Virtual Story Generation:
from TALE-SPIN to the Virtual Storyteller

Laura Langohr
Seminar on Computational Creativity
Department of Computer Science
University of Helsinki
September 22, 2011

Abstract. In this seminar report we describe two story generation systems: TALE-SPIN, one of the earliest approaches, and the Virtual Storyteller, one of the latest approaches. Though the former produces compelling stories, we show why the latter produces stories which are more close to human produced stories. We will discuss similarities and differences of the two systems. In particular, we will show that the main idea (character- and goal-orientation) is the same in both, and the main difference is that the Virtual Storyteller distinguishes between the fabula and story text generation, while TALE-SPIN does not. Finally, we present some open questions and ideas on how to improve virtual story generation.

1 Introduction

In this report, we describe how TALE-SPIN [14] (from 1977) and the Virtual Storyteller [22] (developed since 2002) create and present stories, and give an overview of the computational components involved. We discuss similarities and differences of the systems, and mention open questions and ideas on how to improve virtual story generation. Next, we will introduce a number of concepts that define the scope and focus of the two systems.

Storytelling occurs in play, print, and film for centuries. Here, we consider a story (or drama, or narrative) to be a work of any length, created by an author for some audience. Although inventing and telling a story may go hand in hand, we separate the concepts. Hence, we refer to (virtual) storytelling as a computer telling a story based on an outlined plot. In *interactive storytelling* the user interacts with characters and objects of the story which adapts to the user's interactions and therefore is different each time [2].

A problem of interactive storytelling is that many variations are needed for a fundamental effect on the course of events [20], but writing out all of them is intractable [25]. This can be addressed by *(emergent) story generation* where a computer invents the story. Both TALE-SPIN and the Virtual Storyteller are story generation systems which allow a flexible shaping of the plot and character interactions in place of a predetermined one

authored by a human [20]. Such story generation deepens the understanding of what is needed for a good story, increases the feasibilities for diverse user interactions, and allows virtual authorship. Application areas of story generation are, amongst others, commercial games like first-person shooters and education [2]. For example, for educational purposes such as scenario-based training, educational goals are translated into a narrative environment that provides the interactor the chance to practice problematic situations and explore the consequences of possible choices to make [22]. For instance, FearNot! [3] is an anti-bullying educational story generation system which allows children to explore what happens in bullying in an unthreatening environment.

Story generation has to automatically ascertain whether a story is dramatically interesting and whether it conveys a coherent plot structure. A story is interesting, if it captures and maintains the involvement of the audience [6]. Stories with a plot following a dramatically interesting structure (for example, the dramatic pyramid [7]) are of interest to a wide audience [2]. On the other hand, improvisational theater is dramatically interesting without predetermined plot and principles of it can be used in story generation and interactive storytelling [23]. The Virtual Storyteller contains several features derived from improvisational theater. For instance, the environment is not predetermined in the beginning, but is generated while the story emerges.

Further, sub-stories, within the overall story, are essential as they ensure the continued attention of the audience [6]. TALE-SPIN as well as the Virtual Storyteller may produce simple substories. In both systems, the characters are goal-oriented, that is, they try to achieve their goal, which might result in solving first a subgoal and hence a substory about achieving the subgoal.

Coherence involves causality relations (causal connectedness of story events) and the choice of view point [17]. View point choices are important as simultaneous events in different locations are usually told consecutively. In TALE-SPIN, the view point is predefined (by the user), and causality relations are quite simple (every action is written on the standard output). The Virtual Storyteller separates the fabula (the story construction) generation from the storytelling [22]. A fabula is a formal representation of a story, in which the clauses of a story are classified into a small set of elements which are connected by their causal relationships. The fabula of the Virtual Storyteller is independent of viewpoints and therefore the view point (and changes of view point) can be selected after generating the story. Furthermore, if the story is not told right away, more complex causality relations are possible. For instance, (causal) reasons for a specific action, can be added later on.

Next, we present how TALE-SPIN and the Virtual Storyteller generate and present emergent stories in more detail. We start by describing different approaches of storytelling and story generation in Section 2. We then discuss how TALE-SPIN and the Virtual Storyteller generate stories (Sections 3 and 4). In Section 5 we conclude with some

personal notes about possible future work. The key sources for this report were [16, 15] for TALE-SPIN and [22, 26, 23, 21] for the Virtual Storyteller.

2 Related Work

The first attempt to automate story generation was Klein's Novel Writer in 1973 which generated murder stories taking place at a weekend party [10]. Meehan's TALE-SPIN from 1977 and its simplified version Micro TALE-SPIN created animal tales of bear, bee, and other animal characters [14, 16]. Nowadays there are various systems for interactive storytelling (e.g., [13, 4]) as well as for story generation (e.g., [17, 3, 8]). Even though today's story generation systems have several enhancements compared to the early approaches, most of them (including TALE-SPIN and the Virtual Storyteller) share their main approach: character- and goal-directed behavior. That is, character's goals and their outcomes anchor episodic structures.

To achieve a dramatically interesting plot the story should be consistent, well-structured, unexpected, suspenseful, and present believable characters [26]. A believable character is one who seems lifelike, whose actions make sense, and allows empathy by the audience [13]. Requirements for believability are, amongst others, emotions, self-motivation, the pursuit of multiple, simultaneous goals and actions, reacting and responding to stimuli of the environment, and broad capabilities to listen, talk, and act [12].

Therefore, many story generation systems share similar characteristics. For instance, emotion models for characters are used in Oz project's Façade [13] as well as in the Virtual Storyteller. Façade is an interactive story generation system, in which the user plays a friend of a couple and gets entangled in a conflict between them. Oz project's believable agents were the main inspiration source of the Virtual Storyteller's emotion model for characters [26]. Characters which imitate improvisational theater actors are used in the Virtual Theater Project [8] and the Virtual Storyteller. The Virtual Theater Project creates character agents for storytelling or story generation systems.

The Virtual Storyteller and Fabulist [17] differ from most approaches as they make a narratological distinction between the fabula (a formal representation of a story) on the one hand, and the narrative text on the other hand. The Drammar ontology represents story structures as a high-level ontology [5], which the Virtual Storyteller's fabula structure elaborates [21]. Fabulist generates stories in different, prespecified domains, for example, of Jasemine and king Mamoud [17]. These systems have several advantages due to the differentiation between the fabula and the narrative text. The consistency and coherence of the fabula can be ensured when it is generated. The fabula is independent of stylistic concerns (such as viewpoints or time lapses) and therefore can be shaped into a story of any presentation form (narrative text, animation or picture story). Though, story generation systems often lack a formalization of such structures [5]. Next, we discuss

TALE-SPIN which does not distinguish between these layers. Afterwards, we describe the Virtual Storyteller and see the advantages of the distinction for story generation.

3 TALE-SPIN

In this section we describe how Micro TALE-SPIN generates stories using the story of Figure 1 as an example. We next discuss TALE-SPIN and Micro TALE-SPIN in detail, and explain how this story was generated.

TALE-SPIN's and Micro TALE-SPIN's story generation begins with a setting defined by the user. It consists of *characters*, *objects*, (their) *locations* and their *relationships*, which are specified by the user [16]. The original narrator of (Micro) TALE-SPIN generates stories in the animal domain, where characters are, for example, bears and bees. For this seminar report, we made small modifications to generate stories in the domain of little red cap (also known as little red riding hood). Note, that the goal is not to reproduce the fairy tale by Brothers Grimm. Instead, the aim to use this domain for the seminar report is to understand story generation by producing new, compelling stories using the characters and goals as inspiration.

For instance, in the example story of Figure 1 characters (or personaes in the theatralic sense) are little red cap and big bad (wolf). Both of them are in the forest, the

Once upon a time ... Little red cap was near the forest. Big bad was near the forest. The grandmother was near the house. The berries were near the meadow. [...]

One day, Big bad was hungry. Big bad wanted not to be hungry. Big bad wanted to have the grandmother. Big bad wanted to know where the grandmother was. Big bad wanted little red cap to tell big bad where the grandmother was. Big bad decided that if big bad would give little red cap the berries then little red cap might tell big bad where the grandmother was. Big bad wanted little red cap to think that little red cap would tell big bad where the grandmother was if big bad gave little red cap the berries. Big bad wanted to be near little red cap.

Big bad asked little red cap whether little red cap would tell big bad where the grandmother was if big bad gave little red cap the berries. Little red cap told big bad that if big bad would give little red cap the berries then little red cap would tell big bad where the grandmother was. Little red cap decided that if big bad would give little red cap the berries then little red cap would tell big bad that big bad was not bright.

Big bad wanted to have the berries. [...] Big bad went to the meadow. [...] Big bad took the berries. [...] Big bad wanted to be near little red cap. Big bad went to the forest. [...] Big bad gave little red cap the berries. [...] Little red cap had the berries. Little red cap told big bad that big bad was not bright. [...] Big bad thought that big bad did not have the grandmother. Big bad wanted to have little red cap. [...] Big bad took little red cap. [...] Big bad ate little red cap. Big bad was not hungry. The end.

Figure 1: Example story by Micro TALE-SPIN. We produced the story for this report by adapting Micro TALE-SPIN to the little red cap domain. (TALE-SPIN is implemented in Common Lisp: http://www.eliterature.org/images/microtalespin.txt)

grandmother (here, an object) is in the house, and berries are in the meadow. Characters may know only parts of what the narrator knows: In the example, Little red cap knows that the grandmother is in the house, but big bad does not know that. Big bad was specified to be hungry, i.e., the goal not to be hungry anymore.

Relationships (between characters) specify who likes, deceives and/or dominates whom. In TALE-SPIN relationships and personality traits (e.g. kindness) scale from -10 to $+10$ [16]. For instance, positive values of affection represent that someone likes somebody else. Micro TALE-SPIN uses truth-values (yes or no) instead of scales [15]. For instance, for the example story of Figure 1 the following relationships were specified: Big bad likes little red cap, but little red cap dislikes and deceives big bad.

The story consists of story fragments which are specified by *goals*, *subgoals*, and *actions* [16]. Goals model intentions such as hunger, thirst or loneliness, and will induce subgoals and actions (see Figure 3). In the beginning, the user assigns one goal to one of the characters. This (main) character then tries to solve the goal and hence, the story is generated. If a goal, subgoal, or action occurs, it is recorded, output, and all its consequences (i.e., other subgoals and/or actions) are computed, and so forth. For each (sub)goals and actions a problem solving procedure specifies which subgoals and actions follows from it. For example, to satisfy hunger (a goal), a character will try to get some food and eat it [16]. If he knows where some food is, she will eat it (an action). If he does not know he will try to find out (a subgoal), for example, by asking another character like in the story of Figure 1.

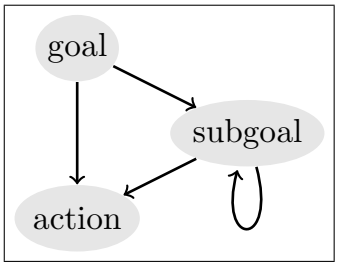

Figure 2: Story fragments of Micro TALE-SPIN. We made this graph for this report based on descriptions in [16].

Subgoals and actions include acquisition of objects, acquisition of information, transportation (of objects or characters including oneself), persuasion, asking, bargaining, threatening, and the transfer of information. Some of these might be dependent on the specified relationships between characters. For instance, requests where a character C has asked another character A to do something are handled as follows. If A dislikes, distrusts, or dominates C, then A will refuse [16]. If A likes C then A will agree. Consider again, the example in Figure 1. Little red cap was set to dislike and deceive big bad. Therefore, she does not tell big bad where the grandmother is and gets eaten herself. If we would have specified that she likes him, she would have told big bad where grandmother is. Hence, big bad would have gone to and eaten the grandmother and not little red cap.

More elaborate language processing would enhance the quality of the generated stories, but were not central to Meehan [16]. Further, the selection of the view point is predefined. The story is told from the view point of the main character, unless the character persuades another character to do something, which will lead the second character to achieve that goal and narrator to change the point of view. In addition, we cannot judge (as we only

run Micro TALE-SPIN by our own), but guess that TALE-SPIN produce more complex stories than Micro TALE-SPIN due to more complex relationship model.

4 The Virtual Storyteller

In this section we describe how the Virtual Storyteller generates the fabula and how the storytelling component then turns it into a story of narrative form. See Figure 3 for an example story generated by the Virtual Storyteller.

4.1 Story Generation

The Virtual Storyteller uses a multi-agent system to generate fabulas [26]. A multi-agent system consists of multiple computing components, called agents [28]. Agents are autonomous of and interact with each other: They make autonomous decisions to reach their goals, given the information about their environment (the story world) and cooperate, coordinate, and negotiate with each other. In the Virtual Storyteller a plot agent, a world agent, and several character agents generate a fabula [22]. Next, we describe the Virtual Storyteller's fabula, its agents and how they interact, and how those interactions produce a fabula.

Little red cap skips to the clearing in the forest. Grandmother bakes the apple pie. Big bad is mean. "Hello," little red cap says to big bad. Grandmother poisons the pie with some cyanide. "Give me that!" says big bad and takes the birthday cake from little red cap. Little red cap starts to cry.

Once upon a time there was a pirate, who was called Billy Bones. He was in the hold of his ship. The water supply was empty and he wanted to fill it. Therefore he opened the hatch. With a ladder the pirate walked to the deck. With the ship he sailed to an island. After he had gone ashore at the island, he filled the water supply with water from a pond.

Figure 3: Two example stories by the Virtual Storyteller. The first one is of the little red cap domain and was produced for a table-top environment in which each sentence is output immediately (source: from M. Theune by e-mail). The second story is from a pirate's domain (source: [22]). Here, the language of the story was improved according to methods described in Section 4.2. For this seminar report we did produce fabulas in the little red cap domain with the Virtual Storyteller, but unfortunately did not succeed to transform them into a narrative text. (The Virtual Storyteller is implemented in Java on top of JADE (Java Agent Development Framework), using SWI-Prolog (Dutch: Sociaal-Wetenschappelijke Informatica Prolog, English: Social Science Informatics Prolog) for its planning and reasoning processes [22]: http://wwwhome.ewi.utwente.nl/∼theune/VS/)

4.1.1 The Fabula

A fabula is a formal representation of a story and models its causal relationships [27]. The Virtual Storyteller captures the fabula (see Figure 4) in a single network, which is independent of viewpoints and therefore offers a global perspective on the fabula [21].

The fabula model of the Virtual Storyteller defines seven types of elements: (1) *setting elements* are aspects of the story world that all characters share and consider to be true (e.g., location of grandmother's house), (2) *goals* are the main drive for a character to act, i.e., a desire to attain, maintain, leave or avoid certain states, activities or objects (e.g., a desire to bring a cake to grandmother), (3) an *outcome* is a concept that relates a character's goal to its fulfillment (positive outcome) or failure (negative outcome), (4) *actions* are changes of the state of the story world, which are performed intentionally by a character (e.g., walking to grandmother's house), (5) *events* are changes of the state of the story world, which are not performed intentionally by a character (e.g., a tree that falls down), (6) *perceptions* describe properties of the story world that a character witnesses, and (7) an *internal element* is everything else that goes on within a character, such as cognitions, emotions, and beliefs [27, 26, 22]. These are only the top elements of a more extensive subsumption hierarchy in which, for instance, goal subsumes attain-goal and avoid-goal and perception subsumes see and hear [21]. Some of these elements have preconditions specifying when they are possible, and effects that specify how the story world changes as a result [22]. For example, goals have preconditions that determine in which cases an agent can adopt the goal, and success- and failure-conditions that determine when the goal is fulfilled or failed.

The elements are connected by four types of causalities: (1) *physical* and (2) *psychological causalities* represent unintentional causes of either physical or mental nature (e.g., the goal to bring a cake to grandmother is psychologically enabled by a perception that the cake is baked, which is physically enabled by the event of baking the cake), (3) a *motivation* represents an intentional cause (e.g., the action of baking a cake is motivated by the goal to bake a cake), and (4) an *enablement* represents any cause where the causing element is making it possible for the resulting element to occur (e.g., the setting element and belief about the existence of grandmother's house and grandmother being in the house enable the goal to bake a cake for grandmother) [22]. The causalities differ in their strength, which is reflected in their presentation: The stronger a causal relationships is, the less explicit mentioning it need [21]. Physical causality is the strongest and might not need to be made explicit in the presentation of the story, followed by motivation and psychological causation. Enablement is the weakest causality.

Episodic structures are commonly found in stories [27]. A goal might lead to the adoption of a new goal [22]. Adopting a new goal may generate another episode as to achieve the new goal might produce an substory before returning to the original goal.

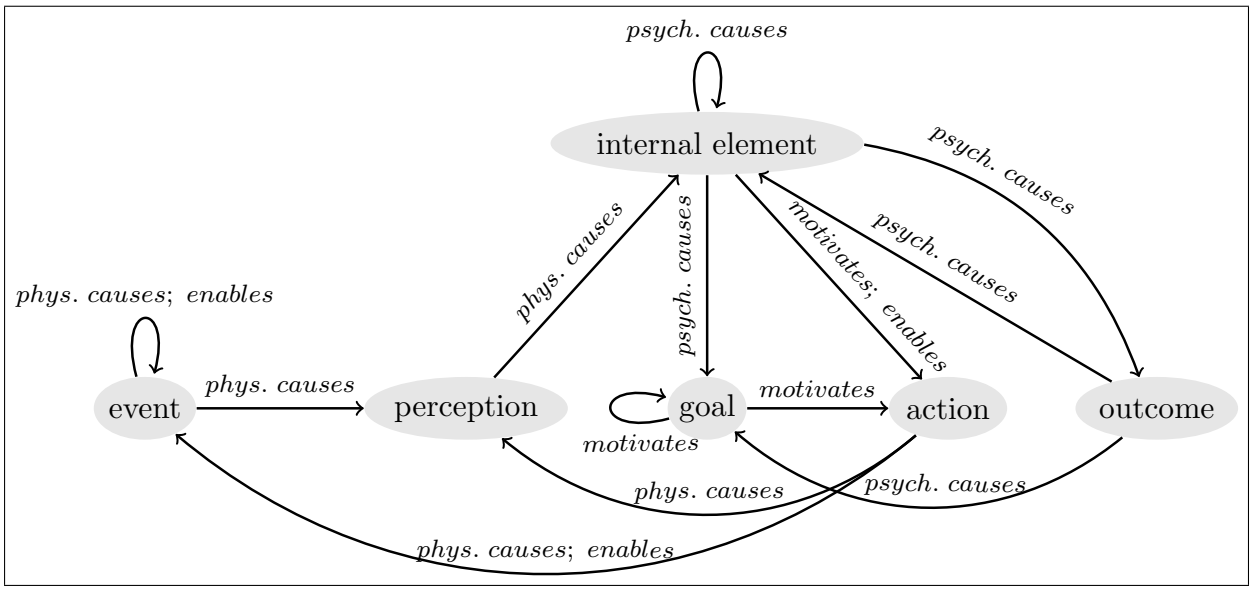

Figure 4: The Virtual Storyteller's fabula model on its top level. (source: [22]).

When an character agent adopts a new goal it checks whether the precondition of this goal is fulfilled by fabula elements that are causally related to earlier goals. In such a case there is a causal connection between the new goal and the enabling or causing of earlier fabula elements. These connections are then recorded in the fabula, and will have an effect on the narrative text generated later (where causal connections are replaced with corresponding rhetorical relations) [20].

A network generated with this model consist of all fabula elements of the story that have either a cause or an effect on other elements [21]. Goals motivate actions or other goals (e.g., the goal of the wolf to eat the grandmother motivates the subgoal of finding her). Actions physically cause perceptions and events (e.g. when a hunter shoots a wolf, this can physically cause the wolf to die, or a hunter crosses a small bridge might cause the bridge to break). Events physically cause perceptions and other events (e.g. a tree that falls on a bridge causes the bridge to collapse). Perceptions psychologically cause internal elements (beliefs). Internal elements (beliefs, emotions, etc.) can enable an action if that belief satisfy the preconditions of the action, psychologically cause other internal elements (emotions), and motivate actions (e.g., an emotion might motivate reflex-like actions like crying or screaming), amongst others. Outcomes can in turn psychologically cause goals or internal elements (e.g., positive/negative outcomes lead to positive/negative emotions and in case of a negative outcome possibly the reinstatement of the failed goal).

The model does not include actions enabling other actions. Actions are planned by the characters (unless they are directly caused by emotions) and hence can be enabled only by beliefs [21]. For example, killing a deer does not enable eating the deer until the hunter sees and believes that the deer is indeed dead. Setting elements are considered as background information and therefore left out of the model.

4.1.2 Plot and World Agents

The plot agent specifies the initial states of the story world, new character agents if necessary, and casts them to play the required roles [22]. The fabula is then generated in rounds, synchronized by the plot agent. In each round the plot agent requests each character agent to select an action. Each character agent goes through a deliberation cycle and either responds with an action or refusal. An action request is refused if the character agent is already performing an action, or cannot come up with any action. The plot agent can forbid character agents actions if they do not fit into the general plot structure [26]. When an action finishes, all character agents receive its effects in form of perceptions, and the performing character agent is informed whether an action was successful or failed [22]. The plot agent captures the fabula elements that the character agents produce and organizes them in a causal network for the subsequent storytelling.

The world agent manages the virtual world representation, executes actions and events as they occur, and sends back the results to the plot agent [22]. The initial story world setting is based on an ontology describing the concepts of the story world. In the world of little red cap, this ontology contains concepts like cake and girl and relationships relevant to the domain, such as owns and likes.

4.1.3 Character Agents

Character agents simulate virtual characters and their life in a story world [22]. The character agents are semi-autonomous as the plot agent might provide them with specific goals [26]. The character agents pursue their goals and interact with each other if this potentially help them to achieve their goal [22]. Thereby, they produce an event sequence, which is captured by the plot agent as fabula.

The character agents imitate principles that improvisational theater actors follow, which make a successful piece of improvisational theater more than a random interaction [23]. Improvisational theater creates unplanned but compelling stories. There are four main principles used in improvisational theater: focus on emotional impact, have ideals, accept all offers, and frame the story world. We next describe how these principles are related to story generation and interactive storytelling in the Virtual Storyteller.

Emotions. Emotions are important as events in the story should affect the character and otherwise the story becomes boring quickly [23]. Therefore, character agents should be emotional. In general, characters should be believable, i.e. personality-rich [12]. One way to enable believable characters is to provide them with emotions [26].

In the Virtual Storyteller, characters' emotional states are represented by intensities of positive and their opposite negative emotions such as hope-fear and love-hate [26]. Intensities are represented by a natural number on a scale from -100 to 100.

The emotional state changes in reaction to events, actions and objects [26]. For instance, an agent experiences hope, if it comes close to achieving its goal, distress, if its current goal proves to be unachievable, pride, if it performs an action meeting its own standards, and joy, if it meets another agent which it likes. The character agent's personality model and parameters define whom it likes as well as the intensity and duration of the emotional effects. For example, a character may be frightened more easily than others, but may forget its fright more quickly.

Based on its emotional state a character agent may develop certain action tendencies, (such as friendly or aggressive behavior), which in turn influence the importance to its goals [26]. The character's personality model defines the relations between the emotions and the action tendencies. For example, strong fear could lead to passive behavior in one character, thus increasing the importance of the goal to flee, and to aggressive behavior in another, thus increasing the importance of attacking the object of fear.

Ideals. Ideals can be seen as a character defining a goal [23]. They allow conflicts to occur, a main principle underlying drama [25]. Disturbing an ideal will almost always prompt an interesting action in response, and hence promote storytelling. In the Virtual Storyteller, each character agent has global and/or episodic goals which will be adopted with a certain probability [26]. Episodic goals have a very high chance (90 out of 100) of being selected for pursuit throughout the episode. Emotional goals may also have a high probability, enabling character's emotion-induced action tendencies and leading to variation and unexpectedness in the generated stories. For example, a cowardly wolf with the episodic goal of killing the grandmother could be equally likely to adopt the goal of running away from her because of its fear.

Once a character has carried out the first step towards achieving some goal, that goal will be maintained until it is reached [26]. This prohibits too much switching between goals. Though, there are two exceptions: The goal might become unreachable by changes in the story world or the importance of some other goal (e.g., fleeing from an enemy) might exceed a pre-specified high value, due to an increase of the character's emotions.

A character's sudden change in behavior may seem incomprehensible to the user [19]. Such a sudden change might happen if the character adopts a new goal [26]. Therefore, the Virtual Storyteller's narrator explicitly describe the character's motivation.

Offers. Offers are possibilities to respond emotionally, prompts to do something, or possibilities to add new information about a player's view of either himself or the story world. They guide actors into situations where they can achieve emotional impact or interesting ideals [23]. Therefore, offers should be accepted and not blocked. Improvisational actors usually accept all offers, because the offers are also visible to the audience. Virtual characters can communicate and negotiate about offers hidden from the audience. Therefore,

a virtual character can decide if it can believably accommodate it.

Framing the story world by late commitment. When a scene in improvisational theater is started, the environment is not set [23]. Instead, to frame the story world actors make up their environment and define what it contains through offers, called *late commitment*. That is, the effects of an action are filled in the story as they appear, making the initial setup small and basic [22]. Hence, improvisational actors can construct the episode retroactive according to the needs of the scene [23].

In many story generation approaches the course of events is determined implicitly by the initial state of the story world and its characters [23]. In contrast, the Virtual Storyteller uses *framing operators* to add new knowledge as setting elements to the story world state whilst creating the illusion that the knowledge has already always been present [22]. Information that might be added include physical states and location of objects as well as social states such a relationships of characters. All characters have to accept the effects of a framing operator and pretend they have always been true ensure consistency. That is, the effects are not treated as new information and hence, emotional responses are prevented.

The Virtual Storyteller's character agents use a partial-order planner to determine which actions to perform to reach a particular goal [22]. A partial-order planner is an algorithm that searches for a solution for a problem, consisting of descriptions of the initial state, the goal, and possible actions [18]. A solution is a set of actions, possible conditions about the order of the actions, causal links (actions that meet preconditions of other actions), and no open precondition. The Virtual Storyteller's partial-order planner has one modification: If it cannot fulfil a precondition using the actions it knows, it can select either an event or a framing operator [22].

However, plans which contain actions that are only in the plan to fulfil an open precondition of a framing operator or event are not allowed [22]. Otherwise, character agents might execute non-authentic actions. For example, a wolf making a plan to get rid of his hunger by going to a nearby house just to accidentally discover the grandmother, does not seem to be very authentic. Though, if such an action was not added to the plan for this reason the plan is allowed. For example, if the wolf went to the house to hide from the hunter, the wolf's action to go to the house is allowed.

The introduction of new aspects of the story world as a result of late commitment is constrained by the story presentation medium and genre expectations [23]. In a visual medium, the displayed virtual world constrains the possibilities of introducing new aspects. For example, in a realistic 3D environment, it might appear strange when a gun pops up out of nowhere the moment the hunter wants to kill the wolf. A cartoon-like visual medium might offer more and textual medium the most flexibility. One can always rely on the potential presence of the gun unless its presence was specifically denied.

In addition, the story generation can use example story pieces, written from a plot perspective by a human author, to inspire the decisions of character and plot agents [24]. However, using cases as basis for problem solving does not guarantee to find solutions that are coherent to the audience. For example, if little red cap and big bad are generalized to be organisms, then the system might come up with a solution where little red cap eats the grandmother to satisfy its hunger, because it uses a case where a wolf did the same.

4.2 Storytelling

To turn the fabula into a story of narrative form, the storytelling component (or narrator) constructs a global structure of the story, removes unnecessary elements from the fabula, replaces the story elements in the document plan with abstract sentence structures, and produces referring expression [26, 22]. Finally these structures are converted into actual sentences. Next, we describe these components in more detail.

4.2.1 The Global Structure of a Story

The global structure of the story is constructed by converting the fabula into a document plan [22]. The document plan is a binary tree in which nodes contain the story elements from the fabula, and edges rhetorical relations between the elements. The edges are labeled by rhetorical terms (e.g., "then", "because", or "thereby") that illustrate the type of the causal relationship of the two story elements. Fabula elements that do not need to be explicitly expressed in the story, are removed. Such elements include outcomes and belief that are obvious or inferable to the reader.

Setting element are introduced by the cue "Once upon a time". Background information, like names and locations of the characters and properties of the story world, are added to the document plan. They provide elaborations, expressed by relative clauses, for example, "who was called". Edges between to consecutive actions are labeled by temporal terms such as "then" and "after". Motivation and psychological cause edges are labeled by terms like "because" and "therefore". Edges of enablement and physical causes are labeled by terms such as "so that" and "thereby". If two story elements together cause another story element, their rhetorical relation is additive and labeled with an "and". The "and" is also used if no other relation applies to the two story elements.

4.2.2 Generating Abstract Sentences, Natural Language Text and Speech

The story elements in the document plan are replaced with abstract sentence structures called dependency trees [22]. In dependency trees nodes represent lexical units and edges relationships between these units, and which adjacent node is dominant over the other one [9]. Each fabula element is transformed into a dependency tree [22]. For example, the characters and objects of an action specify subject and object, while instruments are

expressed by a prepositional argument. See, for example, the last sentence in the story in Figure 3: "[...], he filled the water supply with the water from a pond."

To achieve a coherent text and not a sequence of simple sentences, more complex sentences are produced [22]. Therefore, the narrator combines some dependency trees with each other. The edge label determined whether and how two dependency trees can be combined. For example, a construction such as "Little red cap cried and little red cap screamed" will be reduced to "Little red cried and screamed". Referring expressions to characters and objects are then generated using a pronoun such as "she", "he" or "it" in case the usage does neither produce confusion nor prevents variation [22].

Articles are put in front of the noun. Generally, an indefinite article "a" is used if the entity is mentioned for the first time, and a definite article "the" when the entity has been mentioned before. Exceptions are entities that has already been evoked by the mention of another object it is related to (e.g., if it was specified that each ship has a deck) [22].

Finally, natural language in text form is generated [22]. Using knowledge about syntax and morphology of the specific language, the words in each dependency tree are put into the correct order, nouns, adjectives and verbs are inflected, and punctuation is added. The story text can then be presented as text, in form of speech of a talking face [11], or presented on a multi-user tabletop interface [1]. Other visualizations, such as 3D graphical environments are theoretically possible, too [22]. Unfortunately, the Virtual Storyteller currently produces stories only in Dutch. Though, only the narrator component needs to be adopted for producing stories in another language.

5 Conclusion

We described how TALE-SPIN and the Virtual Storyteller create and present stories, and give an overview of the computational components involved. The presented example stories have simple, but dramatically interesting structures.

Even though most of the patterns (e.g., what to do to satisfy a specific goal) are directly specified in both systems, TALE-SPIN and the Virtual Storyteller, they produce various kinds of stories (within a single domain). From each story produced the user can draw a personal conclusion. For example, from the story of Figure 1 one could draw the conclusion, that one should better do what the bad guy is asking for, in order to prevent a negative outcome for oneself.

Both TALE-SPIN and the Virtual Storyteller share their main approach: character- and goal-directed behavior. Character's goals and their outcomes anchor the story and its episodic structures. As episodic structures are commonly found in stories [27], the character- and goal-directed behavior results in simple, but compelling stories. Though, the Virtual Storyteller is able to produce stories more close to human produced stories as it distinguishes between the fabula and narrative text generation. The fabula is in-

dependent of stylistic concerns (such as view points or time lapses) and therefore the stylistic choices can be made when generating the narrative text. Further, it allows to add or remove retroactive information pieces. For instance, the existence and location of objects (e.g. of an island) might be added to and fabula elements that do not need to be explicitly expressed in the story might be deleted. In addition, more complex sentences with temporal terms, references, etc. are possible due to the distinction. For example, referring, motivation, and temporal expressions, amongst others, can be added, because the previous and later fabula elements (goals, actions, events, ...) are known. In contrast, TALE-SPIN outputs each goal, subgoal, and action once it is produced and hence adding such expressions is quite limited (and was not included in TALE-SPIN). Furthermore, the stories generated by the Virtual Storyteller are undetermined as goals are adopted with a certain probability, whereas in TALE-SPIN the story produced will be the same until some emotion, relationship, goal, etc. will be changed by the user.

Moreover, there are several open questions and issues to address: (1) The stories produced are quite short and simple in comparison to, for example, Brother Grimm's fairy tales. How can more complex and longer stories be produced? (2) Dialogues are, if at all present, quite simple. How can a story generation system produce more complex dialogues? Further, in none of the stories character's thoughts are described. How could a system model thoughts and incorporate them in the generation of stories? (3) Both TALE-SPIN and the Virtual Storyteller considered only two dimensional emotions (e.g., hope-fear, or love-hate). Could more complex emotions be considered and would they improve the quality of the produced stories? (4) Currently goals, action plans (which aim to satisfy some goal), knowledge about the environment etc. are predefined in the system. How could they be generated automatically? Would this help to produce more creative stories as the story generation would be less user/programmer dependent? (5) Both systems use character's emotions, but the story's emotional effect on the audience is not considered. What is needed to make a story funny or sad? (6) There are some approaches to incorporate story generation into games. For instance, Façade can downloaded [13], the Virtual Storyteller's multi-user tabletop interface as well as FearNot! has been tested by children [1, 3]. However, we are not aware of any broad usage of story generation in commercial games. (7) Grammar and word choice (e.g., the usage of synonyms and antonyms) could still be improved [22]. (8) In addition, one could employ shifts in perspective, foreshadowing, and flashbacks [22].

From the point of creativity, both, TALE-SPIN and the Virtual Storyteller deepen the understanding of what is needed to generate an interesting story, and hence, maybe also creative. Simultaneously, they produce creative artefacts.

Acknowledgements. We would like to thank Mariët Theune for her comments by e-mail and for making the narrator component accessible to us.

References

[1] T. Alofs, M. Theune, and I. Swartjes. A tabletop interface for multi-user interaction with a storytelling system. In *International Conference on Intelligent Technologies for Interactive Entertainment (INTETAIN '11)*, 2011. To appear.

[2] M. Arinbjarnar, H. Barber, and D. Kudenko. A critical review of interactive drama systems. In *Adaptive and Emergent Behaviour and Compex Systems (AISB'09) Symposium on Artificial Intelligence and Games*, pages 15–26.

[3] R.S. Aylett, S. Louchart, J. Dias, A. Paiva, and M. Vala. FearNot! — an experiment in emergent narrative. In T. Panayiotopoulos, J. Gratch, R. Aylett, D. Ballin, P. Olivier, and T. Rist, editors, *Intelligent Virtual Agents*, pages 305–316. Springer, 2005.

[4] M. Cavazza, F. Charles, and S.J. Mead. Character-based interactive storytelling. *IEEE Intelligent Systems*, 17:17–24, 2002.

[5] R. Damiano, V. Lombardo, and A. Pizzo. Formal encoding of drama ontology. In G. Subsol, editor, *Virtual Storytelling. Using Virtual Reality Technologies for Storytelling*, LNCS 3805, pages 95–104. Springer, 2005.

[6] M. Esslin. *An Anatomy of Drama*. Sphere Books, 1976.

[7] G. Freytag. *Die Technik des Dramas*. S. Hirzel, 1863.

[8] B. Hayes-Roth, R. van Gent, and D. Huber. Acting in character. In R. Trappl and P. Petta, editors, *Creating Personalities for Synthetic Actors*, LNCS 1195, pages 92–112. Springer, 1997.

[9] P. Hellwig. Dependency unification grammar. In *International Conference on Computational Linguistics (COLING '86)*, pages 195–198, 1986.

[10] S. Klein, J. F. Aeschliman, D. Balsiger, S. L. Converse, C. Court, M. Foster, R. Lao, J. D. Oakley, and J. Smith. Automatic novel writing: A status report. Technical Report 186, Computer Science Department, University of Wisconsin Madison, 1973.

[11] F.L.A. Knoppel, A.S. Tigelaar, D.O. Bos, T. Alofs, and Z. Ruttkay. Trackside DEIRA: a dynamic engaging intelligent reporter agent. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS '08)*, pages 112–119, 2008.

[12] A.B. Loyall. *Believable Agents: Building Interactive Personalities*. PhD thesis, Carnegie Mellon University, 1997. CMU-CS-97-123.

[13] M. Mateas. *Interactive Drama, Art and Artificial Intelligence*. PhD thesis, Carnegie Mellon University, 2002. CMU-CS-02-206.

[14] J.R. Meehan. TALE-SPIN, an interactive program that writes stories. In *International Joint Conferences on Artificial Intelligence (IJCAI '77)*, pages 91–98, 1977.

[15] J.R. Meehan. Micro TALE-SPIN. In R.C. Shank and C.K. Riesbeck, editors, *Inside Computer Understanding: Five Programs Plus Miniatures*, pages 227–258. Lawrence Erlbaum Associates, 1981.

[16] J.R. Meehan. TALE-SPIN. In R.C. Shank and C.K. Riesbeck, editors, *Inside Computer Understanding: Five Programs Plus Miniatures*, pages 197–226. Lawrence Erlbaum Associates, 1981.

[17] M. O. Riedl. *Narrative Planning: Balancing Plot and Character*. PhD thesis, North Carolina State University, 2004.

[18] S.J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, 2nd edition, 2003.

[19] P. Sengers. Designing comprehensible agents. In *International Joint Conferences on Artificial Intelligence (IJCAI '99)*, pages 1227–1232, 1999.

[20] I. Swartjes. *Whose Story is it Anyway? How Improv Informs Agency and Authorship of Emergent Narrative*. PhD thesis, University of Twente, 2010. SIKS Dissertation Series No. 2010-20.

[21] I. Swartjes and M. Theune. A fabula model for emergent narrative. In S. Göbel, R. Malkewitz, and I. Iurgel, editors, *Technologies for Interactive Digital Storytelling and Entertainment*, LNCS 4326, pages 49–60. Springer, 2006.

[22] I. Swartjes and M. Theune. The Virtual Storyteller: Story generation by simulation. In *Benelux Conference on Artificial Intelligence (BNAIC '08)*, pages 257–265. University of Twente, 2008.

[23] I. Swartjes and J. Vromen. Emergent story generation: Lessons from improvisational theater. In *Conference on Artificial Intelligence (AAAI 07) Fall Symposium on Intelligent Narrative Technologies*, FS-07-05, pages 146–149. AAAI, 2007.

[24] I. Swartjes, J. Vromen, and N. Bloom. Narrative inspiration: Using case based problem solving for emergent story generation. In *Creativity and Cognition Conference (CC '07)*, pages 21–28, 2007.

[25] N. Szilas. Interactive drama on computer: Beyond linear narrative. In *Conference on Artificial Intelligence (AAAI '99) Fall Symposium on Narrative Intelligence*, 1999.

[26] M. Theune, S. Rensen, R. op den Akker, D. Heylen, and A. Nijholt. Emotional characters for automatic plot creation. In S. Göbel, U. Spierling, A. Hoffmann, I. Iurgel, O. Schneider, J. Dechau, and A. Feix, editors, *Technologies for Interactive Digital Storytelling and Entertainment*, LNCS 3105, pages 95–100. Springer, 2004.

[27] T. Trabasso, P. van den Broek, and S.Y. Sub. Logical necessity and transitivity of causal relations in stories. *Discourse Processes*, 12(1):1–25, 1989.

[28] M.J. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley and Sons, 2nd edition, 2009.