

Let’s Pretend I Had a Sword

Late Commitment in Emergent Narrative

Ivo Swartjes, Edze Kruizinga, and Mariët Theune

Human Media Interaction group, University of Twente, Enschede, the Netherlands
[swartjes|kruizinga|theune]@cs.utwente.nl

Abstract. We describe how autonomous character agents that inhabit a story world can make (out-of-character) decisions about its details, filling in the story world as they go. We describe how we model these kind of late commitment decisions, and discuss how we use them to support action selection and to justify the adoption of character goals. Although a rigorous evaluation remains future work, we have implemented the approach presented here and have performed some exploratory testing.

1 Introduction

The Virtual Storyteller [1] is a story generation framework based on the concept of emergent narrative (i.e., narrative through unscripted characters [2]). The aim of the project is to investigate possibilities of and issues with emergent narrative, in order to get insight into the conditions necessary for stories to emerge. The characters are enacted by intelligent agents, the design of which is informed by improvisational theatre [3]. We work from the hypothesis that character agents can produce better stories if they are able to make decisions in-character (IC, i.e., what does my character want, feel and do?) as well as out-of-character (OOC, i.e., what would be good for the story development?). For their IC decisions, the character agents run cognitive processes like cognitive appraisal and action planning. Some first steps towards endowing virtual agents with OOC decision processes have been made in the field, e.g., the “double appraisal” approach of Louchart, where characters choose actions based on their IC motivations as well as on the OOC aim to emotionally impact other characters [4]. In this paper, we elaborate on another OOC decision process, which we have called *late commitment*¹[3]. This concept allows virtual agents to fill in the story world in line with what the story needs. Props are added, characters are deepened, relationships defined, and the backstories of the characters are filled in *during the simulation*, when this is useful for story progression, rather than determined beforehand.

¹ Late commitment should not be confused with the term *least commitment*, commonly used in the planning community. While both have a similar rationale (i.e., keeping options open by delaying choices as long as possible), least commitment refers to deferring decisions about variable bindings and step orderings in the search for a plan, whereas late commitment refers to deferring decisions about the story world setting in the course of the simulation.

2 Late Commitment

Emergent narrative has previously taken an authoring approach whereby fixed, predetermined story world settings and character personality definitions serve as the basis for narrative development. However, such a fixed story world setting already considerably constrains the possible course of events, taking away some of its generative potential. Indeed, this approach can even be problematic since it is unclear from an authoring perspective *how* exactly the story world setting constrains the course of events, making it difficult to decide what props, relationships and backstories to put in the initial setting.

The idea of late commitment is based on the observation that in improvisational theatre, actors do not work with a predetermined, agreed upon story world, but frame it as they go along, continually adding new information by the things they say and do, and accepting the information that others contribute. To illustrate, children do this in improvised play (sometimes even very explicit, when they say: ‘let’s pretend I was the mother, OK?’ [5]). In improvisational theatre this is never explicit but always implicitly conveyed through IC communication (‘hello, *daughter*’), because OOC communication would disrupt the performance and the audience’s suspension of disbelief.

When virtual characters are given means to modify the virtual story world, they can make similar framing decisions, but unlike real actors, they can have explicit OOC communication with each other ‘under the hood’, which remains invisible within the story world. We model these decisions using *framing operators*, STRIPS-like operators that embody OOC ‘let’s pretend...’ activities. The execution of a framing operator leads to the assertion of its effects, but must create the illusion that they have always been true within the story world.

A similar idea has been briefly explored from a plot-centric approach by Riedl and Young. Their Initial State Revision (ISR) planner [6] is a story planner with support for a partially indeterminate set of initial facts, aimed to increase the space of possible plans (and consequently, the space of possible stories).

3 The Use of Late Commitment in the Virtual Storyteller

We have implemented support for the selection and execution of framing operators in the Virtual Storyteller and are currently experimenting with the authoring of a story domain about pirates, making extensive use of late commitment. We currently use late commitment in two cognitive processes: (1) goal management and (2) action selection. Both processes make use of a partial order planner, modified to allow the use of framing operators. First, goal management has been extended to allow the Character Agents to justify adoption of new goals when they have no goals to pursue. Character agents can adopt goals when the preconditions of these goals are met. For instance, adopting a goal of plundering another ship might require that there is another ship in sight, and that the character adopting the goal is a pirate captain. If these preconditions do not apply, they can be given to the character’s planner as OOC goals, and can be achieved

by framing operators that introduce a ship in sight, and endow its character with the role of pirate captain. Second, action selection has been extended to enable agents to create plans for their goals. If the captain has adopted a goal to find out whether the approaching ship is friend or foe, he can make a plan involving looking through binoculars. If the story world does not contain binoculars yet, they can be framed to be in the captain’s cabin, affording a plan for the captain to go to his cabin to get them.

Preliminary exploratory testing revealed several issues with planning and execution of framing operators, which led to modifications of the agent architecture. When constructing a partial order plan using framing operators, care is taken that no IC actions are selected to satisfy preconditions of framing operators, which would yield a problem with observed character motivation (e.g., a character trying to become captain because it needs binoculars). Furthermore, if the planner selects an action whose preconditions contradict the effects of a framing operator, it always has to be ordered *after* the framing operator, to maintain the illusion that the effects of the framing operator were already true before any step in the plan.

There are also several consistency issues involved with the execution of a framing operator. For instance, if the illusion is to be maintained that certain effects have always been the case, all characters must unconditionally accept all of the framing operator’s effects, and the characters must be able to believably refrain from responding emotionally to these effects. To this end, a framing operator selected by one of the characters is first proposed to all other characters. If any character cannot consistently accept the framing operator, it rejects the proposal and the framing operator will not be executed.

Example. We will briefly illustrate the use of framing operators for goal management and action selection using a simple example domain about pirates. In this domain, there are two pirates: Anne Bonney and Billy Bones. One of the character goals defined in this domain is a goal to kill one’s enemy. The preconditions for adopting this goal specify that one must hate the character one wants to kill, but we haven’t defined any of the characters to hate the other. Instead, we have created a **HateYou** framing operator, that can be used to actively justify the goal. The framing operator defines that characters can just happen to hate each other². Anne can use it to pretend that she has always hated Billy, so that she can adopt the goal. Based on this goal, she plans to stab Billy, but for this she needs to be carrying a rapier. Again, we have not specified this in the initial setting, but added a framing operator called **CarryRapier** to the domain, usable by Anne to pretend she happens to be carrying a rapier. The preconditions contextualize the operator by specifying that (in our domain) it is believable that a pirate might happen to be carrying a rapier. Furthermore, they ensure that inconsistent situations are avoided, e.g., the situation in which the rapier Anne pretends to be carrying was already located somewhere else. Its effects state that the pirate is indeed carrying the rapier.

² Artistically, this is quite shallow, but it suits the example.

The example illustrates that we can abstain from thinking about whether Anne should have a rapier at the start of the story, and whether she should hate Billy, because the decision whether it makes sense to incorporate these facts is now made within the cognitive processes of the agents. The concepts of carrying a rapier and hating people must of course still be authored, but whether or not they end up in the simulation now depends on the decisions the characters make.

4 Conclusions and Future Work

The technique of late commitment contributes to agents that can introduce story material intelligently, i.e., when useful for the development of a story. We have introduced the notion of *framing operator*, referring to an OOC activity by one of the character agents to fill in previously unspecified knowledge. The benefit of late commitment is that it enables specifying the story world in terms of how it could be rather than how it is, and as such takes away the responsibility of the author to predict which exact properties the story world needs to have for a particular course of events. Furthermore, it offers the characters more flexibility in their reasoning processes; the characters can to a certain extent *choose* to adopt goals and enable actions, by filling in the world around them.

We have implemented a model of the concept of late commitment and are using it to support action selection and goal adoption. We have addressed some consistency issues that were revealed by exploratory testing, but a more rigorous evaluation remains future work. Furthermore, we believe the use of late commitment is not limited to these two character processes. The next step is to investigate how other processes can benefit from its use. For instance, framing operators might be employed to cause desired perceptions (e.g., seeing a hidden door can be enabled by introducing a hidden door in sight of a character), make story world events possible, or facilitate desired emotional reactions.

References

1. Swartjes, I., Theune, M.: The virtual storyteller: Story generation by simulation. In: Proceedings of the 20th Belgian-Netherlands Conference on Artificial Intelligence (BNAIC). (2008)
2. Aylett, R., Louchart, S., Dias, J., Paiva, A., Vala, M., Woods, S., Hall, L.: Unscripted narrative for affectively driven characters. *IEEE Computer Graphics and Applications* **26**(4) (2006) 42–52
3. Swartjes, I., Vromen, J.: Emergent story generation: Lessons from improvisational theater. In: *Intelligent Narrative Technologies: Papers from the AAAI Fall Symposium*. Number FS-07-05 in AAAI Fall Symposium Series (2007)
4. Louchart, S., Aylett, R.: From synthetic characters to virtual actors. In: *Proceedings Third Artificial Intelligence and Interactive Digital Entertainment (AIIDE-07)*. (2007)
5. Sawyer, R.K.: Improvisation and narrative. *Narrative Inquiry* **12**(2) (2002) 319–349
6. Riedl, M.O., Young, R.M.: Open-world planning for story generation. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. (2005) 1719–1720