

Mixed Narrative and Dialog Content Planning Based on BDI Agents

Carlos León, Samer Hassan, Pablo Gervás, and Juan Pavón

Departamento de Ingeniería del Software e Inteligencia Artificial
Universidad Complutense de Madrid
{cleon,samer}@fdi.ucm.es, {pgervas,jpavon}@sip.ucm.es

Abstract. There exist various narrative systems, focused on different parts of the complex process of story generation. Some of them are oriented to *content planning*, and some to *sentence planning*, with different properties and characteristics. In this paper we propose a system based on BDI agents that generates stories (creating content, performing *content planning* and simple *sentence planning*) with narrative parts and dialogs. The content for the story is generated in a multiagent social simulation system, and the *content planning* is based on rules and a state space search algorithm based on the system representation of the reader's perception of the story.

1 Introduction

It is possible to find in the literature several types of narrative systems. These systems try to emulate the human activity of creating readable texts from a set of stored facts or data, organised in several possible ways. There are proposals mainly focused on narrative generation for storytelling, with different characteristics.

In storytelling, dialogs carry much information not present in simple narrative text. Dialogs show many aspects of the characters in a very different way that descriptive sentences do, because literal sentences in dialogs show exactly what the character says, and the form it expresses the content. Then, for a story to be fully descriptive, it is necessary to include dialogs that show interaction based on communication between the characters.

There exist various approaches to story generation, and they are focused on different stages of the generation process. There are systems that propose alternatives for *content planning*, like those in [1,2], but they do not offer dialog generation, and do not address *sentence planning*. Other systems, like [3], can handle dialog generation, but they do not perform *content planning* operations, as they are only focused on *sentence planning*.

In this paper we propose a system that addresses *content planning* for *dialogs* together with *narrative text*, in a coherent way. The work presented is divided in two main applications: a multiagent system that simulates social interaction between a set of characters in a fantastic medieval domain, modelled as intelligent agents with *beliefs*, *desires* and *intentions*, and an automatic story generation system, that receives the set of facts that happened in the simulation, and creates a textual representation of the main events.

2 Previous Work

In this section we briefly review research relevant for the work presented in this paper: BDI model, Multiagent Social Systems, and Natural Language Generation.

2.1 Social Systems and BDI Model

Social phenomena are extremely complicated and unpredictable, since they involve complex interaction and mutual interdependence networks. A social system consists of a collection of individuals that interact among them, evolving autonomously and motivated by their own beliefs and personal goals, and the circumstances of their social environment.

A multi-agent system (MAS) consists of a set of autonomous software entities (the agents) that interact among them and with their environment. Autonomy means that agents are active entities that can take their own decisions. The agent paradigm assimilates quite well to the individual in a social system, so it can be used to simulate them, exploring the complexity of social dynamics. In fact, there are numerous works in agent theory on organisational issues of MAS. Also, theories from the field of Psychology have been incorporated to design agent behaviour, the most extended being the Beliefs-Desires-Intentions (BDI) model, in the work of [4].

The MAS described in this paper has been developed as an extension of an existing one by enriching it with additional features to support dialogs between agents. In this MAS, as explained in [5], the agents have been developed with several main attributes: from simple ones such as gender or age, to complex ones, like for example ideology or educational level. The population in the agents' society (as in real societies) also experiments demographic changes: individuals are subject to a life-cycle: they get married, reproduce and die. Moreover, the agents/individuals can build and be part of relational groups with other agents: they can communicate with other close agents, leading to friendship relationships determined by the rate of similarity. Or, on the other hand, they can build family nuclei as children are born close to their parents.

The system has an underlying sociological model, which can be configured (i.e. changing the mean of male average age of death). It deals with hundreds of agents, all of them interacting with their neighbourhood in each time-step of the execution. As we simulate thousands of time-steps, so many dynamic factors are involved, and as any other complex systems, it is not subject to laws, but to trends, which can affect individuals in a probabilistic way.

As we will see in a further section, these simple agents will be transformed into BDI model agents. In this model, Beliefs represent the knowledge of the agent about his world, including his own actual state and the environment state. To use the term "belief" implies that what the agent believes does not have to be necessarily true (and in fact it can change through time). It is "what I know and believe". The Desires (objectives) represent the state that the agent is trying to reach, the situations that it is seeking: "What I want". The Intentions (plans) are the means that the agent choose to accomplish its objectives, what the agent has chosen to do: its goal. The BDI of a single agent is its mental state.

We can refer to multiple works that use this psychological modelling perspective. A review comparison between several modelling techniques can be found in [6]. A very useful approach to sociologic agents using BDI, with deep and sophisticated social behaviour, was studied by [7]. A common problem using BDI is the huge gap between a clear theoretical model and its implementation. An effort on this subject was done in [8], trying to clarify how to turn BDI theory concepts to the practical complex system.

2.2 The Multiagent System: Fantastic Society

The ideas expressed in section 2.1 concerning social simulations using multi-agent systems are the core of action from which we have built the whole narrative system. Several changes to the original MAS have to be made in the perspective of execution to be able to generate “life logs” of the individuals, which will be the basis for the texts describing the storyline. It is necessary to shift the point of view from trends data acquisition to vital biographies. We do not need numerical data, but semantic content that can be interpreted by the rules as we interpret them, because we want the story generation to be as close as possible to what humans might have done faced with similar sets of events.

In this framework, it was necessary to adapt the designed MAS to a new environment: a Fantasy Medieval World far from the previous Post-Modern context. This deep change of context meant the introduction of several minor changes, and a twist in the meaning and interpretation of some facts. This evolution is explained in depth in [9].

Thus, now the agents have a name, an inheritable last name, a race (elf, orc, dwarf...). For each agent now exists a random possibility of dying, allowing the possibility that we can relate this early death to the betrayal of a friend, poisoning by a wife, a mysterious accident...

Following the cited objective of emulating real life behaviours, in this new MAS context dependent relationships and life events has been introduced: usual life events were not exciting enough to build a fantasy adventure. And so, an individual can have friends and enemies. Along his path, he can get married and have children, but he also can, randomly, suffer several spells, kill monsters, get lost in mazes or dark forests, or find treasures and magic objects in dangerous dungeons. In this way we can build a more interesting story, with several characters that evolve and interact among them.

2.3 Natural Language Generation

Natural Language Generation is important for a study of storytelling because it involves both a model of the task that need to be carried out to generate a valid text - therefore partially modelling the activity of an author - and a model of the story as linguistic artifact - a story model. The general process of text generation is defined in [10] as taking place in several stages, during which the conceptual input is progressively refined by adding information that will shape the final text. During the initial stages the concepts and messages that will appear in the final

content are decided and these messages are organised into a specific order and structure (*content planning*). A number of stages after that deal with *sentence planning*, where each message resulting from the previous stage is progressively enriched with all the linguistic information required to realize it. A final stage of *surface realization* assembles all the relevant pieces into linguistically and typographically correct text.

The Natural Language Generation work presented in this paper is mainly centered around *content planning*. The subtask of building a set of messages from the input is always heavily dependent on the particular domain of operation, and tightly coupled with the particular kind of input being processed. A slightly more generic subtask determines the ordering and rhetorical relations of the logical messages, hereafter called facts, that the generated document is intended to convey. Related work with *content planning* this paper can be found on [11].

Research on storytelling has been a classic since the early days of Artificial Intelligence. In recent times, two approaches stand out as worth mentioning in contrast with the one followed in this paper. MEXICA [2] follows a cognitive approach to build story plots about the ancient inhabitants of Mexico City, paying special attention to evolution of the emotional links and the tensions between the characters of a story. A different approach [1] relies on planning techniques to build stories from an initial world state and a set of goals to be fulfilled by the end of the story. Both are comparable to the *content planning* presented here - rather than the application as a whole - in the sense that they concentrate solely on building a conceptual representation, with only schematic transcription as text intended to make the plots understandable. None of them includes dialogs as part of the content they handle. The STORYBOOK system [3] - at the opposite extreme - considers only the production of text from an input corresponding to the conceptual representation of a given plot. It relies on having an external planner that defines the outline of the intended story, and it carries out elaborated sentence planning to produce input for a surface realizer, and it is capable of dealing with fragments of dialogue represented conceptually in the input.

3 BDI Model in Storytelling

In this section we explain the ideas we have followed for developing a system based on BDI agents. We describe the agents, and how their beliefs, desires and intentions guide their behaviour and the story generation. In the next two Sections (4 and 5), we explain how this BDI model is used in the multi-agent system and the *content planner*.

We have introduced new changes in the system described in Section 2.2, changing the structure of the original simple agents. In this way the interactions will be more complex and we will have much more contents to tell in our stories. Thus, we take as base system the previously defined fantastic world, with its events and enemies. The new objective is to make agents evolve in time (internally, not just in terms of the relationships between them). First, we force agents'

Table 1. Example rules from transition between reader’s perception of the story

Fact	Agent’s mental state
know_where	Agent now knows some location (belief)
want_find	Agent wants to find something (desire)
do_trick	Agent tricks somebody (intention)

characteristics to change internally depending on the events that happened in their lives. For example, if an agent finds a treasure, his economy will be dramatically increased.

But the really deep change is the use of a BDI model inside the agents. Now the agents will have “What I know/believe” (like “I know where a treasure is”), “What I want” (“I want a treasure”) and “What I am going to do” (“I am going to the mountain to find the treasure”). With these ideas, we have agents with specific knowledge, with objectives to accomplish and with actions to do. Table 1 shows a little example of how facts in the story affect the BDI state of the character, and thus, reader’s perception of the story.

To guide the content planner, we consider that in each stage of the discourse, as the reader sees it, we have to know, at least, the information about characters in the story, at every stage of the discourse. This information evolves as the story is being told, in such a way that the reader’s concept of each character changes during the narration. For this purpose we store, in each stage of the narration, the BDI state of the characters as the reader should have understood it from what has been communicated so far. In Section 5 more detail about this is given.

4 Agent Planning

With the BDI model, each agent is “more intelligent”, taking his own decisions, and building a real story. Thus, for example, if a character wants (desire) to kill a dragon, he will ask his friends about information related to dragons (asking if they have such beliefs). When he finds someone that knows something, his friend will tell it to him (that agent will throw a “say” event), and he will have the new belief (for ex. “where to find a dragon”), changing his mental state (and throwing a “know” event, that as all events will be recorded in the XML file). When he discovers (after gathering the info between his friends) enough information about dragons, (their weak points, their types and where he can find one) he will be able to generate the intentions associated to that desire. Those intentions could be “travel to the north”, “go to the highest mountain” and “kill the dragon”. When he generates those events (“do” events), his desire will be satisfied and will disappear.

This ideas guide the planning and the generation of the characters’ story, during the execution of the multiagent system. We can see that the planning of the agents is very simple, but it is enough for the prototype to generate coherent and linked content, in the sense that the facts in the story are not just random and unlinked events.

5 Content Planner

In this section we present the *content planner* of the story generation system. This generator receives data from the Multiagent System described in Section 2.2. The Multiagent System outputs a XML file that stores the full log of every agent in the simulation, each log representing the set of facts of a character in the story, and the *content planner* imports that file.

Along with the facts that happened during the simulation, we need relations between them, like “at the same time”, or “because of that”. The relations between facts of the story are computed when the XML from the multiagent system is loaded. Knowing the domain rules (when some fact is the cause of another, or when two statements are consecutive facts), we can infer when two facts are related by a *causal relation*, a *temporal relation*, and so on. This is done by hard-coding the domain model in the XML loader, creating the relations during the import stage. Of course, this approach is far from being general, and it is possible to infer the relations between facts in a more generalistic manner, although it is not addressed in this work.

To handle dialogs in the narration, we consider that each element of communication between two agents, like “saying where” or “asking how” is indeed another fact, like “jumping” or “going somewhere”. Each one of these facts has particular relations with other facts, that can be dialog facts, or narrative facts. With this approach, dialogs and narrative parts can be mixed inside the text, thus getting a richer narrative output.

5.1 System Representation of the Reader’s Perception of the Story

Of course, nowadays the task of modelling the human brain is far from being possible. There are many scientific (and not only scientific) disciplines involved in such task, and to propose a model is not the objective of this work. However, we have made a lot of relaxations in the formalism of the model, and we have adjusted it to be very particular to the domain of this study. In this manner, it is possible to approximate to a kind of mental state model that does not try to emulate human’s.

We present an algorithm for *content planning* based on the reader’s perception of the story. In the system we present in this paper, the process of creating the discourse of a story takes into account what the potential reader of the text thinks and knows, using the model of the characters, about the world we are describing, and tries to give the reader the information he should receive at each stage of the story. The system chooses, in this way, which set of facts are going to appear in the final realization, and which relations connect them.

With this model of the characters, the *content planner* decides to choose some relations or others, depending on the BDI state of some characters. In this way, the state of the characters during the story guides the generation. We have created rules for computing the new state, and they are dependent on the previous state and the new sentence that the reader is going to read. These rules are based on the BDI structure and rules of the agents, as explained in Section 3.

5.2 State Space Search

Once we have defined a perception model and the guidelines that hold the information we need for writing the story, we have the basic data structures we need to work. The next step, then, is to create the discourse. We propose a simple approach based on state space search algorithms. What we have done is to define a initial state, with no information, and start, from it, a **backtracking** algorithm that explores the solution space, by creating different stories, using *relations between statements* as operators. Figure 1 depicts these ideas.

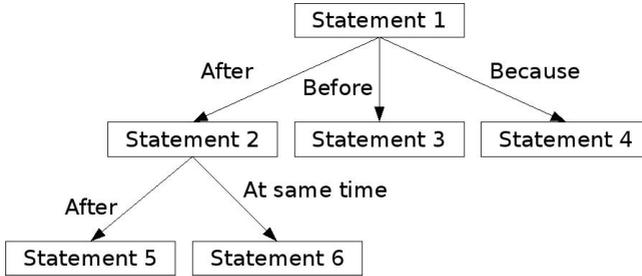


Fig. 1. Search tree for *content planning*

5.3 Generation Goals

When running the space state search, many possible stories are generated. However, we need a method to chose which of them is the best one. For this purpose we define *objectives*, that are data structures holding several attributes of the story draft. These attributes are the *linearity* of the text, or the level of sequentiality the story has; the *theatricality* of the story, counting the percentage of dialogue parts in the text; and the *causality*, or how important it is for the final story to contain “cause–effect” relations. This values are computed with rules and simple algorithms.

Then we can compute the similarity between a story generated during the search and an *objective* previously established by the user. The system will select that story with higher similarity to the *objective*. This similarity is computed by calculating the euclidean distance between the values of the attributes previously explained.

6 An Example

We can see an example of the *content planner* execution, from a set of ten agents during fifty years in Figure 2. It is a fragment of a generated story. It is important to note that *sentence planning* and *surface realization* are not addressed in a formal way. Instead, our proposal is to create simple template based text creation. From the discourse that has been generated, we apply a simple algorithm that fills the gaps in text templates.

It was a man. And His name was Deron. And His last name was Cairnbreaker. And Deron Cairnbreaker desired to become a great wizard. After that, the spell of memory was cast upon Deron Cairnbreaker. Because of that, its education decreased. After that, Deron Cairnbreaker and Parbagar Greatcutter talked:

- Do you know who has the one ring?

- Yes, I can tell you who has the one ring - said Deron Cairnbreaker, and it told where.

- Are you sure? Then I'll go and talk with it. - said Parbagar Greatcutter - Farewell.

Before that, Deron Cairnbreaker and Georgia Houston talked:

- Do you know where can I find another wizard?

- Yes, I do. I will tell you. - said Deron Cairnbreaker. Then, Deron Cairnbreaker showed the place.

- Ok, now I have this useful information. - said Georgia Houston - Thank you!

Fig. 2. Example of a story

However, these templates are not only filled with the statement that is going to be added: Not only the sentence is important for the template: also the reader's mental state influences the final text output. What the reader thinks about a character can change the way we produce the sentences. For example, if the mental state says that the reader knows nothing about the gender of a character, we can not use pronominal references that would require implicit mention of gender.

7 Comparison with Other Systems

The work in HERODOTUS [9] presents a *content planning* system which performs content determination assigning an heuristic value (*interest*) to facts, removing those facts whose *interest* falls below a given threshold, and discourse planning using templates. Compared to the system presented in this paper, the template system is too rigid, and it only allows to create a defined fixed set of types of stories. Although the rules it uses for *content planning* give some flexibility to the system, the resulting discourse is not good enough. Creating new stories is hard, and the system is not "very intelligent". There is much effort on *content determination*, and the main advantage was the ability of filtering a huge amount of information, and displaying it in a more human readable manner.

In MEXICA [2], emotional links are followed for planning the story and the present work follows a mental state model of the reader, although these concepts are not the same, there are some similar ideas between them. The work of Riedl and Young [1] shows a *content planning* system that creates stories from an initial state, and trying to reach some goals in the final story. None of these works addresses *sentence planning*.

The work presented adds dialog generation, mixed with the narrative content of the story, and simple final text generation. With dialogs we try to create richer and more complex stories, because dialogs show aspects of the story not present in plain narrations. Dialogs usually describe characters and action in such a way

that the resulting story, together with narrated parts, becomes more interesting for the reader.

8 Conclusions and Future Work

A multiagent social system with intelligent agents based on BDI model has been presented. This system simulates the life of several fantastic characters living and communicating between them. Also, a storytelling system able to generate text with narrative content and dialogs has been shown and explained. This two parts, taken as a whole system, try to create stories with focus on characters interacting.

The BDI model for the planning system is still not good enough. It follows very simple ideas, and thus the agent system and the story generator do not produce very good output. However, there are many possible solutions. An easy but useful change could be to add a percentage number for each Belief, Desire and Intention. For the Beliefs, the number would represent its certainty: how sure the agent is about what he knows. In this way we could introduce the possibility of being confused, doubting or even lying.

Another way of introducing complexity in the system is by improving the agents' dialogs, thus making them able to negotiate in a more complex way. Now agents have just basic conversations: if one agent asks, the other agent answers, if he knows, what the first one wants to know. But this could be highly more sophisticated. They may ask about their own interests, see what they have in common, and share their knowledge about those things.

The natural text generation system has some advantages with respect to other systems, as it has been shown in the comparison (Section 7), but it has also many disadvantages that must be taken into account to evolve the system. First, it is necessary, following this research line, to improve the model of the mental state. Although creating a mental model is a very ambitious task, it is possible to develop simpler prototypes, thus obtaining some aspects of the main characteristics of the understanding of a text by the reader.

A more effective reasoner is necessary. The algorithm presented is very inefficient, and it would be very interesting to add new reasoning capabilities, perhaps with the use of knowledge systems, like ontologies.

Another main point to be improved in the narrative system is the *sentence planner*. It has been addressed from a very simple and direct point of view, but better quality in the process can give much better output. Also, this *sentence planner* is fully domain dependent, and it only can produce text for the multiagent system logs. Concerning the work of STORYBOOK [3], which has dialogs, perhaps it would be interesting to connect the output of our system with that storytelling system.

Acknowledgements

This research is funded by the Spanish Ministry of Education and Science (projects TIN2006-14433-C02-01 and TIN2005-08501-C03-01).

References

1. Riedl, M., Young, R.M.: From linear story generation to branching story graphs. *IEEE Journal of Computer Graphics and Applications*, 23–31 (2006)
2. Pérez y Pérez, R., Sharples, M.: Mexica: a computer model of a cognitive account of creative writing. *Experim. and Theo. Artif. Intell.* 13(2), 119–139 (2001)
3. Callaway, C.B., Lester, J.C.: Narrative prose generation. *Artif. Intell.* 139(2), 213–252 (2002)
4. Bratman, M.E.: *Intentions, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, USA (1987)
5. Pavon, J., Arroyo, M., Hassan, S., Sansores, C.: Simulacion de sistemas sociales con agentes software. In: *CMPI-2006. Actas del Campus Multidisciplinar en Percepcion e Inteligencia*, vol. I, pp. 389–400 (2006)
6. Sanz, J.J.G.: Metodologías para el desarrollo de sistemas multi-agente. *Inteligencia Artificial, Revista Iberoamericana de Inteligencia Artificial* 18, 51–63 (2003)
7. Dignum, F., Morley, D., Sonenberg, L., Cavedon, L.: Towards socially sophisticated BDI agents. In: *ICMAS 2000. Proceedings of the Fourth International Conference on MultiAgent Systems*, Boston, USA, pp. 111–118 (2000)
8. Rao, A., Georgeff, M.: BDI agents: From theory to practice. In: *Proceedings of the International Conference on Multi-Agent Systems* (1995)
9. León, C., Hassan, S., Gervas, P.: From the event log of a social simulation to narrative discourse: Content planning in story generation. In: *AISB Annual Convention*, Newcastle University, Newcastle upon Tyne, UK, pp. 402–409 (April 2007)
10. Reiter, E., Dale, R.: *Building Natural Language Generation Systems*. Cambridge University Press, Cambridge (2000)
11. Young, R.M., Moore, J.D., Pollack, M.E.: Towards a principled representation of discourse plans. In: *Proceedings of the Sixteenth Conference of the Cognitive Science Society*, Atlanta, GA (1994)