Failing Believably: Toward Drama Management with Autonomous Actors in Interactive Narratives

Mark O. Riedl[1], Andrew Stern[2]

[1] University of Southern California, Institute for Creative Technologies, 13274 Fiji Way, Marina del Rey, California, USA
[2] Procedural Arts LLC, Portland, Oregon, USA
riedl AT ict.usc.edu, andrew AT proceduralarts.com

Abstract. Interactive Narrative is an approach to interactive entertainment that enables the player to make decisions that directly affect the direction and/or outcome of the narrative experience being delivered by the computer system. One common interactive narrative technique is to use a drama manager to achieve a specific narrative experience. To achieve character believability, autonomous character agents can be used in conjunction with drama management. In this paper, we describe the problem of *failing believably* in which character believability and drama management come into conflict and character agents must intelligently produce behaviors that explain away schizophrenic behavior. We describe technologies for implementing semi-autonomous believable agents that can fail believably.

1 Introduction

Storytelling in one form or another is a central component in a vast majority of modern computer games. *Interactive Narrative* is an approach to interactive entertainment that enables the player to make decisions that directly affect the direction and/or outcome of the narrative experience being delivered by the computer system. Often this means the player assumes the role of a story world character, possibly the protagonist. Story and character are intertwined; story cannot exist without character, and character alone often implies narrative. Consequently, computational storytelling researchers have long placed emphasis on creating believable agents (c.f. [2], [13], [5], [11], and [8]). A believable agent [2; 13] is an autonomous agent that enacts the role of a story world character by manifesting through its behavior personality, emotion, beliefs, desires, intentionality, and a host of other traits that lead to the "illusion of life." Believability is an ill-defined concept. Within the context of the paper, we shall define believability as not performing actions that break the user's suspension of disbelief that a character is not an automaton.

Autonomous embodied agents are commonly used for achieving believable characters in interactive systems because autonomy affords reactivity, interactivity, and local grounding of agent decisions in observable behavior. Aylett [1] notes that interactive narratives that rely on autonomous agents alone do not always produce narrative experiences with recognizable structure or coherent, meaningful outcomes. Interactive narrative researchers often turn to the concept of drama management to ensure coherent, recognizable narrative experiences in their systems (c.f. [18], [14], [24], [22], [16], [21], and [8]). A drama manager [12] is an intelligent, disembodied agent that manages the virtual world in order to bring about a narrative experience for the user. Typically, this is achieved is by providing the drama manager with some sort of pre-specified branching or non-branching storyline.

In our research, we adopt the use of both autonomous character agents and drama management. We believe the careful combination of these technologies affords a richer, more interactive narrative experience. Believable agents, through dialogue and reactivity, help immerse the user and invoke the sense of user self-agency. The drama manager provides a coherent narrative structure to the user's experience in the virtual world. We adopt a *generative* technique of drama management pioneered by [18; 25]. The generative approach affords a drama manager more flexibility to accommodate the user by dynamically re-generating narrative content to incorporate user actions into the plot representation. That is, the drama manager adapts its narrative automatically – changing the direction, in terms of future plot points, and/or outcome – to the actions of the user. The trade-off for using this approach is that narrative content cannot be known a priori, the consequence of which is that special attention must be paid to the coordination between the drama manager and believable agents. Specifically, one cannot design believable agents to avoid performing actions that conflict with the drama manager's storytelling goals because the narrative structure is not known at design time.

In this paper, we propose an approach to coordination between drama managers and autonomous believable agents that avoids or mitigates the effects of the inherent tension between believable agents and the drama manager. In the following sections, we describe circumstances where coordination between believable agents and a drama manager can result in character agent abandoning or abruptly changing its goals or inexplicably contradicting its previous actions, all of which may be perceived by a user as schizophrenic [20] behavior. Traditionally, agent goal failure is a topic addressed in artificial intelligence literature by techniques for avoiding failure or by re-planning when failure occurs. But agent goal failure that can be explained in the context of the narrative and virtual world may actually be desirable; for example, [5] uses goal failure to create comedic and/or dramatic situations. However, *inexplicable* goal failure results in a loss of agent believability. Explainability of the behavior of an agent is an integral component in user comprehension of autonomous agent performance [23].

The problem we are addressing is the realization that sometimes a believable agent under drama management is required to fail to achieve a goal, in the interest of preserving the drama manager's plotline. We use the term *failing believably* to refer to a technique whereby a believable agent can given an a priori unknown plotline (a) intelligently and believably avoid situations where the agent and drama manager come into conflict, and (b) enact a plausible and believable explanation for goal failure when conflict with the drama manager is unavoidable.

2 A Generative Drama Manager for Directing Autonomous Characters

In our research, we adopt a generative technique of drama management pioneered by [18; 25]. The generative approach allows a drama manager to accommodate user actions by incorporating user actions into the plot representation. When the user – knowingly or unknowingly – performs actions that change the story world in a way that makes it impossible or nonsensical to achieve future plot points, the drama manager adapts or repairs the plotline by re-generating part or all of the remaining plot structure. One restriction we enforce is that computer-controlled characters are disallowed from performing actions that cause the drama manager to adapt or repair the plot. This restriction is based on the assumption that the original plot is considered ideal and that every alteration to the plot degrades its effectiveness [18].

Scenario adaptation requires that the drama manager's plot be represented computationally so it can be reasoned over by an artificial intelligence system. Following [18], we represent plots as partially-ordered plans. A plan contains steps – events that change the state of the world – and annotations that explicitly mark the temporal and causal relationships between all steps in the plan, defining a partial order indicating the steps' order of execution. Causal annotations, called *causal links*, specifically are used to denote causal relationships between the steps in the plan. A causal link relates the effect of one plan step to a precondition of another plan step that is temporally constrained to occur later than the first step. Fig. 1 shows part of a plot plan taken from [19], used to train situation awareness to military leaders. Plan steps (boxes) represent a sequence of partially ordered plot points to be achieved in the story world if no branching occurs. The arcs represent causal links denoting conditions in the virtual world that are established by the successful achievement of the temporally earlier plot point and required for successful achievement of the temporally later plot point. For example, many plot points require that certain characters are not detained and not incapacitated.

The scenario is expected to progress as follows. Two characters, Hassan and Saleh, are merchants and sectarian rivals. Hassan is to act friendly towards the player while Saleh is to act hostilely towards the player. Later, Hassan acquires a bomb and plants it. However, Hassan is observed planting the bomb by a third character, Ali. The bomb goes off but it is a dud and no one gets hurt. Once the event takes place, Hassan falsely accuses Saleh of being the villain while Ali righteously accuses Hassan. The objective of the scenario is to challenge the user to develop more sophisticated situational understanding in order to peacefully resolve conflicts.

If the user performs an action in the virtual world that changes the world state in such a way that a causal link condition is negated, the drama manager re-generates the plot plan to reestablish the negated world condition or to remove jeopardized future plot points and find a new plot sequence. For example, suppose the trainee finds the planted bomb disarms it. Were this to occur the causal link specifying that the bomb must be armed before it can go off (as a dud) is negated. The drama manager, detecting the threat to the plot, re-generates the remaining plot. One simple adaptation is to have another character plant a second bomb, although the current state of the drama manager is capable of significantly more complex adaptations (see [19]

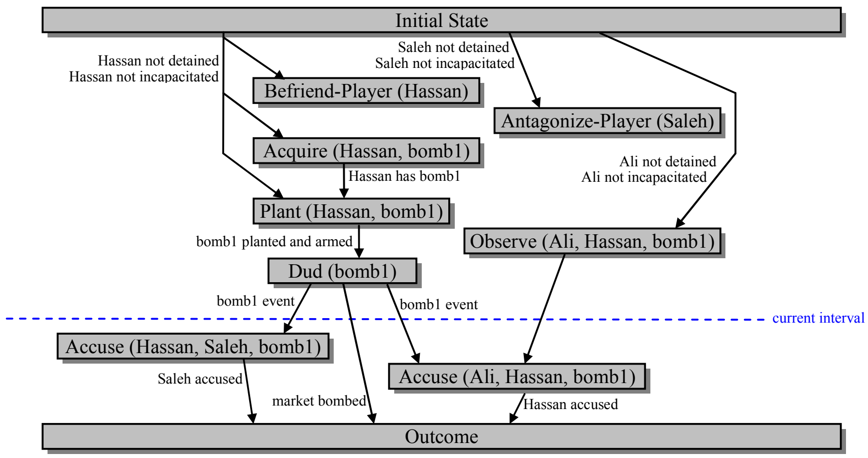

Fig. 1. A causally connected graph of plot points.

for additional plot adaptations). Only actions that threaten causal links require adaptation of the plot. All other user actions are considered consistent with the plot as currently structured. Interactions with autonomous character agents that do not threaten the plot account for most of the variability in the user's experience.

By representing narrative as a plan, the narrative can be regenerated using specialized narrative planners such as [17]. This approach has been successfully applied to entertainment and educational applications (e.g. [18; 25]) and to interactive narrative-based training (e.g. [19]). For more information on the drama manager and the application, see [19].

2.1 Directing Autonomous Character Agents

Drama management techniques for interactive storytelling systems are often used to direct the behaviors of character agents so that they do the things that are required to drive a plot forward. But one aspect of drama management that has been largely overlooked is directing character agents to prevent undesirable world states. The issue at hand is that character agents with any degree of autonomy may make decisions to perform actions that change the state of the virtual world in such a way that the drama manager's intentions for future plot become hampered. A character agent acting autonomously may choose to perform an action that makes future plot progression impossible or nonsensical. For example, a character agent may destroy or make inaccessible an essential resource or in some other way change the state of the world so that it is incompatible with the plot sequence.

Drama managers do not always have to act proscriptively against autonomous character agents. In game-like systems where the drama manager's linear or branching storyline is known in advance by system designers, autonomous agents can be designed to perform behaviors that are guaranteed never to cause world state changes that will violate the drama manager's plotline. However, if the drama manager's story is not known before hand, can be dynamically adapted (c.f. [19]), or is generated automatically by the system on a per-session basis (c.f. [18; 25]), then it is impossible to design autonomous believable agents that are guaranteed never to conflict with the drama manager's storytelling goals.

Blumberg and Galyean [4] describe a way of controlling the behaviors of autonomous virtual agents through weighted preferences. Not only can control of agents be *prescriptive* – "do X" – but also *proscriptive* – "I don't care what you do as long as you don't do Y." One system, Virtual Storyteller, [22] produces stories using autonomous agents and a drama manager that applies prescriptive and proscriptive techniques; each autonomous character agent must ask permission of the drama manager before performing an action. Likewise, Szilas [21] proposes a model where the drama manager participates in the agent's intelligent action selection process.

Proscriptive control of autonomous believable agents presents the following challenge: How does one proscribe a believable agent without adversely impacting that character believability? At the heart of the issue is that an autonomous believable agent – designed to make local decisions in order to appear believable – can determine at any given moment that the most believable action to perform is one that will change the world in a way that negatively impacts the drama manager's ability to progress through the desired plot structure.

For example, in the context of the narrative plan shown in Fig. 1, suppose that Hassan and Saleh are bitter rivals and Saleh forms the goal to incapacitate Hassan. This may be seen as reasonable given the relationship between the two characters. Unlike the user, character agents are prohibited from performing actions that negate causal links – in this case the condition that Hassan is not incapacitated. Causal links extend over a period of time. If Saleh acts during an interval in which the condition that Hassan not be incapacitated is proscribed, then Saleh should not be able to execute any plan to incapacitate Hassan and may be forced to abandon the goal immediately. Suppose however, that the character agents have been operating for some amount of time before the drama manager begins managing the world and that Saleh has formed the goal before the drama manager proscribes against the condition. Further suppose that the Saleh agent has already gone through the motions of picking up a weapon and heading over to Hassan. When the drama manager proscribes against incapacitating Hassan, the Saleh agent will be forced to abandon the goal, in effect *failing*. The concept of failing believably is to append the aborted sequence of actions – the menacing of Hassan – with a continuing sequence of actions that is not inconsistent with past actions and also justifies to the user why Saleh did not attempt to complete the goal. In the simplest form, failing believably could mean (possibly pretending) to be called away on a more important issue. A more sophisticated variation on failing believably could involve a joint performance (for a discussion of joint behaviors in interactive drama systems, see [15; 16]) in which Hassan pleads with and convinces Saleh to leave him unharmed.

2.2 Generating Directives

Plot points and their causal relationships are used to generate directives to believable, autonomous agents. Directives take the following forms:

- *Prescriptive directive*: Direction to achieve some world state that is desirable to the drama manager and moves the plot forward. Prescriptive directives should be of sufficiently high level that character agents are free to achieve the directive in a way that best suites the believability of that character.
- *Proscriptive directive*: Direction to avoid particular states in the virtual world that are contradictory to the intermediate requirements of the drama manager's plot representation.

The drama manager must derive directives from its internal plot representation. Prescriptive directives are derived directly from plot points. That is, a plot point specifies a world state that must be achieved. The drama manager determines which story character is responsible for achieving this world state and formats the plot point as a goal for the character agent to achieve. It is often the case that the character responsible for a plot point is encoded directly into the plot point as the entity that is changed or creates the change.

Proscriptive directives are more problematic because they do not correspond to any particular plot point. Instead, proscriptive directives are derived from the causal relationships between plot points. Specifically, for every plot point in the drama manager's representation of expected user experience, there are certain world state conditions that must be maintained. These conditions are established by plot points that occur earlier in the plot structure or by the initial world state. If an autonomous agent were to perform an action that negates such a condition, future plot points would become impossible or, at the very least, cease to make sense.

Fig. 1 shows a hypothetical plot plan. The dashed line delineates a point in time in the plot plan where plot points above have been achieved and plot points below have yet to be achieved. The causal links that cross the current interval line are the conditions that must be maintained for the plot to continue. In the example, character agents will be directed to avoid actions that negate the truth values of conditions *(planted bomb1)* and *(armed bomb1)*. This is especially important as the agent playing the role of Ali, having seen the bomb being planted, may believably want to choose to do something about it. As character agents receive positive directives, achieve plot points, and report back successes, the current interval shifts downward and the set of world state conditions that must be maintained changes, further constraining or freeing up the actions available to character agents.

3 Believable, Directable Semi-Autonomous Character Agents

Agents that must operate in dynamic environments often rely on robust teleo-reactive [3] plan execution systems such as RAPS [9], PRS [10], Soar (when not used to model psychomotor skills), and ABL [15] (used by the *Façade* [16] interactive drama) that take advantage of hierarchical representations of behavior. The advantage of these systems is that they can readily re-plan behaviors when circumstances change by failing the currently operating plan (also referred to in agent literature as methods, productions, task networks, or behaviors) and choosing a different decomposition of an abstract high-level goal.

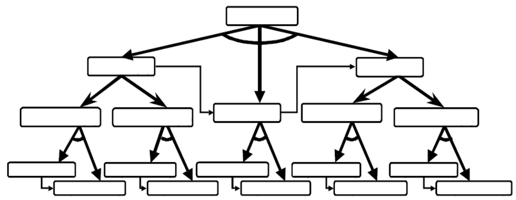

Fig. 2. Hierarchical AND/OR plan tree.

One can represent the space of possible agent behaviors of a hierarchically-based plan execution system as an AND/OR graph [5]. Nodes are plans that, unless they are leaf nodes, have one or more child nodes. If a node is an AND-node, then the child nodes are ordered and must all be achieved in the specified order for the parent node to succeed. If a node is an OR-node, then the child nodes represent alternative methods for achieving the parent node and only one must be chosen. Preconditions on the children of OR-nodes determine which alternative methods of achieving the parent plan are legal given the current executing environment. Fig. 2 shows a conceptual hierarchical arrangement of goals. The boxes denote goals at different degrees of abstraction. Downward arrows show how plans are decomposed. If the arrows are joined by an arc, then that node is an AND-node. Otherwise, it is an OR-node. Horizontal arrows denote ordering constraints among children of an AND-node.

Agent behavior is determined in real-time by taking the current executing abstract plan and decomposing it. If the plan is represented by an OR-node, decomposition is a process of selecting one of the alternative sub-plans. When a chosen node is a leaf, a corresponding primitive action is executed by the agent in the world. Because of OR-nodes, it is impossible in many circumstances to predict what primitive actions an agent will perform from moment to moment. Key to believable agents in general, and those that must fail believably in specific, is for agents to have a lot of alternative OR-nodes to select from every time in needs to decompose a top-level goal into primitive actions.

3.1 Handling Prescription: Mixing Autonomy and Narrative Direction

Because of our high-level approach to drama management plot representation, positive directives are of the form of goals for a character agent to bring about a particular world state. Prescriptive directives may come infrequently, however. Character agents must fill up the duration between prescriptive directives by autonomously choosing goals and acting to achieve them. When a character agent receives a prescriptive directive from the drama manager, the agent must incorporate that new goal into it's reasoning in a way that does not violate believability, e.g. avoid schizophrenically [20] switching from one goal to another seemingly unrelated goal.

There are many intelligent design tricks that can be used to seamlessly integrate autonomous goals and directed goals. Our character agents use tables of resource requirements to determine whether a currently pursued autonomous behavior is compatible with a direction from the drama manager. These tables are compiled at agent design time. From this information and prioritization information, a character agent can determine if it can:

- Merge its current autonomous goal and the goal of the prescriptive directive, achieving both simultaneously through clever selection of decompositions;
- Complete the current autonomous goal and then switch to the task of achieving the prescriptive directive; or
- Abort or suspend the autonomous behavior to pursue the goal of the prescriptive direction.

In the last case, to avoid the appearance of the agent schizophrenically abandoning a goal without reason, we provide our character agents with transition behaviors. A transition behavior is a sequence of believable primitive actions that an agent can perform to justify or explain why it would stop doing one thing and start doing something else. This satisfies the requirement of explainable behavior.

For example suppose agent Hassan is performing the role of a merchant, enacting the types of behaviors one would expect a merchant to perform – autonomously interacting with customers and keeping shop – when it receives a prescriptive direction to acquire a bomb. The agent aborts its current behavior with a transition behavior in which the agent (possibly pretends to or not) receive an urgent phone call and closes the shop. The agent then begins executing actions to complete the directive.

3.2 Handling Proscription: Plan Coordination with Restoration Modes

The logic for avoiding certain world states is potentially more complicated than achieving particular world states; avoidance of world states requires autonomous agents to be aware of the implications and side effects of their actions. For an autonomous agent without the ability to perform sophisticated look-ahead it also means that the agent might encounter situations where it needs to adjust its course of action and do so without creating the appearance of schizophrenically changing goals. However, the standard agent architecture does not lend itself to sophisticated look-ahead because at any given moment, a teleo-reactive agent cannot predict which OR-decomposition the agent might pick when decomposing a behavior plan.

Handling proscriptive directives involves addressing two issues:

1. Given alternative plan decompositions, an agent must decompose plans in such a way that it never performs a primitive action that has effects that conflict with the drama manager's desired plot; and
2. Failing believably in the case that and agent's performing of a primitive action that conflicts with the drama manager's desired plot is imminent and inevitable.

The solution we have adopted is to treat the problem as one of plan coordination. That is, the drama manager has a plan for how the user's experience should progress, and the character agents have reactive plans for how to behave believably in the world. Because of the particular arrangement between the drama manager and the character agents, the drama manager's plan takes precedence and the agents must adapt their behavior selection to coordinate with the drama manager's plot plan. This must be done without sacrificing believability. That is, character agents must reactively plan without appearing schizophrenic. The next two sections respectively describe (a) how techniques for plan coordination can address the first issue, and (b) how the plan coordination algorithms can be adapted to support believability, addressing the second issue.

Plan Coordination. There is an existing body of research on the coordination of hierarchical plans of multiple agents. We base our approach on the plan synergy algorithm [6; 7] because it directly supports multiple teleo-reactive agents with hierarchical plans expressed as AND/OR trees of goals.

In the plan synergy approach [6; 7], the plans of the agents at run-time must be coordinated and that coordination should occur as close to the leaves of the plan trees as possible for the crispest coordination possible. The approach assumes that nodes are plans that achieve goals and that they have pre-conditions, post-conditions, and in-conditions. Pre-conditions are conditions that must hold before an agent can execute a plan. Post-conditions are conditions that must hold after an agent completes the execution of a plan. In-conditions are conditions that will be true at some point during the execution of a plan. The synergy approach uses a two phase process. The first phase is an offline process that preprocesses the AND/OR trees of plans for each agents. Each node is annotated with summary information that describe the pre-conditions, post-conditions, and in-conditions of that node plus the pre-, post-, and in-conditions of all sub-plans. Conditions are marked *must* or *may*. A *must* annotation denotes that a pre/post/in-condition will definitely be true during the execution of a plan and all sub-plans. A *may* annotation denotes that a pre/post/in-condition could be true during the execution of a plan and some sub-plans. The preprocessing algorithm is described in [6]. The second phase uses the annotated AND/OR trees of multiple agents to constrain each agent so that they avoid acting in a conflicting manner.

In our work on integrating reactive believable agents and high-level plot plans, we use the synergy preprocessing phase to annotate agent behavior libraries with *must* and *may* condition annotations. We discard all but *must* condition annotations because in our case the drama manager's plot plan is assumed to be static and character agents are constrained to enact behaviors that do not conflict with the plot plan. If a plan is annotated with a *may* condition, then the agent can find an OR sub-plan that does not cause/require that condition to become true. We use the *must* annotations to restrict certain sub-plans from executing if there is a proscriptive directive from the drama manager and that sub-plan is annotated to indicate that it conflicts with the directive. That is, if a proscriptive directive matches a *must* condition on a sub-plan, the sub-plan cannot be chosen for decomposition and the agent must select an alternative decomposition. Character agents trying to achieve a particular top-level goal will only pick methods for achieving that goal that do not conflict with any negative directive. If all methods of achieving a particular top-level goal conflict with a negative directive, the character agent will be forced to choose a different but equally plausible and believable top-level goal before execution begins.

It should be noted that many agent behavior specification languages such as ABL [15] do not encode information about post-conditions or in-conditions. To compensate, we maintain a dual STRIPS-like representation of all ABL behaviors that are part of a plan to achieve a top-level goal. We run the annotation process on this alternative representation and merge the *must* condition annotations derived from this process with agent behavior libraries to produce new behavior libraries that support failing believably. The input to this process is a vocabulary of predicates used by the drama manager to represent the state of the story world. The new behavior libraries contain modified behaviors with additional preconditions and context conditions that cause any particular behavior to (a) not be able to execute or (b) immediately fail if certain negative directives are being enforced by the drama manager.

Restoration Modes. There exists the possibility that an agent has already begun a behavior plan that already has or will cause some world state condition to become true in direct contradiction of the drama manager's plot plan. This situation can happen under two circumstances. First, the condition is not prohibited by the Story Director at the time the agent selects the behavior plan because the drama manager had not yet transitioned into an interval of the plot plan in which a particular condition must be maintained. Second, a user action forces the drama manager to branch to an alternative plot plan that has different protected conditions. We do not use look ahead to prevent character agents from performing actions that will conflict with future directives because agents will have to look down all possible branches of the plot regardless of plausibility. In either event, the agent can be executing a plan that has caused or will inevitably conflict with the drama manager's plot plan.

When a top-level goal is forced to fail after execution has begun because the character agent cannot continue to pursue that goal without conflicting with the drama manager's plot plan, we encounter the potential for character believability to falter. That is, a character has suddenly stopped pursuing one goal and has switched to another goal without apparent reason. To avoid the appearance of such schizophrenic behavior we provide our character agents with special goal handler daemons that note when a goal fails. If the goal fails because of a proscriptive directive, the goal hander daemon triggers a *restoration mode*. A restoration mode is a hierarchically organized behavior that is ordinarily inaccessible to the character agent and provides an enacted behavior that visibly explains why the agent would cease one plan and adopt a new plan. If the agent has already negated the condition at the time that the proscriptive directive is given, the failure mode also re-establishes the condition.

For example, suppose there is a proscription against any characters besides Hassan and Ali to be in a particular area when Hassan plants the bomb. If Saleh is already in the area, the agent must produce actions that explain why it can no longer be in the area. At a minimum, the agent may suddenly remember that he forgot something, or coordinate with another agent that asks it to run an errand. To invoke less of a sense of *deus ex machina*, the agent deliberates and generates a plausible explanation for why he'd suddenly prefer to avoid this region or do something else somewhere else; this may be related to some specific personality quirk of the agent.

It should be noted that restoration modes are not immune from further proscriptive directives (the exception is the negative directive for which the restoration mode was invoked in the first place; it can disregard that particular directive since the mode exists to protect and restore the world in accordance with that directive). When the restoration mode is complete, the story world will be restored to a state that conforms to the drama manager's plot plan. Thus restoration modes are necessarily more sophisticated than the transition behaviors used to handle prescriptive directives. When the restoration mode completes, the character agent activates its goal arbitration process to autonomously choose a new top-level goal.

4 Conclusions

Failing believably is a term we use to describe the notion that autonomous agents should be designed to believably avoid inexplicable goal failures caused by the complex coordination of agents and a proscriptive drama manager. We present techniques for endowing actor agents with the ability to fail believably, that is, to generate believable behavior to explain to an observing player a reason why an agent must not proceed with, abort, or undo previous behavior plans. Our approach provides a framework in which teleo-reactive [3], believable agents can invoke believability-preserving behaviors in the face of goal failure. Although the agent framework automates the process, creative design of behaviors, transitions, and restoration modes is still in the hands of agent designers.

Failing believably becomes important when proscription by a drama manager can have an impact on agent believability. A drama manager must be able to proscribe agents from taking particular actions that make it impossible for the drama manager to drive the narrative forward. For systems such as [25] and [19] in which the desired narrative experience is a priori unknown, or can be adapted over time, or generated from scratch, it is impossible to design autonomous believable characters that can be guaranteed never to choose behaviors that will not conflict with the drama manager's global narrative plot specification. It is inherent to such systems that believable character goals may produce conflict with a drama manager's storytelling goals. However, both are important. Drama management affords a coherent narrative structure to the user's experience that can be adapted to the user's actions. The use of autonomous agents affords interactivity and believability and helps immerse the user in a world where he or she has the perception of minute-to-minute self-agency.

Acknowledgements

The project or effort described here has been sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

References

[1] Aylett, R.: Emergent Narrative, Social Immersion and "Storification." *Proc. of the 1st Int. Workshop on Narrative and Interactive Learning Environments* (2000).
[2] Bates, J.: Virtual Reality, Art, and Entertainment. *Presence: The Journal of Teleoperators and Virtual Environments*, 1 (1992), 133-138.
[3] Benson, S. Nilsson, N.: Reacting, Planning, and Learning in an Autonomous Agent. In Furukawa, Michie, Muggleton (eds.): *Machine Intelligence*, 14. Clarendon (1995).
[4] Blumberg, B., Galyean, T.: Multi-Level Control for Animated Autonomous Agents: Do the Right Thing… Oh, not that… In: R. Trappl, P. Petta (eds.): *Creating Personalities for Synthetic Actors*. Springer (1997).

[5] Cavazza, M., Charles, F., Meade, S.J.: Planning Characters' Behaviour in Interactive Storytelling. *Journal of Visualization and Computer Animation*, 13 (2002), 121-131.
[6] Clement, B.J., Durfee, E.H.: Top-Down Search for Coordinating the Hierarchical Plans of Multiple Agents. *Proc. of the 3rd Int. Conf. on Autonomous Agents* (1999).
[7] Cox, J.S., Durfee, E.H.: Discovering and Exploiting Synergy between Hierarchical Planning Agents. *Proc. of the 2nd Int. Joint Conf. on Autonomous Agents and Multi Agent Systems* (2003).
[8] Figueiredo, R., Dias, J., Aylett, R., Paiva, A.: Shaping Emergent Narratives for a Pedagogical Application. *Proc. of the 4th Int. Conf. on Narrative and Interactive Learning Environments* (2006).
[9] Firby, J.: *Adaptive Execution in Complex Dynamic Domains*. Ph.D. Dissertation, Yale University (1989).
[10] Georgeff, M, Lansky, A.: Procedural Knowledge. *Proc. of IEEE Special Issue on Knowledge Representation*, 74 (1986), 1383-1398.
[11] Gratch, J., Marsella, S.: Tears and Fears: Modeling Emotions and Emotional Behaviors in Synthetic Agents. *Proc. of the 5th Int. Conf. on Autonomous Agents* (2001).
[12] Kelso, M., Weyhrauch, P., Bates, J.: Dramatic Presence. *Presence: The Journal of Teleoperators and Virtual Environments*, 2 (2005).
[13] Loyall, A.B.: *Believable Agents: Building Interactive Personalities*. Ph.D. Dissertation, School of Computer Science, Carnegie Mellon University (1997).
[14] Magerko, B., Laird, J.E., Assanie, M., Kerfoot, A., Stokes, D.: AI characters and directors for intereactive computer games. *Proc. of the 16th Innovative Applications of Artificial Intelligence Conference* (2004).
[15] Mateas, M., Stern, A.: A Behavior Language: Joint Action and Behavior Idioms. In: H. Prendinger & M. Ishizuka (eds.): *Life-like Characters: Tools, Affective Functions and Applications*. Springer (2004).
[16] Mateas, M., Stern, A.: Structuring Content in the Façade Interactive Drama Architecture. *Proc. of the 1st Conf. on AI and Interactive Digital Entertainment* (2005).
[17] Riedl, M.O. *Narrative Generation: Balancing Plot and Character*. Ph.D. Dissertation, North Carolina State University.
[18] Riedl, M.O, Saretto, C.J., Young, R.M.: Managing Interaction Between Users and Agents in a Multi-Agent Storytelling Environment. *Proc. of the 2nd Int. Conf. on Autonomous Agents and Multi Agent Systems* (2003).
[19] Riedl, M.O., Stern, A.: Believable Agents and Intelligent Story Adaptation for Interactive Storytelling. *Proc. of the 3rd Int. Conf. on Technologies for Interactive Digital Storytelling and Entertainment* (2006).
[20] Sengers, P.: Schizophrenia and Narrative in Artificial Agents. In: M. Mateas and P. Sengers (eds.): *Narrative Intelligence*. John Benjamins (2003).
[21] Szilas, N.: The Future of Interactive Drama. *Proc. of the 2nd Australasian Conf. on Interactive Entertainment* (2005).
[22] Theune, M., Faas, S., Nijholt, A., Heylen, D.: The Virtual Storyteller: Story Creation by Intelligent Agents. *Proc. of the 1st Int. Conf. for Interactive Digital Storytelling and Entertainment* (2003).
[23] van Lent, M., Fisher, W., Mancuso, M.: An Explainable Artificial Intelligence System for Small-Unit Tactical Behavior. *Proc. of the 16th Innovative Applications of Artificial Intelligence Conference* (2004).
[24] Weyhrauch, P.: *Guiding Interactive Fiction*. Ph.D. Dissertation, School of Computer Science, Carnegie Mellon University (1997).
[25] Young, R.M., Riedl, M.O., Branly, M., Jhala, A., Martin, R.J., & Saretto, C.J.: An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development*, 1 (2004), 51-70.