

Towards Integrating Plot and Character for Interactive Drama

Michael Mateas

Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213
michaelm@cs.cmu.edu

Andrew Stern

www.interactivestory.net
andrew@interactivestory.net

Introduction

Interactive drama concerns itself with building dramatically interesting virtual worlds inhabited by computer-controlled characters, within which the user (hereafter referred to as the player) experiences a story from a first person perspective (Bates 1992). Over the past decade there has been a fair amount of research into believable agents, that is, autonomous characters exhibiting rich personalities, emotions, and social interactions (Mateas 1997; Bates, Loyall and Reilly 1992; Blumberg 1996; Hayes-Roth, van Gent and Huber 1997; Lester and Stone 1997; Stern, Frank, and Resner 1998). There has been comparatively little work, however, exploring how the local, reactive behavior of believable agents can be integrated with the more global, deliberative nature of a story plot, so as to build interactive, dramatic worlds (Weyrauch 1997; Blumberg and Galyean 1995). The authors are currently engaged in a two to three year collaboration to build an interactive story world integrating believable agents and interactive plot. This paper provides a brief description of the project goals and design requirements, discusses the problem of autonomy in the context of story-based believable agents, and finally describes an architecture that uses the dramatic beat as a structural principle to integrate plot and character.

Design requirements

The design requirements for the project are divided into two categories: project requirements and story requirements.

Project requirements

The project requirements are the overarching goals for the project, independent of the particular interactive story expressed within the system.

Artistically complete. The player should have a complete, artistically whole experience. The system should not be a piece of interactive drama technology without a finished story, nor only a fragment of a story.

The experience should stand on its own as a piece of art, independent of any technical innovations made by the project.

Animated characters. The characters will be represented as real-time animated figures that can emote, have personality and can speak.

Interface. The player will experience the world from a first-person 3D perspective. The viewpoint is controlled with the keyboard and mouse. The perspective may occasionally automatically shift to a third-person perspective to show action that is difficult to show from first-person.

Dialog. Dialog will be the primary mechanism by which a player interacts with characters and influences how the story unfolds. To achieve dialog, the player types out text that is visible on screen; the computer characters' dialog is spoken speech with simultaneously displayed text. The conversation discourse is real-time; that is, if the player is typing, it is as if they are speaking those words in (pseudo) real-time. The system should be very robust when responding to inappropriate and unintelligible input. Although the characters' dialog and "intelligence" are narrowly focused around the topic of the story, the characters have a large variety of responses to off-the-wall remarks from the player. (For example, if the player says "Do you ever go camping?", the characters can respond with "We hate the outdoors".)

Interactivity and plot. The player's actions should have a significant influence on what events occur in the plot, which are left out, and how the story ends. The plot should be generative enough that it supports replayability. Only after playing the experience 6 or 7 times should the player begin to feel they have "exhausted" the interactive story. In fact, full appreciation of the experience requires the story be played multiple times.

Change in the plot should not be traceable to distinct branch points; the player will not be offered an occasional small number of obvious choices that force the plot in a different direction. Rather, the plot should be smoothly mutable, varying in response to some global state which is itself a function of the many small actions performed by the player throughout the experience.

Even when the same plot plays out multiple times, the details of how the plot plays out, that is, the exact timing

of events and the lines of dialog spoken, should vary both as a function of the player's interaction and in response to "harmless" random variation, that is, random variation that expresses the same thing in different ways.

Distributable. The system will be implemented on a platform that is reasonably distributable, with the intention of getting the interactive experience into the hands of as many people as possible. It should not just be an interesting demo in a closed door lab, but be experienced by people in the real world. Ultimately, this is the only way to validate the ideas.

Story Requirements

The story requirements describe the properties that the story itself should have. These are not intended to be absolute requirements; that is, this is not a description of the properties that all interactive stories must have. Rather, these requirements are the set of assumptions grounding the design of this particular interactive story we intend to build.

Short one-act play. Any one run of the scenario should take the player 10 to 15 minutes to complete. We focus on a short story for a couple of reasons. Building an interactive story has all the difficulties of writing and producing a non-interactive story (film or play) plus all the difficulty of supporting true player agency in the story. In exploring this new interactive art form it makes sense to first work with a distilled form of the problem, exploring scenarios with the minimum structure required to support dramatically interesting interaction. In addition, a short one-act play is an extreme, contrarian response to the many hours of game play celebrated in the design of contemporary computer games. Instead of providing the player with 40 to 60 hours of episodic action and endless wandering in a huge world, we want to design an experience that provides the player with 10 to 15 minutes of emotionally intense, tightly unified, dramatic action. The story should have the intensity, economy and catharsis of traditional drama.

Relationships. Rather than being about manipulating magical objects, fighting monsters, and rescuing princesses, the story should be about the emotional entanglements of human relationships. We are interested in interactive experiences that appeal to the adult, non-computer geek, movie-and-theater-going public.

Three characters. The story should have three characters, two controlled by the computer and one controlled by the player. Three is the minimum number of characters needed to support complex social interaction without placing the responsibility on the player to continually move the story forward. If the player is shy or confused about interacting, the two computer controlled characters can conspire to set up dramatic situations, all the while trying to get the player involved.

The player should be the protagonist. Ideally the player should experience the change in the protagonist as a

personal journey. The player should be more than an "interactive observer," not simply poking at the two computer controlled characters to see how they change.

Embodied interaction should matter. Though dialog should be a significant (perhaps the primary) mechanism for character interaction, it should not be the sole mechanism. Embodied interaction, such as moving from one location to another, picking up an object, or touching a character, should play a role in the action. These physical actions should carry emotional and symbolic weight, and should have a real influence on the characters and their evolving interaction. The physical representation of the characters and their environment should support action significant to the plot.

Action takes place in a single location. This provides unity of space and forces a focus on plot and character interaction.

The player should not be over-constrained by a role. The amount of non-interactive exposition describing the player's role should be minimal. The player should not have the feeling of playing a role, of actively having to think about how the character they are playing would react. Rather, the player should be able to be themselves as they explore the dramatic situation. Any role-related scripting of the interactor (Murray 1998) should occur as a natural by-product of their interaction in the world. The player should "ease into" their role; the role should be the "natural" way to act in the environment, given the dramatic situation.

The Story

The particular story we plan to build, which satisfies the project and story requirements, is a domestic drama in which a married couple has invited the player over for dinner. (Assume for the moment that the player's character is male.) Grace and Trip are apparently a model couple, socially and financially successful, well-liked by all. Grace and Trip both know the player from work. Trip and the player are friends; Grace and the player have gotten to know each other fairly recently. Shortly after arriving at their house for dinner, Grace confesses to the player that she has fallen in love with him. Throughout the rest of the evening, the player discovers that Grace and Trip's marriage is actually falling apart. Their marriage has been sour for years; deep differences, buried frustrations and unspoken infidelities have killed their love for each other. How the veneer of their marriage cracks, what is revealed, and the final disposition of Grace and Trip's marriage, and Grace and the player's relationship, depends on the actions of the player.

The above story description assumes a male player. Ideally the player will be able to choose whether they wish to be a male or female player (important to support the "player should not be over-constrained by a role" story requirement). In the case of a female player, the story would play itself out symmetrically, with Trip confessing his love for the player. For the purposes of this

story, we are assuming heterosexual relationships. Ideally, sexual orientation would be selectable by the player as well.

Given these project and story requirements, many technology issues are raised, including interface issues, integrating plot and character, and supporting dramatic dialog. The rest of this paper will focus on the particular issue of integration of plot and character.

Autonomy and Story-Based Believable Agents

Most work in believable agents has been organized around the metaphor of strong autonomy. Such an agent chooses its next action based on local perception of its environment plus internal state corresponding to the goals and possibly the emotional state of the agent. All decision making is organized around the accomplishment of the individual, private, goals of the agent. Using autonomy as a metaphor driving the design of believable agents works well for believable agent applications in which a single agent is facilitating a task, such as instructing a student (Lester & Stone 1997), or giving a presentation (Andre, Rist, and Mueller 1998), or in entertainment applications in which a user develops a long-term relationship with the characters by "hanging-out" with them (Stern, Frank, and Resner 1998). But for believable agents used as characters in a story world, strong autonomy becomes problematic. Characters in a story world are there not to believably convey their personalities but rather to have the right characteristics to take the actions required to move the story forward. That is, knowing which action to take at any given time depends not just on the private internal state of the agent plus current world state, but also on the current story state. And the current story state includes information about all the characters involved in the story, plus the entire past history of the interaction considered *as a story*, that is, as a sequence of actions building on each other and moving towards some end. The global nature of story state is inconsistent with the notion of an autonomous character that makes decisions based only on private goal and emotion state and local sensing of the environment.

Only a small amount of work has been done on the integration of story and character. This work has preserved the strong autonomy of the characters by architecturally dividing the responsibility for state maintenance between a drama manager, which is responsible for maintaining story state, and the believable agents, which are responsible for maintaining character state and making the moment-by-moment behavior decisions (Weyhrauch 1997; Blumberg and Galyean 1995). These two components communicate via a narrow-bandwidth, one-directional interface flowing from drama manager to agent. The messages sent across this interface consist of goals that characters should assume or perhaps specific actions they should perform. The character is still responsible for most of the decision making. Occasionally the drama manager will modify one or more of the

characters behaviors (by giving them a new goal or directly instigating a behavior) so as to move the plot along. In the absence of the drama manager, the character would still perform its normal autonomous behavior. The idea seems to be that one can author fully autonomous believable agents which are able to convey their personalities in the absence of any story, drop them into a story world being managed by a drama manager, and now have those characters participate in the story under the drama manager's guidance.

This architecture makes several assumptions regarding the nature of interactive drama and believable agents: drama manager decisions are infrequent, the internal structure of the believable agents can be reasonably decoupled from their interaction with the drama manager, and multiple-character coordination is handled within the agents. Let's explore each of these assumptions.

Infrequent guidance of strongly autonomous believable agents means that most of the time, behavior selection for the believable agents will occur locally, without reference to any (global) story state. The drama manager will intervene to move the story forward at specific points; the rest of the time the story will be "drifting," that is, action will be occurring without explicit attention to story movement. Weyhrauch (Weyhrauch 1997) does state that his drama manager was designed for managing the sequencing of plot points, that is, for guiding characters so as to initiate the appropriate next scene necessary to make the next plot point happen (whatever plot point has been decided by the drama manager). Within a scene, some other architectural component, a "scene manager," would be necessary to manage the playing out of the individual scene. And this is where the assumption of infrequent, low-bandwidth guidance becomes violated. As is described in the next section, the smallest unit of story structure within a scene is the beat, a single action/reaction pair. The scene-level drama manager will thus need to continuously guide the autonomous decision making of the agent. This frequent guidance from the drama manager will be complicated by the fact that low-bandwidth guidance (such as giving a believable agent a new goal) will interact strongly with the moment-by-moment internal state of the agent, such as the set of currently active goals and behaviors, leading to surprising, and usually unwanted, behavior. In order to reliably guide an agent, the scene-level drama manager will have to engage in higher-bandwidth guidance involving the active manipulation of internal agent state (e.g. editing the currently active goal tree). Authoring strongly autonomous characters for story-worlds is not only extra, unneeded work (given that scene-level guidance will need to intervene frequently), but actively makes guidance more difficult, in that the drama manager will have to compensate for the internal decision-making processes (and associated state) of the agent.

Thinking of a believable agent as an autonomous, independent character leads to a style of agent authoring focusing on the goals, motivations, behaviors and

emotional states of the agent independent of their participation within a story context or their interactions with other agents. The internal structure of these agents is decoupled from consideration of how they will be guided by a drama manager. But, as mentioned above, any goal or behavior level guidance will strongly interact with the agent's internal decision making processes and state. Reliable guidance will be greatly facilitated by building hooks into the agents, that is, goals and behaviors that are specifically designed to be activated by the drama manager, and which have been carefully crafted so as to override the agent's autonomous behavior in an appropriate manner. But to the extent that authoring story-based believable agents requires special attention to guideability, this brings into question how useful it is to think of the believable agents as "autonomous" in the first place.

As the drama manager provides guidance, it will often be the case that the manager will need to carefully coordinate multiple characters so as to make the next story event happen. For example, it may be important for two characters to argue in such a way as to reveal specific information at a certain moment in the story. In a sense the real goal of these two characters is to conspire towards the revelation of a specific piece of information by arguing with each other. But an author who thinks of the characters as autonomous will tend to focus on the individual character goals, not story-level goals. To make a story-level goal happen, the character author will have to somehow coordinate the individual character goals and behaviors so that as the characters individually react to each other, the resulting interaction "just happens" to achieve the story goal. An alternative to this is to back away from the stance of strong autonomy and provide special goals and behaviors within the individual agents that the drama manager can activate to create coordinated behavior (a specific instance of providing hooks as described above). But even if the character author provides these special coordination hooks, coordination is still being handled at the individual goal and behavior level, in an ad-hoc way, on a case-by-case basis. What one really wants is a way to directly express coordinated character action at a level above the individual characters.

At this point the assumptions made by an interactive drama architecture consisting of a drama manager guiding strongly autonomous agents have been found problematic. The next section presents a sketch of a plot and character architecture that addresses these problems.

Integrating Plot and Character with the Dramatic Beat

In dramatic writing, stories are thought of as consisting of events that turn (change) values (McKee 1997). A value is a property of an individual or relationship, such as trust, love, hope (or hopelessness), etc. In fact, a story event is precisely any activity that turns a value. If there is activity

– characters running around, lots of witty dialog, buildings and bridges exploding, and so on – but this activity is not turning a value, then there is no story event, no dramatic action. Thus one of the primary goals of an interactive drama system should be to make sure that all activity turns values, and is thus a story event. Of course these values should be changed in such a way as to make some plot arc happen that enacts the story premise. The premise is the controlling idea of the story (McKee 1997), such as "Goodness triumphs when we outwit evil", or "To be happy you must be true to yourself".

Major value changes occur in each scene. Each scene is a large-scale story event (but in the case of our short one-act story, not necessarily as lengthy as a scene in a feature film or full-length play). In our story, an example of a scene would be "Grace confesses her love for the player". Scenes are composed of beats, the smallest unit of value change. Any activity below the level of the beat is not associated with value change. Roughly, a beat consists of an action/reaction pair between characters. For example, in the case where action is being carried by dialog, a beat could simply consist of one character speaking a line of dialog, and another character reacting. Generally speaking, in the interest of maintaining economy and intensity, a beat should not last longer than a few actions or lines of dialog.

Scenes and Beats as Architectural Entities

Given that the drama manager's primary goal is to make sure that activity in the story world is dramatic action, and thus turns values, it makes sense to have the drama manager use scenes and beats as architectural entities.

In computational terms, a scene consists of preconditions, a description of the value(s) intended to be changed by the scene (e.g. love between Grace and the player moves from low to high), a (potentially large) collection of beats with which to construct the scene, and a description of the arc that the value(s) changed by the scene should follow within the scene. The scene precondition tests whether the scene is appropriate given the current story and character state. The story state consists of the current story values and other global state such as active conversational topics, physical locations occupied by the characters, etc. To decide which scene to attempt to make happen next, the drama manager examines the list of unused scenes and chooses the one that has a satisfied precondition and whose value change best matches the shape of the global plot arc.

Once a scene has been selected, the drama manager tries to make the scene play out by selecting beats that change values appropriately. A beat consists of preconditions, a description of the values changed by the beat, success and failure conditions, and a joint plan to be executed by the characters. Like the preconditions on scenes, preconditions on beats also test story and character state for beat appropriateness. The success and failure conditions are tests that indicate when a beat has succeeded or failed and, for polymorphic beats, indicate

which specific beat should be considered to have occurred given how the beat was terminated (this will be described in more detail below). The joint plan coordinates the characters in order to carry out the specific beat.

The Function of Beats

Beats serve several functions within the architecture. First, beats are the smallest unit of dramatic value change. They are the fundamental building blocks of the interactive story. Second, beats are the fundamental unit of character guidance. The beat defines the granularity of plot/character interaction. Finally, the beat is the fundamental unit of player interaction. The beat is the smallest granularity at which the player can engage in meaningful (having meaning for the story) interaction. A player's activity is interpreted as having affected the story only to the extent that this activity participates in a beat.

Polymorphic Beats

The player's activity within a beat will often determine exactly which values are changed by a beat and by how much. For example, imagine that Trip becomes uncomfortable with the current conversation - perhaps at this moment in the story Grace is beginning to reveal problems in their relationship - and he tries to change the topic, perhaps by offering to get the player another drink. The combination of Grace's line of dialog (revealing a problem in their relationship), Trip's line of dialog (attempting to change the topic), and the player's response is a beat. Now if the player responds by accepting Trip's offer for a drink, the attempt to change the topic was successful, Trip may now feel a closer bond to the player, Grace may feel frustrated and angry with both Trip and the player, and the degree to which relationship problems have been revealed does not increase. We might label such a beat "Grace fails to discuss her marriage" or equivalently "Trip successfully changes topic away from marriage." On the other hand, if the player directly responds to Grace's line, either ignoring Trip, or perhaps chastising Trip for trivializing what Grace said, then the attempt to change the topic was unsuccessful, Trip's affiliation with the player may decrease and Grace's increase, and the degree to which relationship problems have been revealed increases. We might label this beat "Grace successfully brings up troubles with marriage." Before the player reacts to Grace and Trip, the drama manager does not know which beat will actually occur. This beat is a polymorphic beat. The drama manager selects this beat based on a range of effects that might occur. While the beat is executing, it is labeled "open." Once the player "closes" the beat by responding, the drama manager can now update the story history (a specific beat has now occurred) and the rest of the story state (dramatic values, etc.).

Joint Plans

Associated with each beat is a joint plan that guides the character behavior during that beat. Instead of directly initiating an existing goal or behavior within the character, the drama manager hands the characters new plans (behaviors) to be carried out during this beat. These plans are joint plans: they describe the coordinated activity required of all the characters in order to carry out the beat. As discussed in section 3, it is possible to write individual character behaviors that use ad-hoc communication (either in the form of sensing, or some form of direct, out-of-band message passing) to achieve multi-character coordination. It is difficult, however, for a behavior author to understand ahead of time all the synchronization problems that can occur; as unforeseen synchronization problems appear during play-testing, repeated patching and re-authoring of the behaviors will be necessary. In addition, the behavior author will have to separately solve the coordination problems of each new behavior involving multiple characters. However, multi-agent coordination frameworks such as joint intentions theory (Cohen and Levesque 1991) or shared plans (Grosz and Kraus 1996) provide a systematic analysis of all the synchronization issues that arise when agents jointly carry out plans. Tambe (Tambe 1997) has built an agent architecture providing direct support for joint plans. His architecture uses the more formal analyses of joint intentions and shared plans theory to provide the communication requirements for maintaining coordination; when a joint plan is being carried out, the architecture automatically takes care of all the necessary message passing. We propose modifying the reactive planning language Hap (Loyall and Bates 1991; Loyall 1997), a language specifically designed for the authoring of believable agents, to include this coordination framework.

Beats will hand the characters joint plans to carry out which have been designed to accomplish the beat. This means that most (perhaps all) of the high level goals and plans that drive a character will no longer be located within the character at all, but rather will be parceled out among the beats. Given that the purpose of character activity within a story world is to create dramatic action, this is an appropriate way of distributing the characters' behavior. The beat is precisely the smallest unit of dramatic action (the smallest unit that turns values). The character behavior is now organized around the dramatic functions that the behavior serves, rather than organized around a conception of the character independent of the dramatic action (a conception thus requiring the drama manager to coerce the character into serving the action). Since the joint plans associated with beats are still reactive plans, there is no loss of character reactivity to a rapidly changing environment. Low-level goals and behaviors (e.g. locomotion, ways to express emotion, personality moves, etc.) will still be contained within individual characters. These low-level behaviors provide a

library of character-specific actions that are available to the higher-level behaviors handed down by the beats.

A Response to the Problem of Autonomy

In the section "Autonomy and Story-based Believable Agents" we critiqued interactive drama architectures that consist of strongly autonomous characters guided by a drama manager. In this section we discuss how our proposed architecture addresses these issues.

In our architecture, the individual characters are no longer strongly autonomous. In the absence of the drama manager, the characters will not take action (or perhaps will only have very simple reactions to the environment). The beat level of the drama manager provides frequent guidance to the characters by giving them reactive joint plans to carry out. These frequent, beat-level decisions are made based on the global story state. Multiple characters are coordinated at the beat level; character authors are not forced to provide ad-hoc coordination within individual characters. Since the characters contain only low-level goals and behaviors, there is no complex character state complicating drama manager guidance. There is no longer a tension between authoring self-contained autonomous characters that have independent motivations, and providing those characters with the appropriate "hooks" to support guidance by an external process. Instead, the characters become libraries of character-specific ways of accomplishing low-level tasks; all higher-level motivation is provided by the drama manager. Thus this architecture addresses the tension between autonomy and dramatic guidance by backing away from strong autonomy on the part of characters and instead having dramatic guidance be responsible for most high-level character behavior.

Conclusion

In this paper we described the project goals of a new interactive drama project being undertaken by the authors. A major goal of this project is to integrate character and story into a complete dramatic world. We then explored the assumptions underlying architectures which propose that story worlds should consist of strongly autonomous believable agents guided by a drama manager, and found those assumptions problematic. Finally, we gave a brief sketch of our interactive drama architecture which addresses these problems. This architecture operationalizes structures found in the theory of dramatic writing, particularly the notions of changing dramatic values, and organizing dramatic value change around the scene and the beat.

Bibliography

Andre, E., Rist, T., Mueller, J. 1998. Integrating Reactive and Scripted Behaviors in a Life-Like Presentation Agent. *Proc. of the Second International Conference on Autonomous Agents* (Agents '98), pp. 261-268.

Bates, J. 1992. Virtual Reality, Art, and Entertainment. Presence: *The Journal of Teleoperators and Virtual Environments* 1(1): 133-138.

Bates, J., Loyall, A. B., and Reilly, W. S. 1992. Integrating Reactivity, Goals, and Emotion in a Broad Agent. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Bloomington, Indiana, July 1992.

Blumberg, B. 1996. Old Tricks, New Dogs: Ethology and Interactive Creatures. Ph.D. Dissertation. MIT Media Lab.

Blumberg, B. and Galyean, T. 1995. Multi-level Direction of Autonomous Creatures for Real-Time Virtual Environments. In *Proceedings of SIGGRAPH 95*.

Cohen, P. and Levesque, H. 1991. Teamwork. *Nous*, 35.

Grosz, B. and Kraus, S. 1996. Collaborative plans for complex group actions. *Artificial Intelligence*, 86, 269 - 358.

Hayes-Roth, B., van Gent, R. and Huber, D. 1997. Acting in character. In R. Trappl and P. Petta (Eds.), *Creating Personalities for Synthetic Actors*. Berlin, New York: Springer.

Lester, J., Stone, B. 1997. Increasing Believability in Animated Pedagogical Agents. *Proceedings of the First International Conference on Autonomous Agents*. Marina del Rey, CA, USA, 16-21.

Loyall, A. B. 1997. Believable Agents. Ph.D. thesis, Tech report CMU-CS-97-123, Carnegie Mellon University.

Loyall, A. B., and Bates, J. 1991. Hap: A Reactive, Adaptive Architecture for Agents. Technical Report CMU-CS-91-147. Department of Computer Science. Carnegie Mellon University.

Mateas, M. 1999. An Oz-Centric Review of Interactive Drama and Believable Agents. In M. Wooldridge and M. Veloso, (Eds.), *AI Today: Recent Trends and Developments*. Lecture Notes in AI 1600. Berlin, New York: Springer.

McKee, R. 1997. *Story: Substance, Structure, Style, and the Principles of Screenwriting*. New York, NY: HarperCollins.

Murray, Janet. 1998. *Hamlet on the Holodeck*. Cambridge, MA: MIT Press.

Stern, A.; Frank, A.; and Resner, B. 1998. Virtual Petz: A hybrid approach to creating autonomous, lifelike Dogz and Catz. In *Proceedings of the Second International Conference on Autonomous Agents*, 334-335. Menlo Park, Calif.: AAAI Press

Tambe, M. 1997. Towards Flexible Teamwork. *Journal of Artificial Intelligence Research* (7) 83-124.

Weyhrauch, P. 1997. Guiding Interactive Drama. Ph.D. thesis, Tech report CMU-CS-97-109, Carnegie Mellon University.