universität wien

Faculty of Computer Science

Technical Report #1/2013 (local report no. 2013001)

AUTOMOVIE: Ideas towards a complete interactive movie scripting and rendering system

Helmut Hlavacs, Giuseppe Maggiore, Yohann Pitrey

Research Group Entertainment Computing

AUTOMOVIE
Ideas towards a complete interactive movie scripting and rendering system

Helmut Hlavacs, Giuseppe Maggiore, Yohann Pitrey

May 13, 2013

Abstract

This technical report presents ideas for a system allowing a user to interactively and automatically generate full length movies. According to a set of user requirements, the system seeks a satisfying set of events that together build a story. The planner included in the system is based on a layered architecture that is able to handle the transition from story-scale and narrative goals, down to more concrete actions. Then, a rendering module takes care of the actual graphical creation of the movie.

We first present our views on the system requirements and provide a model of its architecture. Our demonstration is backed up by a literature review on automated storytelling and rendering. We then analyse the strengths and weaknesses of existing work and show that further efforts are needed in order to address the full system requirements.

1 Introduction

Movies and TV series attract billions of viewers and create interest worldwide. For more than a century writers, actors, directors and other professions in the movie industry constantly work on producing movies or series. Even though nowadays more and more processes can be carried out automatically, no completely automated movie generator is available. However, science has discovered many fundamental principles of movies such as storytelling and drama, and computer graphics are gradually able to render photo realistic scenes in real time.

These hardware developments now raise the question whether it is possible to extend movie production in a way that movies or TV series can be produced fully automatically and in real-time. We call such a movie *automovie*. Especially TV series, due to their repetitive nature, seem to fit perfectly to an automated production.

Due to the current status of knowledge and research in this domain, we hypothesize that competitive results are likely to be made within the next years. However, in this technical report, we state our ideas on how it actually could be done. We focus mainly on the automatic generation of movie scripts. Our demonstration is backed up by a literature review on automated script generation. We first refer to existing work in this domain along with the demonstration, then we provide a summary of the covered parts in the form of a table. This gives us an estimate of the areas where significant effort is still needed and allows us to orient our future work.

Since the envisioned scenario uses the movie script as input for a real-time rendering software, the script would not necessarily look like traditional scripts, but rather be a machine-readable data file, possibly encoded by using formats such as XML, which could even be further compressed by some text compressor. This idea of storing movies not frame by frame, but rather as compressed textual or binary description, leads to the fascinating possibility of dramatically reducing the size of a movie file, ranging from many Gigabytes for high definition content, to a much smaller amount of data, enriched with interactive components, for instance allowing the user to choose how the movie might proceed, or to make changes to the movie script on the fly.

2 State of the world

We base our ideas on a virtual world, acting as a container for all elements of the story. The virtual world is the medium on which the story is to be built. We first define a model for the world, based on objects and relations. A state of the world can either represent the current state or a future state to be reached at one point in the story. Subsets of the world are also useful. They can be used to represent a part of the world that is relevant at some point in the story. As will be developed in the following of the paper, milestones to be reached at some point of the story can be defined as subsets of the world.

2.1 Objects

The world is composed of objects [1, 3, 6, 7, 8, 11]. These objects can have a physical instantiation or not. Objects with a physical instantiation have physical characteristics such as shape, mass, colour, texture, *etc*. They obey some physical rules specified by a physics simulator, such as gravity and friction.

Objects with no physical instantiation can be institutions or abstract ideas. Entities such as companies and concepts of the world can be modelled as non-physical.

2.2 Characters

Characters, also called agents, are particular objects of the world that have a dynamic and conscious behaviour. They can perform actions in order to modify the current state of the world. Characters can be human or animals. For the sake of the story, the concept can be extended beyond these two categories (*e. g.* magical creatures, machines, *etc*). Characters are driven by internal goals and influenced by internal states, and have their own personality [1, 2, 3, 4, 5, 6, 9, 10]. Goals are constraints that the character seeks to fulfill, while states alter the reactions of the character to its environment. The personality of a character can be defined as a combination of goals and internal states. The internal traits of a character can evolve under the influence of the environment. For instance, seeing its actions being jeopardized by external events can make a character become less optimistic or more frustrated. Different traits might have different strength in the character's personality and be harder or easier to change. A highly optimistic person might remain happy under bad circumstances, whereas a depressive personality might become sad under small inconveniences.

The personality of a character can be modelled for instance using a combination of the five basic components of a personality, such as presented by Ekman in [12]. A set of attributes can be linked to the character representing its levels of happiness, anger, fear and disgust.

Due to higher level decisions by the system, a character can see its plans jeopardized or delayed. Each character has a level of resiliency to unexpected events. A character with high resiliency will be quite flexible and insensitive to external events. On the opposite, a character with a low resiliency will be deeply affected by external events. Resiliency can be conditioned by the good or bad aspect of the unexpected event. It can also be influenced by the physiological and psychological internal states of the character.

Each character has a physical and physiological setup that determines its appearance and the way it can evolve in the world [6, 9, 11]. For human characters, the physical setup determines the visual appearance of the character, with physical features such as size, ethnic origin, amount of body fat, type of voice, *etc*. The physiological setup determines the way the character moves. An old man will for instance not move the same way as a young child. Distinct characters might have different abilities for sports, different ways of walking, or expressing itself.

Additionally, a character's reactions are influenced by its psychological setup [2, 3, 4, 5, 6, 8, 9, 10, 11]. The psychological setup can be seen as a combination of emotions. Emotions are an important part of a character's setup as they influence its reactions and encourage the expression of different aspects of its personality. The emotional state of a character can be modelled as a combination of basic emotions, such as suggested in the work of Ekman (*i. e.* anger, disgust, fear, happiness, sadness and surprise) [12]. Actions and events have an impact on a character's emotions and can in turn modify its reactions. A given combination of chemicals might influence the character's reactions in distinct ways. For instance a character with a high level of anger might react to a threat by striking back, whereas a character with a high level of fear might choose to flee from the threat.

One key notion of a character's perception of the world can be represented by personal belief [4, 8, 9, 11]. Two types of belief can be identified: knowledge and expectation.

Knowledge, or static belief, is a subset of the past and present world, as perceived by the character. This subset of the world can be true or false, as it is only the way it is perceived through the character's eyes. Knowledge is a rather static concept, that can be modified by observations.

Expectation, or dynamic belief, includes the character's predictions about the way the world is changing. It allows a character to anticipate events, and plan its actions ahead. As this part of perception is subject to validation by current observations of the character, it is highly uncertain and can change in a faster way than knowledge. Naturally, expectation influences the way the character reacts to the world.

2.3 Relations

Objects are connected to each other using relations [8, 11]. Relations can be physical, to represent the structure of the world. Relations can also be symbolic, such as love, friendship, or profession. They can depend on other relations in order to be established or maintained. A relation can require a set of other relations to be established in order to exist itself.

Relations can exist between characters and characters [1, 2, 4, 5, 6, 8, 9, 11]. A character can create or break a relation with another character. Relations between characters need not be in both directions. A character can have relations to itself, such as being angry at or proud of itself [2, 11]. Characters can establish relations with surrounding objects [1, 6, 11]. A relation between a character and an object can be physical or symbolic. For instance, sitting on a chair is a physical relation, while working in a company is a symbolic relation between a character and a non-physical object.

3 Possible actions by the agents

Actions are preformed by characters. They modify the state of the world. They can have pre-conditions and require other objects or characters in order to be executed. There are multiple levels of actions. High-level actions are decomposed into simpler, lower-level actions, until only a succession of atomic actions is reached. The planner handles this process, such as explained in the following section.

Actions performed by characters can have an impact on both objects and characters [6, 8, 11]. The impact on characters can be in terms of emotions [4, 5, 6, 9, 10, 11] or in terms of physical integrity. Actions can also have an impact on relations [4, 5, 6, 8, 9, 10, 11]. Stealing an object from a character might question the relation of ownership between the character and the object. Stealing the object from the character might also increase its level of anger or sadness.

4 Planner

The planner is a central module of the system, handling the transition from the current state of the world to a target state. Its main role is to find a succession of actions allowing the world to change from one state to another in a coherent way. Some existing approaches do not use a planner, but some heuristic approach [6, 7, 9, 10]. However, we believe that using a planner is the best option, as it allows implementation of very useful features such as the ones we describe in this section [1, 2, 3, 4, 5, 8, 11].

4.1 Hierarchy of actions

The planner can be built around a layered architecture [1, 3, 6, 7, 8, 10, 11]. The higher layers comprise story-level considerations, such as the various chapters of the story. The lower layers handle increasingly concrete sequences of events, down to simple actions. Each layer has a set of goals to reach and constraints to respect. The goals are formulated as subsets of the state of the world. They represent milestones to be reached. The constraints are not directly written, but are inferred by the planner in order to maintain the coherence of the story.

A given layer constrains the next lower layer. The constraints of the higher layer become the goals of the lower layer. For instance, if the goal of the higher layer is that characters A and B fall in love, the constraints of the lower layer is that the succession of actions lead to A falling in love with B.

4.2 Constrained randomness

In the story generation, choices can be randomized in order to make the system more flexible and avoid the same scenarii repeating over time. However, pure randomness might lead to incoherent decisions regarding the story [1, 3, 6, 9]. Therefore, the random aspect of the system has to be constrained. A set of variables have to be considered when choosing the best action, by affecting probabilities to the different available options.

In addition to having a consequence on emotions, relations and characters, actions have an impact on the dramatic colour of the story, by influencing concepts such as suspense, violence, romance or humour. For instance, preventing a character to perform an obvious action might increase suspense, whereas making a character hit another might increase violence.

Characters' internal state also has to be taken into account when making a random decision. As the planner has a full knowledge of the surrounding world, it is possible to anticipate the impact of a given set of actions on the plans of a character. The planner can take this into account in order to drive the story along the chosen line. For example, knowing that a character wants to escape from some place can allow the planner to decide that it will prevent the escape, in order to increase the level of suspense of the story.

4.3 Resiliency

The system can be made interactive during the story presentation. That is, the user might be allowed to modify the story as it goes, by manipulating the decisions made by the planner. In this case, the planner must take these modifications into account and replan the sequence of actions in order to freeway the chosen goals [1, 2, 3, 6, 10].

Two main approaches regarding the planners behaviour towards external intervention can be outlined. On one hand, the planner can try and find the best set of actions in order to reach the specified goals for the story. In this case, external decisions can severely change the plans at the next step, as the planner does not take into account the possibility that an external decision can act against it. On the other hand, the planner can consider that some external decisions will occur, and choose to perform actions that are more resilient in the possibility of change.

4.4 User profiling and prediction

In case the system is made interactive along the story presentation, a highly interesting feature is user profiling. The system learns progressively what kind of choices are favoured by the user, and adapts the future development of the story accordingly [7]. This aspect of the system will be discussed in more details in the section on online interaction with the player.

4.5 Searching for a suitable solution

The whole problem of storytelling can be seen as an optimization problem, where the planner has to choose the right succession of actions, in order to create a suitable and entertaining story. In order to make the system usable in realistic entertainment scenarii, real-time or close-to-real-time constraints might be applied [1, 2, 3, 4, 5, 6, 9, 10]. Indeed, the complexity of the system should yield acceptable loading times. In a possible extension, the player could control the complexity of the system, in order to find an acceptable tradeoff between responsiveness and flexibility.

5 Authoring tools

The goal of the system is not to take the task of storytelling completely away from the script writer. Instead, we rely on the writers expertise or desire to adjust the dramatic parameters. To this end, authoring tools need to be provided for the writer to help the system understand the guidelines of the story. This type of interaction is performed offline, *i. e.* before the story is actually generated. The writer could be the same person as the final viewer, or be someone different.

5.1 Specify states of the world

In order to control the development of the story in terms of milestones to be reached, the writer can describe subsets of the world state. First, the writer can describe the initial state of the world in which the story will play [4, 5, 8]. To do so, either a procedural or descriptive specification should be provided. This would represent an extension of the existing work, as it has never been included in a story generator. The writer provides high level parameters describing the contents of the world, the characters involved in the story, the relevant relations between the elements of the world.

The author can also describe intermediate states of the world, that the script has to fulfill at some point in the story. To the best of our knowledge, this aspect has not been studied in the past. These intermediate states can be expressed as a set of conditions to be met, such as required relations between characters and/or objects. The final state of the world that the story should reach can also be provided by the author [8]. The story ends only when the final substate of the world is reached.

5.2 Well-formedness of a story

One very important input from the story writer is the guidelines for well-formedness of the story. As the end user might have a limited or no expertise in storytelling, a writer who has this kind of expertise could provide high level structure information, so that the story follows a globally coherent path.

Some narration models can be used to define the struggle if the story [4, 5, 8, 10]. The global optimization function for the planner is specified in this step. Several parameters have to be defined. The coloration of the story can be specified by the author. Here, the balance between suspense, violence, romance and humour, is determined. A romantic comedy is likely to have a high level of humour and romance, whereas an action movie might have a higher level of violence and suspense.

The balance between determinism and randomness is also set buy the writer [10]. If the writer wants to generate a very stable story that is likely to repeat similar patterns, randomness can be set to a low level. On the opposite, in order to make the system generate more unique stories, the randomness can be set to a higher level.

5.3 Parameters for replanning/recovering

The writer also has the possibility to specify the characteristics of the planner in terms of resilience and replanning [10]. A resilient planner will replan its actions when the user interacts with the system, whereas a non resilient planner will try to find a coherent chain of actions in order to meet the previously established goals.

5.4 Allowed interactions

The writer can also specify what kind of interactions the player will be allowed to have with the system [10]. For instance, the writer can forbid the user to modify the story during the presentation. The interaction modalities can also be defined by the writer. He can choose to allow only predefined actions, or leave the user free to interact in a free manner.

The writer can also allow the user to go back in time and modify prior decisions, or constrain the user to modify only the current action [10].

6 Player interactions during presentation

In addition to offline interaction, the system might provide the possibility for the final user to interact with the story while watching the movie.

6.1 Player actions

The player might be allowed to modify the script itself, in order to alter the course of events. In this case, the script should appear in a readable and editable way for a non-expert user. To the best of our knowledge, this functionality had never been integrated in a story generator. Still, it enables a brand new way of interaction with the user and increase significantly the feeling of interactivity of the system. However, it raises questions about the interaction modality with the end user. One first option would be to allow a fixed set of alterations on the script. The user might then allow two characters to fall in love, or on the contrary, decide to separate them. The two available options would be presented to the user, who would choose one.

Another type of possible interaction with the end user is allowing the world to be directly modified [1, 3, 4, 5, 6]. For example, objects can be moved in order to make them unreachable by a given character.

6.2 Non-linear interaction

If the user is allowed to modify not only the current state of the story, but also past and future events, a whole new aspect of storytelling can be enabled. The system could present him the consequences of a given action on the script out in the world, and the user could decide if he wants to proceed, or to revise his choice. This possibility has not been studied in the literature, as far as we could find out. It raises obvious questions about computational complexity, as the system would have to generate arbitrary ranges of the script and be able to present them to the user in an understandable way.

6.3 Modalities

Various modalities can be considered for user interaction. The simplest medium of interaction is to present a predefined list of options to the end user. The list contains actions that make sense with regard to the story, which decreases the potential out-of-bounds decisions from the user. Finding a list of coherent actions might however represent a challenge in itself.

An alternative would be to let the user build sentences using a semi automated approach. The structure of the action would be provided, letting the user fill blanks. A list of available characters and objects could for instance be presented, together with a list of available actions. It would then be possible to restrain the available interactions to a coherent set.

Semi automated writing could of course be extended to free writing, which is a more natural way of expression [4, 5, 6]. The challenges related to this functionality are obviously related to the highly variable structure of written language and their interpretation.

A yet more natural but more challenging medium is speech recognition [5]. Considering the current state of speech recognition, this type of interaction might however be manageable under the same limitations as expressed for semi automated writing.

Finally, some approaches constrain the user interaction to use a traditional game controller [1, 3, 5, 6, 7, 11]. This way, the possible interactions are limited to the capabilities of the controller, and might be easier to handle than speech or free writing.

7 Graphical rendering

Current developments in off-the-shelf computer graphics hardware recently allowed the photorealistic rendering of humans, especially faces, as well as objects, *etc*[1] in real-time. Key components of these developments are massive parallelism inside graphics cards, using thousands of general purpose computing cores on one card. The current Shader Model 5.0, as for instance built into DirectX 11 compatible cards, also allows hardware based tesselation, *i. e.* refining object surfaces by hardware. New developments in computer graphics have revealed better lighting models, such a subsurface scattering and screen space ambient occlusion, which allow a much more realistic rendering result in real-time. Current AAA-game titles already show impressive graphics, including *e. g. Battlefield 3* or *Dead Space 3*, and made by leading game engines like Crytec's CryEngine [2] or the Fox Engine [3].

Graphics assets should be stored on a database available at the rendering node. However, in order to create a large variety of different characters, objects or scenes, we distinguish between static and dynamic assets. Static assets are those that do not change their appearance, like textures, traffic signs *etc*. Dynamic assets however should be stored together with a set of parameters. Such morphable models are known to be able to capture faces or bodies. Since each set of concrete parameter values determines a different appearance, random characters can be easily created by picking random numbers from pre-determined value intervals. We therefore propose to find appropriate morphable models for most of the world geometry, including, characters, vehicles, furniture, houses, *etc*.

Complementary to graphics is the question of simulating the physical aspects of the world. Though movements in principle can be pre-determined by manually creating a large amount of static animations, we envision a more elaborate principle making use of physics simulation engines as used in modern computer games. Middleware such as *Bullet 3D*, *Havok* or *PhysX* can be used to create believable interactions between objects and characters in real-time. For driving virtual charaters we envision a set of controllers that act as muscles in the limbs. Such controllers can be automatically created off-line, for instance by using genetic algorithms. Other moving objects such as vehicles are driven by virtual motors and force generators. This way, new and surprising dynamics can be created online.

[1] http://blogs.nvidia.com/2013/03/a-demo-thats-truly-a-head-of-its-time/
[2] http://www.crytek.com/cryengine
[3] http://www.geek.com/games/hideo-kojima-unveils-his-new-fox-game-engine-1386361/

7.1 Text, Speech and Music

Spoken text is still a problem in a fully synthesized environment. As a first solution we envision a database with text pieces, statically atttached to script actions. Of course, when creating the script, text pieces must be picked at random and assembled into a larger unit of text. Once text to be spoken is known, text-to-speech (TTS) systems can then be used to synthesize the language. Here we emphasize that emotion is an integral part, and only emotional TTS systems are useful. Also it is possible to analyse spoken text to compute the visual representation of a human mouth. Such *visemes* look the same on each face, and a mapping between face muscles and the resulting visemes can be inferred using models such as proposed by Ekman [12]. By using an appropriate face muscle model and a skin simulation, any kind of facial expression and viseme can be synthesized, and synthesized language and mouth movements can be synchronized to result in credible human speech.

Likewise, various research groups world wide have succeeded in synthesizing artificial music. Music synthesis can be generated by learning from real music of different styles. After a learning phase, the music synthesizer creates similar music, including some random patterns to make it sound new. Current results are good enough to be perceived as pleasing and it is to be expected that in a few years, movie scores can be automatically created in realtime, following the story plot in a way similar to current movies.

8 Summary, open questions and future work

As stated in the introduction, this paper focuses on the script generation step of the design of an interactive movie generator. Through our demonstration, we presented our views on the structure of this part of the system and referred to existing work when it was possible. In order to provide a summary of our literature review, Figure 1 presents the cited contributions facing the important components of the scripting system. Each cell contains a *V* if the corresponding approach includes the element in its design. We added a column named *#Hits* showing how many approaches cover a given element.

This table enables a rapid reading of the c vintage of the system. First, it appears quite obviously that no approach covers the whole system as we described it. This is a technical element to explain why no existing approach has yet been largely successful at generating movies in a large scale. It lets us hypothesize that a proper combination of all these elements might have better chances of success among industrials and eventually among end users.

In more details, one can identify elements on which work is still particularly needed, as they have rarely been covered, or not covered at all. First of all, authoring tools show the lowest amounts of hits. A meaningful contribution would therefore be to provide innovative ways for the story writing expert to create basic movie structures that could be adapted to the user's desires. Intermediate states of the world assist to be a very interesting issue, as they allow the writer to restrain the story to a given frame, while allowing some randomness to be added, in order to create always innovative stories. A framework allowing the writer to specify the structure of the word in a procedural way would also constitute a very interesting asset.

Secondly, user profiling has not received repeated coverage, even though it appears to us as a dominant feature in terms of interactivity and customisation. Learning how to analyse the user's actions and how to anticipate them over time should therefore receive more efforts.

Allowing the user to modify the script in a textual or similar way has not been proposed in the articles we identified as the major contributions in the domain. This feature adds a brand new aspect to story telling, allowing the user to tune the generated script itself, giving him a new kind of control on the story. Nevertheless, it raises difficult questions on how to present the script to the user and how to structure the interactions.

Non linear interaction has not been covered either in the literature. This affect us also very innovative and challenging. It enables the user to travel in time within the story, allowing him to evaluate the impact of his choices on the development of the story. Designing such a tool could allow the system to be used not only in entertainment scenarii, but also in didactic situations where the user would be allowed to test different ways of reacting to events and learn from this experience. One could imagine using such a system to assist learning about social skills or overcoming fears.

Finally, we want to mention that the rendering part of the system is itself a challenging topic. Building a credible graphical movie out of a textual description of the script is a complex task and requires more work. One of our main tasks in our future work will be to provide more efforts for this part in terms of literature review and to identify the technical and theoretical difficulties.

	# Hits	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]	[11]
State of the world		-	-	-	-	-	-	-	-	-	-	-
Modeling of objects (physical and non)	6	V	-	V	-	-	V	V	V	-	-	V
Characters personality	8	V	V	V	V	V	V	-	-	V	V	-
Characters goals	8	V	V	V	V	V	V	-	-	V	V	-
Characters physical setup	3	-	-	-	-	-	V	-	-	V	-	V
Characters psychological/emotional setup	9	-	V	V	V	V	V	-	V	V	V	V
Characters beliefs (static and dynamic)	4	-	V	-	-	-	-	-	V	V	-	V
Relations between objects	2	-	-	-	-	-	-	-	V	-	-	V
Relations between characters	8	V	V	-	V	V	V	-	V	V	-	V
Relations between characters and themselves	2	-	V	-	-	-	-	-	-	-	-	V
Relations between characters and objects	3	V	-	-	-	-	V	-	-	-	-	V
Possible actions by the agents	0	-	-	-	-	-	-	-	-	-	-	-
Impact on characters and objects	3	-	-	-	-	-	V	-	V	-	-	V
Impact on relations	7	-	-	-	V	V	V	-	V	V	V	V
Impact on emotions	6	-	-	-	V	V	V	-	-	V	V	V
Planner	7	V	V	V	V	V	-	-	V	-	-	V
Hierarchy of actions	7	V	-	V	-	-	V	V	V	-	V	V
Constrained randomness	4	V	-	V	-	-	V	-	-	V	-	-
Resiliency to breakage and interaction	5	V	V	V	-	-	V	-	-	-	V	-
User profiling and prediction	1	-	-	-	-	-	-	V	-	-	-	-
Planner evaluates consequences of actions	8	V	V	V	V	V	-	-	-	V	V	V
Plan in real time	8	V	V	V	V	V	V	-	-	V	V	-
Authoring tools		-	-	-	-	-	-	-	-	-	-	-
Initial state of the world	3	-	-	-	V	V	-	-	V	-	-	-
Procedural world description	0	-	-	-	-	-	-	-	-	-	-	-
Intermediate states of the world	0	-	-	-	-	-	-	-	-	-	-	-
Final/desired state of the world	1	-	-	-	-	-	-	-	V	-	-	-
Hierarchy of actions	4	-	-	-	V	V	-	-	V	-	V	-
Narratological model	4	-	-	-	V	V	-	-	V	-	V	-
Randomization parameters	1	-	-	-	-	-	-	-	-	-	V	-
Parameters for replanning/recovering	1	-	-	-	-	-	-	-	-	-	V	-
Allowed actions	1	-	-	-	-	-	-	-	-	-	V	-
Allowed interaction modalities	1	-	-	-	-	-	-	-	-	-	V	-
Allowed linear/non-linear navigation	1	-	-	-	-	-	-	-	-	-	V	-
Possible actions by the user		-	-	-	-	-	-	-	-	-	-	-
Actions that modify the script	0	-	-	-	-	-	-	-	-	-	-	-
Actions that modify the world	5	V	-	V	V	V	V	-	-	-	-	-
Non-linear interactions	0	-	-	-	-	-	-	-	-	-	-	-
Text parsing/recognition	3	-	-	-	V	V	V	-	-	-	-	-
Speech recognition	1	-	-	-	-	V	-	-	-	-	-	-
Traditional game controllers	6	V	-	V	-	V	V	V	-	-	-	V

Figure 1: Summary of the parts of the system covered by the relevant literature.

References

[1] Riedl, M.O. and Stern, A. Believable Agents and Intelligent Story Adaptation for Interactive Storytelling. Technologies for Interactive Digital Storytelling and Entertainment 4326, 1 - 12 (2006).

[2] Marc Cavazza, Fred Charles: Dialogue Generation in Character-based Interactive Storytelling. AIIDE 2005: 21-26

[3] Marc Cavazza, Fred Charles, and Steven J. Mead. 2002. Emergent situations in interactive storytelling. In Proceedings of the 2002 ACM symposium on Applied computing (SAC '02). ACM, New York, NY, USA, 1080-1085.

[4] Pizzi, D., Charles, F., Lugrin, J.-luc and Cavazza, M. Interactive Storytelling with Literary Feelings 2 Emotional Planning for Interactive Storytelling. Language 4738, 630-641 (2007).

[5] Cavazza, M., Lugrin, J.L., Pizzi, D. and Charles, F. Madame Bovary on the Holodeck : Immersive Interactive Storytelling. Virtual Reality 651-660 (2007).

[6] Mateas, M. and Stern, A. Towards Integrating Plot and Character for Interactive Drama. In Working notes of the Social Intelligent Agents The Human in the Loop Symposium AAAI Fall Symposium Series Menlo Park 2002, 113-118 (2000).

[7] Thue, D., Bulitko, V., Spetch, M. and Wasylishen, E. Interactive Storytelling : A Player Modelling Approach. Artificial Intelligence Associatio, 43-48 (2007).

[8] Porteous, J. Cavazza, M. Controlling Narrative Generation with Planning Trajectories: The Role of Constraints. Interactive Storytelling 5915, 234-245 (2009).

[9] Peinado, F., Cavazza, M. and Pizzi, D. Revisiting Character-Based Affective Storytelling under a Narrative BDI Framework. Lecture Notes in Computer Science 1-6 (2008).

[10] Spierling, U., Grasbon, D., Braun, N. and Iurgel, I. Setting the scene: playing digital director in interactive storytelling and creation. Computers and Graphics 26, 31-44 (2002).

[11] Szilas, N., IDtension: A narrative engine for interactive drama. Proc. of the 1st Int. Conf. on Technologies for Interactive Digital Storytelling and Entertainment (2003)

[12] Ekman, Paul (1992). An argument for basic emotions. Cognition and Emotion 6: 169200.