required not only to cooperate with other agents, but also to oppose antagonists' moves in conflictual situations.

In order to implement intelligent interfaces that can interact with humans in a natural way, agent theories are exploited as the basis for building a model of the user's goals and plans starting from his observed actions (speech acts but also non linguistic actions executed in the domain in which the interaction takes place); various techniques have been developed for *plan and intention recognition*, i.e. to understand the agent's goals which have led him to act.

This kind of reasoning is particularly useful in case of communication: as it has been noticed since (Allen, 1983), when a user asks a question to a dialog system, he does not expect just a literal answer, but he wants also that his domain goals are taken into account: for example, from a question concerning the location of a library it is possible to infer that (probably) the user wants to go to the library and borrow some books. Further observations - like a question about the opening hours of the library - can support this hypothesis.

In particular, it is assumed that human behavior proceeds smoothly following a plan, since an agent's intentions tend to be steady, unless some other events occur: in order to relate each new action to the previous context, it is better to examine first the focus of the agent activity (i.e. the current subgoal) and, only if it is not possible to find a connection with the observed action, the interpretation resorts to examining the higher goals of the agent (Carberry, 1990).

```

Get-to-do(A, B, Answer(B, Phone-call)) --> Intend(B, Answer(B, Phone-call))
|
Request(A, B, Answer(B, Phone-call)) --> MB(B, A, Intend(A, Intend(B, Answer(B, Phone-call))))
|
Utterance-act(A, B, "Answer the phone")

```

Figure 1: The representation of speech acts

A plan based model of dialog

In (Ardissono et al., 1998b) we presented a model of dialog where interaction is reduced to the notion of co-operation: a move (either linguistic or not) is coherent with respect to the previous part of the interaction only if it is performed as a means to satisfy one of the goals that the two interactants have (implicitly or explicitly) put forth in the previous part of the interaction.

This definition implies that, in order to take part in an interaction, an agent must reconstruct his partners' goals; moreover, the agent must be able to access his own internal states in order to understand if and how his partners' turns contribute to his own activity. To satisfy these cooperation requirements, we have introduced in this model a declarative representation of agents' behavior that is exploited both to reconstruct the agents' activity and to produce an agent's behavior.

In particular, this agent model is composed of a plan-based representation which concerns not only the knowledge about linguistic and domain actions, but also the way this knowledge is used, i.e. a plan-based representation of how the agent's intentions and plans evolve during his activity (Ardissono and Boella, 1998).

In addition, the same libraries of plans that serve as the knowledge base for inferring others' goals are used as the "set of instructions" for making an agent act in the world. After the execution of actions the agent maintains the instances of the performed plans, i.e. a declarative representation of his own activity, that he can use to reason about himself. In particular, by exploiting the direct correspondence between verbs and actions that we will discuss below, the agent can describe what he has done. Therefore, the narrative understanding ability relies on the same knowledge structures and reasoning mechanisms that are used for recognizing other agents' behavior, producing an agent's own actions and describing the agent's own activity.

Starting from the dialog model and agent architecture, we have implemented a dialog interpretation system that takes in input the sentences or observed actions of two interactants and produce a representation of the dialog. The output is composed of plan instances representing the intentions of the two interactants together with the coherence links that relate their goals. This model has been extended to account for a comprehension of the actions of the characters involved in a story: the characters' actions described by the narrator are given in input to the dialog interpretation system that incrementally produces a representation of the described agents' plans.

The action based representation

The basic idea is that, in order to interpret a verb concerning an action, an action instance must be built. By representing the described event as an action we gain two advantages: first, from the action definition we can get information concerning action preconditions, effects and how the action can be carried out by means of a sequence of steps; second, the plan based representation of the described event can be given in input to the plan recognition mechanism.

The plan recognition phase allows to understand the intentions underlying the execution of the input action, whose knowledge is necessary in order to understand what is going on in the described situation. In fact, the narrator can describe the situation in different manners, without necessarily explaining what the intentions of the described agent are: for example, in order to create suspense about what will happen next. Or he can just describe the agent's activity at a high level, without going into detail.

Let's consider an example, which aims both at giving a flavor of how our approach works and at explaining the current point. Let's suppose that in a situation to be described a character *A* utters to *B* *answer the phone*. This event is interpreted by the dialog interpreter by means of a set of actions that are related by generation relations (Pollack, 1990). First, the observed action of uttering a sentence, then the action of expressing a propositional content (the locutionary act of (Austin, 1962)), which involves the goal of making mutually believed (*MB*) the speaker's intention that *B* performs the action and finally the perlocutionary action of convincing the hearer (*B*) to answer the phone. The representation produced by the dialog interpreter can be sketched in Figure 1 (the vertical lines represent the generation relation, i.e. a "Request" is a means for executing a "Get-to-do action", and the horizontal arrows identify the effects of the actions).

This representation is connected to higher domain-dependent levels, in a scenario that specifies, for example, that the agent *A* wants to know some important information from the (hypothesized) caller.

Given this situation, the narrator *S* can depict it by describing more explicitly any of these interlaced levels, knowing that the final representation of *A*'s intentions built by *R* will be the same in all cases: that *A* wants to know who is calling and what he wants to tell him. *S* can say as well:

1. *A told B to answer the phone* Utterance-act
2. *A requested B to answer the phone* Request
3. *A convinced B to answer the phone* Get-to-do
4. *B wanted to know what the caller wants to say* Satisfy

In case 1, the verb is interpreted as the linguistic action of uttering a sentence (Utterance-act), i.e. the same kind of action that is given in input to the dialog interpreter when a user sentence is typed on the keyboard. In both cases, starting from that utterance action the plan recognition algorithm recognizes the illocutionary act of requesting. This is the level that the verb “request” aims at describing: an instance of the linguistic action of “Request” is created.

What if the story continued with the description of the action *B walked to the telephone*? the intention recognition mechanism formulates the *hypothesis* that *B* has accepted *A*’s request: seemingly, he could be walking to the phone in order to make one of the pre-conditions of the *answer the phone* action true (being close to the phone). Under this interpretation, the story proceeds coherently, since the actions of the two characters are related.

However, this is just an hypothesis, maybe chosen according to the probability assigned to actions by the plan recognition mechanism. Perhaps, the narrator is trying to mislead the reader, and he can go on with some unexpected event: *B unplugged the phone*. If this is the case, the reader has to revise his hypothesis and backtrack to another explanation that accounts for the coherence of the whole story:¹ *B* is not cooperating and is trying to make the goal of *A* (knowing the information that the caller can provide) false. However, the entire story is still coherent, since the actions of *A* and *B* are related by the two contrasting goals that *A* knows or does not know some information.

Beside this level of coherence, there is another one that must be taken into account: the connection among the linguistic actions performed by the narrator. In fact, he is not producing just a sequence of unrelated speech acts concerning a coherent content (the actions of *A* and *B*), but, instead, he is following a detailed discourse plan in which each step is linked to the others by rhetorical relations (Moore and Paris, 1988). In general, the actions are related by the goal of describing a sequence of actions of the character, but also more specific goals can be achieved by discourse plans. In case of *A asked B to answer the phone; he was waiting for an important information* the two sentences are related by a *motivation* relation: the second one explains why *A* is requesting *B* to answer the phone. In *when the*

¹Intentions tend to persist, therefore if *B* changed his mind while approaching the phone the narrator should make the reader aware of this fact, otherwise he knows that the reader will adopt the interpretation where agents do not change their mind, by inferring that *B* already had the intention of unplugging the phone when he started moving towards it: *B walked to the telephone in order to answer, but then he changed his mind and unplugged the phone*.

phone started ringing. A asked B to answer the phone a *circumstance* relation links the subordinate clause to the main one. Finally, rhetorical reasons motivate the possibility, for the narrator, to describe events in a different order from the temporal sequencing, relying on the reader’s ability to reconstruct the right temporal order and the correct relations between the described actions and events.

Moreover, different effects can be achieved by the narrator: in the previous example, he is zooming on *B*’s actions (*B approached the phone*) without explaining what he is going to do in order to create suspense and maybe mislead his reader. Another reason that could induce the narrator to adopt such description style is the goal of describing events in a more natural way, as the reader was observing the actions himself.

In our proposal, both levels of intentions are taken into account: plan recognition is the reasoning mechanism exploited in relating the characters’ actions and the narrator’s ones; the difference lies in the knowledge about actions exploited in the two tasks: speech acts and non-linguistic actions in the first case, discourse plans in the second one.

Is this enough to understand narrative?

In this section we will consider how our model of narrative understanding fits into some of the properties characterizing narrative that (Bruner, 1991) has highlighted (we refer to the definitions reported in (Sengers, 1998), page 145).

- *Narrative diachronicity*: “Narratives do not focus on events moment-by-moment, but on how they relate over time.” The incremental and focusing-based plan recognition techniques and the coherence-seeking model of dialog are inherently oriented towards looking for the relations among the actions over time. As noticed by (Sengers, 1998), This is possible only using a deliberative agent paradigm, as our proposal does (Bratman et al., 1988): deliberative agents generate a structured behavior following a plan-based representation of actions instead of stimulus-response rules, as in case of reactive agents (Brooks, 1991).²
- *Intentional state entailment*: “When people are acting in a narrative the important part is not what the people do, but how they think and feel about what they do.” Clearly, the plan recognition mechanism has the goal to recognize what beliefs and intentions have led an agent to act: it is possible to reach a certain degree of understanding of why people act, or at least to attribute beliefs and intentions to agents.

²(Sengers, 1998) criticizes deliberative agents for their modularity deriving from the divide and conquer techniques applied for building agents. We believe that our approach is more integrated: for example, the natural language interpretation phase, the recognition of other agents behavior and agent behavior generation are all integrated in a plan based approach.

even if the hypotheses are not correct or the described agents are not designed as intention based.

- *Hermeneutic Composability*: "... actions are understood with respect to how they fit into narratives as a whole. ... understanding narrative requires interpretation in a gradual and dialectical process of understanding." As stated before, the plan recognition algorithm works in an incremental manner. As we have seen, previous hypotheses can be withdrawn: this means that the same revision mechanisms that are exploited for modeling misunderstandings in dialog could apply to misunderstandings of narrative (see (Ardissono et al., 1998a)).
- *Canonicity and breach*: "There is a tension in narrative between what we expect to happen, and what actually happens." the reader knows (since "Things cannot go as planned") that agents will maintain their intentions until they are not satisfied, they become impossible to achieve or they are not relevant anymore; therefore the reader has expectations about what agents will do when unexpected events occur.
- *Normativeness*: "Narratives depend strongly on the audience's conventional expectations about plot and behavior." See above.
- *Context sensitivity and negotiability*: "narrative is negotiated between reader and text." See the reference to misunderstandings.

Conclusions

This paper presents an approach to narrative understanding based on the recognition of agents' plans and goals. As we have seen, the approach enables us to match some requirements put forth in the analysis of narratives. Starting from the dialog interpreter, presented in (Ardissono et al., 1998b), we have implemented a story understanding system, that exploits plan recognition to identify both the narrator's discourse plans and the relations between the described events. The work described in the paper is part of a larger research program that ranges from dialog modeling to lexical semantics. In particular, we exploit plan based representation of action verbs in order to explain aspectual phenomena and lexically triggered presuppositions.

We just want to make a final remark about (Sengers, 1998)'s idea to make agents "give off cues that are easy for users to understand as coherent, intentional behavior". We do not think this is necessary in case of deliberative agents: they follow by themselves a coherent intentional behavior. Instead, we believe that Sengers's idea applies well to defining the narrator task: in fact, a coherent agent behavior may appear uncoherent to a user since it is difficult to interpret, unexpected or because he lacks the necessary background to understand what is going on; it is up to the narrator to enhance the appearance of coherence by selecting the right level

of description of the character intentions, trying to say enough but not everything.

References

- Allen, J. (1983). Recognizing intentions from natural language utterances. In Brady, M. and Berwick, R., editors. *Computational models of discourse*, pages 107–166. MIT Press.
- Ardissono, L. and Boella, G. (1998). An agent model for NL dialog interfaces. In *Lecture Notes in Artificial Intelligence n. 1480: Artificial Intelligence: Methodology, Systems and Applications*, pages 14–27. Springer Verlag, Berlin.
- Ardissono, L., Boella, G., and Damiano, R. (1998a). A plan-based model of misunderstandings in cooperative dialogue. *International Journal of Human-Computer Studies*, 48:649–679.
- Ardissono, L., Boella, G., and Lesmo, L. (1998b). An agent architecture for NL dialog modeling. In *Proc. Second Workshop on Human-Computer Conversation*, Bellagio, Italy.
- Austin, J. (1962). *How to Do Things with Words*. Harvard University Press, Cambridge, Mass.
- Bratman, M., Israel, D., and Pollack, M. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4:349–355.
- Brooks, R. (1991). Intelligence without reason. In *Proc. 12th IJCAI*, pages 589–595, Sydney.
- Bruner, J. (1991). The narrative construction of reality. *Critical Inquiry*, 18(1):1–21.
- Carberry, S. (1990). *Plan Recognition in Natural Language Dialogue*. MIT Press.
- Moore, J. and Paris, C. (1988). Constructing coherent text using rhetorical relations. In *Proc. 10th Annual Conf. Cognitive Science Society*.
- Pollack, M. (1990). Plans as complex mental attitudes. In Cohen, P., Morgan, J., and Pollack, M., editors. *Intentions in communication*, pages 77–103. MIT Press.
- Sengers, P. (1998). *Anti-Bozology: Agent Design in Cultural Context*. PhD thesis, Carnegie Mellon University.