

Emotional Characters for Automatic Plot Creation

Mariët Theune, Sander Rensen, Rieks op den Akker, Dirk Heylen, and Anton Nijholt

University of Twente, PO Box 217, 7500 AE Enschede, The Netherlands
{theune|rensen|infrieks|heylen|anijholt}@cs.utwente.nl

Abstract. The Virtual Storyteller is a multi-agent framework for automatic story generation. In this paper we describe how plots emerge from the actions of semi-autonomous character agents, focusing on the influence of the characters' emotions on plot development.

1 Introduction

The Virtual Storyteller is a multi-agent framework for automatic story generation. Its functionality includes plot creation, narration and presentation. Plots are automatically created by semi-autonomous character agents that carry out different actions to achieve their goals. To guarantee that the resulting plots will be well-structured, the characters are guided by a director agent, who provides (some of) them with specific goals, and who can forbid their actions if they do not fit into the general plot structure. The resulting plot is converted into a natural language text by a narrator agent, and presented to the user by an embodied, speaking presentation agent. The global architecture of the Virtual Storyteller is shown in Fig. 1.




Fig. 1. Architecture of the Virtual Storyteller

With the first version of the Virtual Storyteller [12] we aimed at meeting two basic plot requirements: consistency at the level of character actions (through the use of goal-oriented characters) and well-structuredness (through the use of

a director agent). Another storytelling system that combines a character-based approach with the use of a virtual director is Teatrix, which is aimed at aiding story creation by children [3]. In the Oz project [4], a ‘drama manager’ was used to ensure the occurrence of some essential plot points during interactive story creation. The ACONF system [8] uses a different approach: here, a coherent plot structure is achieved by representing the plot as a partial-order plan, which is constructed by a group of expert systems representing the story characters.

Consistency and well-structuredness are essential for achieving a well-formed plot, but do not guarantee that the resulting story will be entertaining. For that, a plot should have additional properties such as unexpectedness, suspense and the presence of believable characters. A common approach to make character agents believable is to provide them with emotions. Most computational models of personality and emotion are based on Ortony, Clore and Collins’ event appraisal theory, known as the OCC-model [6]. Examples are Elliott’s Affective Reasoner [1] and Neal Reilly’s Em system [5]. A simplified version of Em was incorporated in Loyall’s Hap architecture for believable agents [2], used in the Oz project. This work was the main source of inspiration for our recent extension of the Virtual Storyteller with an emotion model for the characters.

In the following sections, first we briefly sketch the emotion model used in the current version of the Virtual Storyteller. Then we discuss the influence of the character’s emotions on plot creation, showing how the characters’ ability to adopt goals that are based on their emotions, rather than dictated by the general storyline, increases their believability and may give rise to unexpected (but nevertheless consistent) plot twists. We illustrate this with an example.

2 Emotions and Actions

In the Virtual Storyteller, the emotional state of an agent is represented by pairs of corresponding positive and negative emotions and their intensities. Emotions that are directed at the agent itself are hope-fear, joy-distress, and pride-shame. Emotions directed at other agents include admiration-reproach, hope-fear, and love-hate. Each emotion pair has an intensity value, represented by a natural number on a scale from -100 to 100. For example, a value of 20 for hope-fear means that the character experiences hope with an intensity of 20 (and no fear); a value of -20 means exactly the opposite.

Following the OCC-model [6], the emotional state of an agent changes in reaction to events, actions and objects. If an agent comes close to achieving its goal, it experiences hope, and if its current goal proves to be unachievable, the agent experiences distress. If an agent performs an action meeting its own standards (which are part of its personality model), the agent experiences pride. If an agent meets another agent which it likes (also part of the personality model), it experiences joy. The intensity and duration of these emotional effects depend on the personality parameters of the agent, which can be changed by the user. For example, some characters may be frightened more easily than others, or may forget their anger more quickly.

Based on its emotional state, a character may develop certain action tendencies such as friendly or unfriendly, aggressive or passive behaviour (cf. the mapping of emotions to action styles in Hap [2]). The action tendencies in their turn influence the importance a character attaches to certain goals. The relation between the emotions and the action tendencies of a specific character is laid down in a number of action tendency scripts, which are part of the character's personality model. For instance, strong fear could lead to passive behaviour in one character (increasing the importance of the goal to flee) and to aggressive behaviour in another (increasing the importance of attacking the object of fear).

3 Plot Creation Using Episodes

The Virtual Storyteller is aimed at the generation of short stories in the fairy tale domain. An advantage of this domain is that most fairytales have a common, fairly simple plot structure. The stories generated by the Virtual Storyteller consist of four global episodes, which are loosely based on Greimas' reformulation of Propp's story functions [7]:

1. initial state of equilibrium (a description of the initial setting of the story)
2. disruption of the state of equilibrium (something bad happens)
3. mission of the hero (the hero tries to restore the state of equilibrium)
4. return to state of equilibrium (or not, if the hero fails)

At the start of the story generation process, the director lays the basis for each of the episodes by selecting four 'episodic scripts' from a database. An episodic script contains information about the setting, goals and constraints of an episode. The setting specifies which locations, objects and characters are present in the episode. The episodic goals are goals which must be reached within the episode by (some of) the characters. Typically, episodic scripts of type 2 (disruption) will involve a villainous character with the episodic goal of performing some evil deed, and episodic scripts of type 3 (mission of the hero) will involve a character with the episodic goal of putting things right again. Finally, the episodic constraints specify certain actions that are not allowed in the episode. For example, it might be disallowed to kill a certain character before its episodic goal is reached. These three episode properties define the limits within which the character agents are allowed to act while constructing the episode.

After the episodic scripts have been selected, a story world is created according to the specified setting, and the character agents are initialized. The actual episode is then created by the characters carrying out actions to reach their individual goals. These goals can be the episodic goals that have been assigned to them by the director, but they can also be goals that were spontaneously adopted, based on the character's emotions (see section 4). Before a character can carry out a planned action, it must ask the director for permission. The director checks if the action violates any episodic constraints, or any global constraints (for instance, the director might forbid numerous repetitions of an action because they would make the story boring). If no constraints are violated, the

character is allowed to carry on with its action, and the story world is updated accordingly. When one of the characters reaches its episodic goal, the episode ends. Because the characters can also adopt goals other than the episodic goal, there are many different ways in which an episodic script can develop before an episodic goal is reached. This variation can be further increased by changing the personality parameters of the characters (see section 2). Still, all variations will respect the global plot structure, because of the fixed order of episodes in combination with the local episodic goals and constraints.

In our current set-up, the episodic scripts are created and combined by a human author.¹ In terms of the ‘scalable autonomy’ approach of Spierling *et al.* [9], this means that there is no autonomy at the story level: the global plot structure is predefined. However, there is high autonomy at the scene (episode) and character levels: the actions carried out within an episode are not specified in advance, and the characters can make their own decisions (guided by the director, which is itself an autonomous agent).

4 Goal Selection: Emotions Taking Over

In the Virtual Storyteller, each character agent has a fixed set of potential goals. The importance that a character attaches to these goals varies depending on its action tendencies (see section 2). In addition, one of its goals may have the status of ‘episodic goal’, in which case it has a very high importance (90 out of 100) throughout the episode. Instead of having a character simply adopt the goal with the highest importance, we use a probabilistic approach to goal selection, where a goal’s importance corresponds to the probability that it will be adopted. In other words, an episodic goal has a 90% chance of being selected for pursuit, but some ‘emotional’ goals may also have a high probability, due to the character’s emotion-induced action tendencies. This means that a cowardly hero with the episodic goal of killing the villain could be equally likely to adopt the goal of running away from the villain because of its fear.

As we will illustrate with some examples below, probabilistic goal selection leads to variation and unexpectedness in the generated stories. However, there is risk of ‘overdoing’ it by having too much random switching between goals. Therefore, if a character has already carried out the first step of a plan towards achieving some goal, that goal will be maintained until it is reached (or made unreachable by changes in the story world). The goal is only abandoned if the character’s emotions run so high that they cause the importance of some other goal (e.g., fleeing from an enemy) to exceed a pre-specified high value. Another risk of our approach is that a character’s sudden change in behaviour, due to the adoption of a new goal, may seem incomprehensible to the user [10]. In the

¹ As our script database grows, it will become more important to have a proper, automatic selection mechanism that only combines compatible scripts. Also, in the future we would like to distinguish more than four types of episodes, by including more of Propp’s functions. The sophisticated, real-time scene selection mechanism of the GEIST system [9] could serve as an example here.

Virtual Storyteller, we can solve this by having the narrator explicitly describe the character's motivation. Currently, this is only done to a very limited extent, and further work on this is required.

5 An Example

Here we illustrate the plot creation process by describing the different realisations of a simple script of type 2(disruption). The script involves two characters, a princess (Diana) and a villain (Brutus), who hate each other. Diana fears Brutus, but Brutus does not fear her. Diana has the episodic goal to kill Brutus, and Brutus has the episodic goal to capture Diana. Initially, the characters are in different locations. In yet another location, a sword can be found. We first discuss what happens given the default, rather stereotypical action tendency scripts of the princess and the villain. These scripts specify (among other things) that Diana is likely to behave passively when experiencing emotions such as fear and hate, whereas these same emotions trigger aggressive behaviour in Brutus.

In most episodes that were generated on the basis of this script, the characters immediately start pursuing their episodic goals: Brutus starts walking towards Diana to capture her, and she sets out to get the sword, needed to kill Brutus. Occasionally, one of the characters adopts another, emotionally motivated goal instead, such as singing out of joy. This occurs rarely, however, because initially the character's emotions have low intensity, so the importance of the associated goals is also relatively low. Whatever goal they start out with, the two characters eventually meet each other. When meeting Brutus, Diana is usually overcome by fear, causing her to abandon her episodic goal (even if she has already found the sword) and to run away instead. Brutus, who will act aggressively even if Diana has the sword, will follow her, either to capture or to attack her (the latter goal being motivated by his aggression). In the end, Diana is inevitably cornered by Brutus without any possibility to flee. If she has the sword, she is very likely to use it now, since killing Brutus is now her highest priority goal that can be attained. Otherwise, she is most likely to scream, because this is the only passive behaviour left open to her. If she does not have the sword, Diana may also go for the less likely goal of attacking Brutus by hitting or kicking him, but Brutus is stronger and will eventually capture her anyway.

We have also tested what happens, given the same episode script, if the action tendency scripts of the two characters are exchanged. In that case, the division of power is more equal, because although Brutus may be stronger than Diana, he is now more fearful and passive than she is. This leads to longer episodes involving more chase sequences and mutual attacks. In the end, as in the first version, the outcome depends on whether Diana manages to get to the sword in time or not. All in all, a large variety of episodes can be generated based on the same script. The probabilistic goal selection may give rise to unexpected character actions, but these actions can always be traced back to the character's emotions. Properly expressing these inner motivations in the narration of the plot will be our next challenge.

6 Conclusions and Future Work

The version of the Virtual Storyteller we have described here has been fully implemented, albeit with a limited set of rules and episodes. However, our experiences so far indicate that even with this limited implementation, a large number of different stories can be created. Further extension of the system with more possible goals, actions and objects is needed to show the full potential of our approach. As pointed out above, another important area of further work is the conversion of the plot into natural language, so that the underlying motivations of the characters are properly put across. With respect to presentation, we are currently working on the automatic generation of ‘storytelling prosody’ and on the conversion of abstract plots to 2D animations, as an alternative presentation mode for the Virtual Storyteller. Finally, in the longer run we want to extend the Virtual Storyteller so that the emotions of the user can be taken into account too, by learning rules to predict the effect on the user of different character actions. These rules can then be used to guide the decisions of the director (cf. the narrative sequencer employed in IDtension [11]).

References

1. C. Elliott. *The Affective Reasoner: A Process Model of Emotions in a Multi-Agent System.*, Ph.D. thesis, Northwestern University, 1992.
2. B. Loyall. *Believable Agents: Building Interactive Personalities.* Ph.D. thesis CMU-CS-97-123, Carnegie Mellon University, 1997.
3. I. Machado, R. Prada, and A. Paiva. Bringing drama into a virtual stage. *Proceedings of the Third International Conference on Collaborative Virtual Environments (CVE 2000)*, ACM Press, 2000.
4. M. Mateas. *An Oz-Centric Review of Interactive Drama and Believable Agents.* Technical Report CMU-CS-97-156, Carnegie Mellon University, 1997.
5. S. Neal Reilly. *Believable Social and Emotional Agents*, Ph.D. thesis CMU-CS-96-138, Carnegie Mellon University, 1996.
6. A. Ortony, G.L. Clore and A. Collins. *The Cognitive Structure of Emotions.* Cambridge University Press, 1988.
7. V. Propp. *Morphology of the Folktale.* University of Texas Press, 1968.
8. M. Riedl and R.M. Young. Character-focused narrative generation for execution in virtual worlds. *Proceedings of the International Conference on Virtual Storytelling*, 47-56, 2003.
9. U. Spierling, D. Grasbon, N. Braun and I. Iurgel. Setting the scene: Playing digital director in interactive storytelling and creation. *Computers & Graphics* 26, 31-44, 2002.
10. P. Sengers. Designing comprehensible agents. *Proceedings of the 16th International Joint Conference of Artificial Intelligence (IJCAI)*, 1227-1232, 1999.
11. N. Szilas. IDTension: A narrative engine for interactive drama. *Proceedings of the Technologies for Interactive Digital Storytelling and Entertainment (TIDSE) Conference*, 187-203, 2003.
12. M. Theune, S. Faas, A. Nijholt and D. Heylen. The Virtual Storyteller: Story creation by intelligent agents. *Proceedings of the Technologies for Interactive Digital Storytelling and Entertainment (TIDSE) Conference*, 204-215, 2003.