

An Oz-Centric Review of Interactive Drama and Believable Agents

[Michael Mateas](#)

Abstract

Believable agents are autonomous agents that exhibit rich personalities. Interactive dramas take place in virtual worlds inhabited by characters (believable agents) with whom an audience interacts. In the course of this interaction, the audience experiences a story (lives a plot arc). This report presents the research philosophy behind the Oz Project, a research group at CMU that has spent the last ten years studying believable agents and interactive drama. The report then surveys current work from an Oz perspective.

Contents

- [Introduction](#)
- [Drama = Character + Story + Presentation](#)
- [Structure of the report](#)
- [Believable agents](#)
 - [The Oz philosophy](#)
 - [Taking character seriously](#)
 - [Classical vs. behavioral AI](#)
 - [Research goals: believable agents vs. AI](#)
 - [Authoring](#)
 - [Oh, what a tangled web we weave...](#)
 - [Related research areas](#)
 - [Artificial life](#)
 - [Humanoids](#)
 - [Embodied characters](#)
 - [Chatterbots](#)
 - [Behavioral animation](#)
 - [Why study believable agents?](#)
- [Interactive story](#)
 - [Interactive story: an oxymoron](#)
 - [Oz drama manager](#)
 - [Design dimensions](#)
 - [Local vs. global control](#)
 - [Granularity of story control](#)
 - [Generation](#)
- [Clip-based work](#)
- [Current explorations](#)
- [Conclusion](#)
- [Bibliography](#)

Introduction

This report provides an overview of research in believable agents and interactive drama. Many of the original sources used in compiling this report can be found on the web; the annotated bibliography provides links to these web resources.

This report unabashedly surveys its topic with the bias of the [Oz project](#) at CMU. The reason for this is threefold. First, I am a member of this research group and have internalized much of their perspective; I won't pretend not to have a viewpoint. Second, there is not much work that is directly related to interactive drama and believable agents; using the Oz project as a center allows me to make sense of more peripheral work. Finally, Oz is the only group giving equal attention to both character (believable agents) and story (interactive drama); an Oz perspective allows me to present character and story as a unified whole.

For much of the content of this report, I am indebted to the members of the Oz project: Joe Bates, Bryan Loyall, Scott Neal Reilly, Phoebe Sengers, and Peter Weyhrauch. Much of my understanding grew out of conversations with them. In particular, the Oz philosophy, which is described below, was developed by the group during 10 years of research.

The first item of business is to define the research goal for believable agents and interactive drama: building worlds with character and story.

Drama = Character + Story + Presentation

Artists building non-interactive dramas (e.g. movies, books) have commented on the importance of both character *and* story for authoring powerful, dramatic experiences. For example, Lajos Egri, in the [Art of Dramatic Writing](#), has this to say about premise (i.e. plot or story).

"No idea, and no situation, was ever strong enough to carry you through to its logical conclusion without a clear-cut premise.

If you have no such premise, you may modify, elaborate, vary your original idea or situation, or even lead yourself into another situation, but you will not know where you are going. You will flounder, rack your brain to invent further situations to round out your play. You may find these situations - and you will still be without a play."

Later, in talking about character, he defines three dimensions every character must have: physiology, sociology and psychology. He has this to say about these three dimensions:

"Analyze any work of art which has withstood the ravages of time, and you will find that it has lived, and will live, because it possesses the three dimensions. Leave out one of the three, and although your plot may be exciting and you may make a fortune, your play will still not be a literary success.

When you read drama criticisms in your daily papers, you encounter certain terminology time and again: dull, unconvincing, stock characters (badly drawn, that is), familiar situations, boring. They all refer to one flaw - the lack of tridimensional characters."

Figure 1, below, shows the high level architecture of the Oz project. This architecture arose out of a desire to treat character and story in interactive drama as *seriously* as do dramatic

artists in traditional media.

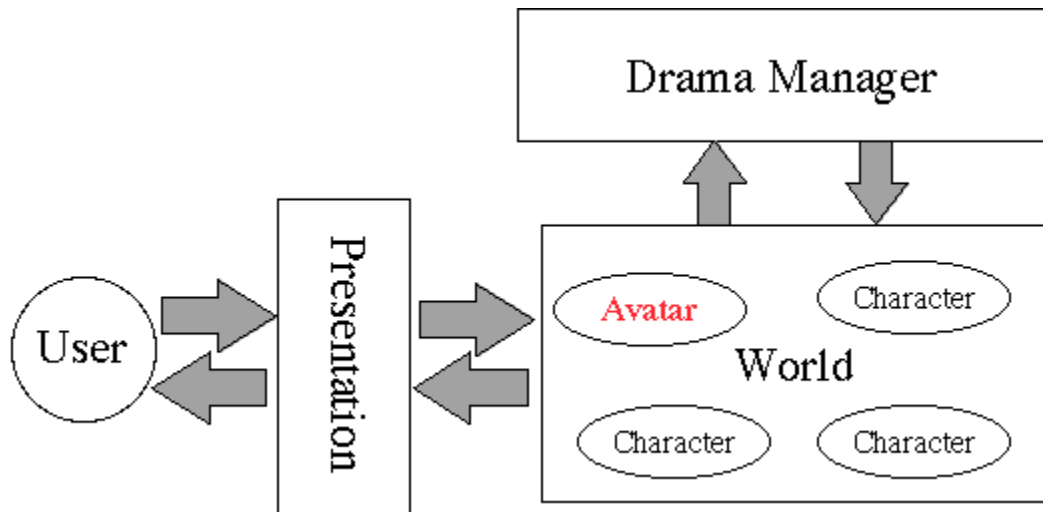


Fig. 1

A simulated world contains characters. These characters exhibit rich personalities, emotion, social behavior, motivations and goals.

The user interacts with this world through some presentation. This presentation may be an objective, third person perspective on the world, or it may introduce various kinds of dramatic filtering - effecting camera angles and point of view in graphical worlds, or changing the style of language used in textual worlds.

The drama manager can see everything happening in the world. It tries to guide the experience of the user in order to make a story happen. This may involve changing the physical world model, inducing characters to pursue a course of action, adding or deleting characters, etc.

Structure of the report

The rest of the report is divided into two main sections, Believable Agents (character) and Interactive Story, and a small section describing interactive video work.

In the believable agents section, I will first describe the Oz research philosophy. Then I will present work related to believable agents (artificial life, virtual humanoids, embodied characters, chatterbots, and behavioral animation) in light of this philosophy. Finally, I will discuss why believable agents is an important and interesting research area.

In the interactive drama section, I will first define the problem of interactive drama in terms of the inherent tension between the concepts of *interaction* and *drama*. After a brief description of the Oz drama manager, I will describe three design dimensions which help structure the design space of interactive drama systems.

In the small section, I'll describe the relationship between the virtual world approach to story and character vs. the interactive video approach.

Believable agents

The Oz philosophy

Taking character seriously

When attempting to marry a technical field like Computer Science with a cultural activity such as story telling, it is extremely easy to become sidetracked from the artistic goals and to begin pursuing purely technical research. This research, while it may be good science, does not lead you closer to building a new kind of cultural experience: engaging, compelling, and hopefully beautiful and profound. Effective techno-artistic research must continuously evaluate whether the technology is serving the artistic and expressive goals. The application of this principle to interactive characters implies that interactive character technology should follow from an understanding of what makes characters believable. And indeed, creators of non-interactive characters have written extensively on what makes a character believable.

Before continuing, it's a good idea to say something about this word *believable*. For many people, the phrase *believable agent* conjures up some notion of an agent that tells the truth, or an agent you can trust. But this is not what is meant at all. *Believable* is a term coming from the character arts. A believable character is one who seems lifelike, whose actions make sense, who allows you to suspend disbelief. This is not the same thing as realism. For example, Bugs Bunny is a believable character, but not a realistic character.



So believability is this good thing that we want characters to have. After examining the writings of several character artists including [The Illusion of Life](#), [Chuck Amuck](#), and [The Art of Dramatic Writing](#), the Oz group defined a set of requirements for believability including the following:

- Personality - Rich personality should infuse everything that a character does, from the way they talk and move to the way they think. What makes characters interesting are their unique ways of doing things. Personality is about the *unique* and *specific*, not the *general*.
- Emotion - Characters exhibit their own emotions and respond to the emotions of others in personality-specific ways.
- Self-motivation - Characters don't just react to the activity of others. They have their own internal drives and desires which they pursue whether or not others are interacting with them.
- Change - Characters grow and change with time, in a manner consistent with their personality.
- Social relationships - Characters engage in detailed interactions with others in a manner consistent with their relationship. In turn, these relationships change as a result of the interaction.
- Illusion of life - This is a collection of requirements such as: pursuing multiple, simultaneous goals and actions, having broad capabilities (e.g. movement, perception, memory, language), and reacting quickly to stimuli in the environment. Traditional character artists do not mention these requirements explicitly, because they often get them for free (from a human actor, or as a deep assumption in animation). But builders of interactive characters must concern themselves explicitly with building agent architectures that support these requirements.

Chapter 2 of [Bryan Loyall's thesis](#) offers a more detailed analysis of the requirements for believability.

Classical vs. behavioral AI

To begin thinking about how to meet the Illusion of Life believability requirement, let's explore the distinction between classical and behavioral AI. In order to make the distinction clear, the following discussion describes the *extreme* classical and behavioral positions. There is certainly work in AI which incorporates aspects of both approaches.

| Classical AI | Behavioral AI |
|---|--|
|  |  |
| narrow/deep | broad/shallow |
| generality | fits an environment |
| disembodied | embodied and situated |
| semantic symbols | state dispersed and uninterpreted |
| sense-plan-act | reactive |

Classical AI concerns itself with building mind, not complete agents. This research program consists of isolating various capabilities of mind (e.g. reasoning, memory, language use, etc.), and building theories and systems to implement a capability in isolation. While it is believed that these disembodied pieces of mind will be put together to form a complete "person", this integration is deferred to the future. Behavioral AI seeks to build complete agents (rather than minds or pieces of minds) that can operate in complex environments. This concern with the environment is one of the key distinguishing characteristics between classical and behavioral AI. Where classical AI attempts to build mental components that duplicate the capabilities of high-level human reasoning in abstract, simplified environments, behavioral AI attempts to build systems with the savvy of insects (or ambitiously, small mammals) in complex environments. Behavioral systems have a *broad* range of *shallow* sensory, decision and action capabilities rather than a single, *narrow*, *deeply* modeled capability.

Classical AI seeks general solutions; *the* theory of language understanding, *the* theory of planning, etc. Behavioral AI starts with the assumption that there is a complex "fit" between an agent and its environment; there may not be generic solutions for all environments (just as many animals don't function well when removed from their environment).

Classical AI divorces mental capabilities from a body; the interface between mind and body is not commonly addressed. Behavioral AI assumes that having a body which is embedded

in a concrete situation is essential for intelligence. Thus, behavioral people don't buy into the Cartesian split. For them, it is the body that defines many of the interaction patterns between the agent and its environment.

Because of AI's historical affinity with symbolic logic, many classic AI systems utilize semantic symbols - that is, pieces of composable syntax which make one-to-one reference to objects and relationships in the world. The state of the world within which the mind operates is represented by a constellation of such symbols. Behavioral AI, because of its concern with environmental coupling, eschews complex symbolic representations; building representations of the environment and keeping them up-to-date is notoriously difficult (e.g. the frame and symbol grounding problems). Some researchers, such as [Brooks](#), maintain the extreme position that *no* symbolic representations should be used (though all these systems employ state - one can get into nasty arguments about what, precisely, constitutes a symbol).


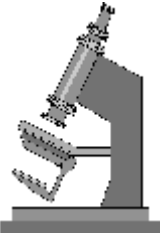
In classical AI, agents tend to operate according to the sense-plan-act cycle. During sensing, the symbolic representation of the state of the world is updated by making inferences from sense information. The agent then constructs a plan to accomplish its current goal in the symbolically represented world by composing a set of operators (primitive operations the agent can perform). Finally, the plan is executed. After the plan completes (or is interrupted because of some unplanned-for contingency), the cycle repeats. Rather than employing the sense-plan-act cycle, behavioral systems are reactive. They are composed of bundles of behaviors, each of which describes some simple action or sequence of actions. Each behavior is appropriate under some environmental and internal conditions. As these conditions constantly change, a complex pattern of behavioral activation occurs, resulting in the agent taking action.

In order to build characters that have the illusion of life, they will need to have broad capabilities to interact with complex environments. This has lead Oz to develop a research philosophy and technology with strong affinities to behavioral AI. The insect-like capability to continuously act in a complex and changing environment is more immediately useful for building lifelike characters than the brain-in-a-vat cogitation of classical AI. The discerning reader, however, may have noticed that Bugs Bunny (or Hamlet, or James Bond, or Charlie Chaplin's Tramp, ...) doesn't seem very similar to either a brain-in-a-vat or an insect. Thus, while behavioral AI begins to give us a handle on the illusion-of-life requirement, the other requirements for believability don't seem to be well served by either camp.

Research goals: believable agents vs. AI

Both behavioral and classical AI share some high level research goals which are at odds with research in believable agents.

| | |
|-------------------|----|
| Believable Agents | AI |
|-------------------|----|

| | |
|---|--|
|  |  |
| Personality | Competence |
| Audience perception | Objective measurement |
| Specificity | Generality |
| Characters | Realism |

For believable agents, personality is king. A character may be smart or dumb, well adapted to its environment or poorly adapted. But regardless of how "smart" a character is at dealing with their environment, everything they do, they do in their own personal style. On the other hand, the focus in AI is on competence. For classical AI, this has often meant competence at complex reasoning and problem solving. For behavioral AI, this has often meant moving around in complex environments without getting stepped on, falling off a ledge, or stuck behind obstacles.

The success of a believable agent is determined by audience perception. If the audience finds the agent believable, the agent is a success. AI tries to measure success objectively. How many problems could the program solve? How long did the robot run around before it got into trouble? How similar is the system's solution to a human's solution? Such audience independent evaluations of research don't make sense for characters.

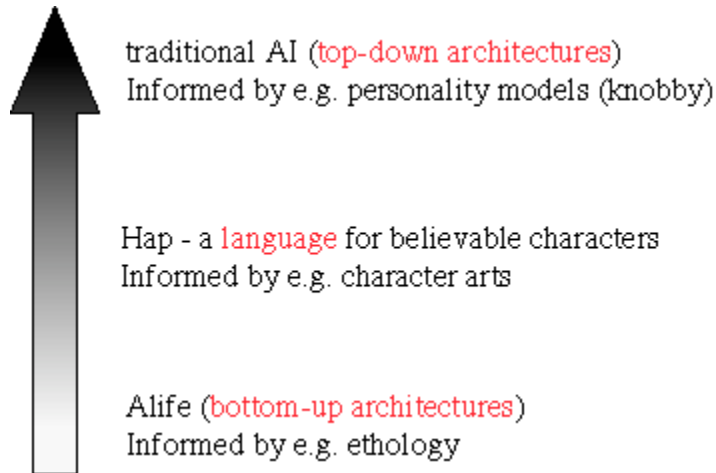
Believable agents stress specificity. Each character is crafted to create the personality the author has in mind. AI, like most sciences, tries to create general and universal knowledge. Even behavioral AI, while stressing the importance of an agent's fit to its environment, seeks general principles by which to describe agent/environment interactions. But for characters, that type of general knowledge doesn't make sense. To what general problem is Mickey Mouse, or Don Quixote, a solution?

Finally, believable agent research is about building characters. Characters are not reality, but rather an artistic abstraction of reality. Much AI research is motivated by realism. A classic AI researcher may claim that their program solves a problem the way human minds really solve the problem; a behavioral AI researcher may claim that their agent *is* a living creature, in that it captures the same environment/agent interactions as an animal.

So, though the need for reactive intelligence gives Oz some affinities with behavioral AI, believable agents are not a problem to which the wholesale import of some AI technology (such as behavioral AI) is the solution. Any technology used for building believable agents will be transformed in the process of making it serve the artistic creation of characters. Thus, believable agents research is not a subfield of AI. Rather it is a stance or viewpoint from which all of AI is reconstructed. Any technology, whether it comes from classical or behavioral AI, or from outside of AI entirely, is fair game for exploration within the Oz context as long as it opens up new expressive and artistic spaces.

Authoring

The desire to pursue the specific rather than the general is strongly connected with the desire to support the direct artistic creation of characters. In traditional media, such as writing, painting, or animation, artists exhibit fine control over their creations. Starting with an idea or vision in her head, the artist uses this fine control to create a representation of her vision in her chosen medium. Similarly, Oz wants to support the same level of artistic control in the creation of believable agents. This approach provides an interesting contrast with both traditional AI and Alife. Below, AI, Alife and Hap (a language developed in the Oz project) are laid out along a spectrum of *explicitness*. This will become clearer shortly.



Traditional AI lies on the high explicitness end of the spectrum. That is, such systems tend to explicitly encode (often in a human-readable form) high level features of the system. For example, suppose you wanted to build James Bond using the traditional AI mindset. First you would think about characters in the abstract. What general theory captures the notion of character? How might this general theory be parameterized (perhaps through infusions of "knowledge") to select specific characters? To inform the work, you might look at the dimensions of personality as described by various personality models in psychology (e.g. introvert-extrovert, thinking-feeling, intuitive-sensing, judging-perceiving). Once a generic architecture has been built, you could then define different characters by setting the right personality knobs. Though the position just described is a bit of an exaggeration, it is not dissimilar to the approach taken by the [Virtual Theater Project](#) at Stanford. For example, in both their CyberCafe and Master/Servant work, they describe using explicit personality dimensions to describe characters. Thus you can actually look in a character's mind and find some symbol denoting whether the character is introverted, intuitive, etc.

Alife lies at the low-explicitness end of the spectrum. A major methodological assumption in Alife work is that you want high-level features (such as introvertedness) to emerge from simple, low-level mechanisms. So how would you go about building James Bond as an Alifer? First, you would demure, saying that Alife technology is not at the stage yet to emerge such high-level behavior. So you might build something else, like a dog. To inform this work, you might look at models developed by biologists, such as ethological models of animal behavior. Then you would build a general architecture capturing an ethological theory of action selection (how animals decide what action to take). Finally, you would instill dog-specific behavior into your general architecture. This approach is not dissimilar to [Bruce Blumberg's](#) approach at the Media Lab in building Silas the dog (though his group's current

work seems more directly focused on building characters rather than building biologically motivated systems).

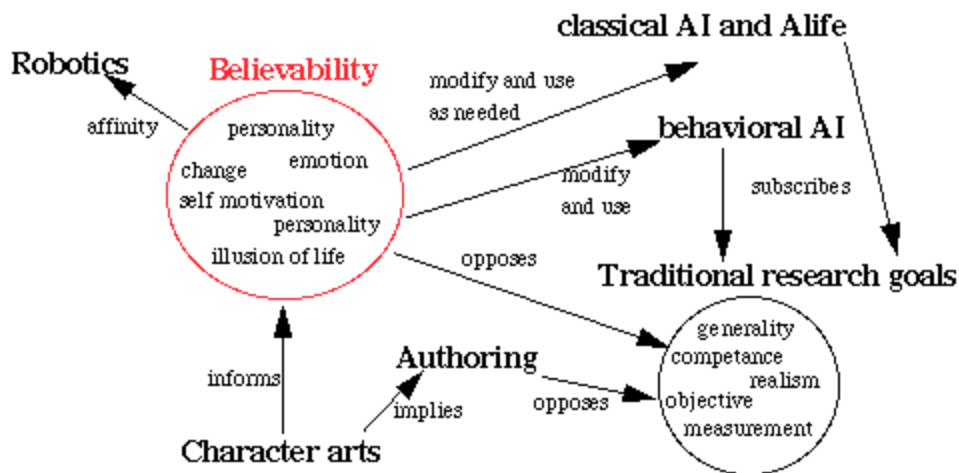
Hap, a language developed in the Oz project for writing believable agents, lies at a midpoint in the spectrum. Hap provides mechanisms that support writing behaviors for characters. A behavior is a chunk of activity; such behaviors can be high-level (a behavior for "having fun"), or low-level (a behavior for moving the character's body when they open a door). If you wanted to build James Bond in Hap, you would identify high-level goals (motivations) that make James Bond who he is. Then you would think of the multiple ways (behaviors) that James Bond might use to accomplish these high level goals. These multiple behaviors probably themselves have subgoals. Any given behavior is only appropriate under certain conditions (what's recently happened, how Bond is feeling, what's happening right now in the world, etc.); these conditions are captured within each behavior. At every level of description, James Bondness can be infused into the character. From how Bond thinks, to how Bond walks, the artist has the control to create the character consistent with their vision.

Both the traditional AI and Alife approaches make architectural commitments; there is some general architecture which characters have to be made to "fit." The traditional AI approach tries to capture high-level mental regularities (e.g. types of personalities). The problem is, how many of these personality knobs are needed to "tune in" a large number of characters? How many personality knobs need to be turned, and how many degrees of freedom does each knob need, in order to allow the creation of Bugs Bunny, Hamlet, The Terminator, Bambi? The differences between these characters seem to far outweigh any similarities. Or to put it another way, is Bugs Bunnyness captured in a few symbols which can be read inside the mind, or is his way-of-being smeared throughout his mind and body?

The Alife approach avoids the use of high level knobs to define personality. Instead, it depends on low-level mechanisms to cause the high-level behavior to emerge. Pragmatically, if you want to build human-like characters, the Alife approach is not understood well enough yet to emerge such high-level behavior. However, this might just be a matter of time. The Hap approach to behavior authoring would then be a useful stop-gap for building characters *today* (we're impatient and don't want to wait) until Alife has developed enough to support such characters. However, from an Oz perspective, there is another problem with Alife: the dependence on emergence. The notion of emergence is that you can't tell what kind of high-level behavior will emerge from low-level mechanisms without actually running the system. But Oz wants to build systems that give artists the control to express their artistic visions. An emergent system removes this control from the artist; the best they can do is make (principled) guesses about mechanism and see what kind of behavior emerges.

Oh, what a tangled web we weave...

The following diagram summarizes the above description of the Oz philosophy.



Taking the character arts seriously leads to requirements for believable agents. The "illusion of life" requirements, namely reactive, situated and embodied behavior, lead Oz to utilize techniques and ideas from behavioral AI. However, work in classic AI and Alife is not automatically rejected on ideological grounds; whatever enriches the space of characters will be *modified* and assimilated. Modification is key: even the behavioral AI ideas, while supporting the "illusion of life", need to be modified in order to support emotion, personality, self-motivation, etc. Believability is not a subfield of AI - it is a stance from which all of AI is transmuted. This is clearly seen in the conflict between believability research goals and traditional research goals. Believability leads Oz to reject the supremacy of the traditional research goals, to which both behavioral and classical AI subscribe. The character arts also point out the importance of artistic control over character creation (authoring). Artistic control opposes traditional research goals as well, particularly generality. Oz wants to build a new canvas and paint brush, not paint-by-number kits. Finally, believability leads to an affinity with robotics. The desire to build believable agents is at heart pragmatic; the agents must live and breath in engaging story worlds. Similarly, roboticists must build systems that act and move effectively in the real world. Thus believability and robotics both share the technical interests of embodied, situated action, as well as a certain pragmatic bent that leads one to pursue what works, regardless of ideological lines.

Related research areas

Now, with the Oz research philosophy in mind, I will explore related research areas. For each of these areas, I will point out similarities and differences with the Oz research program.

Artificial life

The application of artificial life to believable agents is most evident in the design of virtual pets. One of the beauties of virtual pets is that, since they are animals, the audience expectation of the agent's behavior is set at a level the technology can reasonably meet. If the actions of the pet are sometimes confusing, they can be forgiven because we don't always know what animals are doing. Difficult natural language technologies can be avoided because animals don't have language competence. Virtual pets are often cute; cuteness can in itself evoke a strong response in an audience. Two examples of virtual pets are [Dogz](#), and [Creatures](#). Dogz are virtual dogs that you can pet, play with using various toys, and teach

tricks. Over time the user is supposed to develop a relationship with the pet. *Creatures* is a world inhabited by small, cute creatures that autonomously explore the environment. The user is a caretaker; she can provide positive and negative feedback to the creatures, teach them to associate words with objects in the world, and move the creatures around the world. Without the user's intervention, the creatures don't live long.

How do virtual pets relate to Oz believable agents? Pets are one kind of believable agent; Oz wants to build all kinds of characters. Oz focuses on building specific, unique characters. Rather than building dogs, Oz wants to build Pluto, or Goofy. Users interact with virtual pets over extended periods of time; the user builds a relationship with the pet through repeated interaction. Oz believable agents are often designed to be part of a specific story world. Interactions with the character are intended to be intense, but bounded in duration and context by the story world. The notion of repeated interaction with long term characters is certainly an appealing one. It just becomes more difficult to pull off as the character becomes more sophisticated.

Artificial life approaches to building animal characters often rely on modeling of biologically plausible processes. For example, the creatures in [Creatures](#) utilize a neural net for action selection, a model of bio-chemistry for modeling motivations and drives, an artificial genome (with crossover and mutation) for reproduction, and an artificial immune system. [Blumberg's Silas](#) uses an action-selection mechanism motivated by ethology. The intuition behind utilizing such models is that biological systems already exhibit complicated behavior. If structurally similar computational processes are used, this may enable equally complex behavior in artificial systems. However, in Oz, the goal is an artistic abstraction of reality (believable agents), not biologically plausible behavior. By taking a programming language approach to the construction of character (the Hap believable agent language), Oz hopes to avoid premature commitment to an architecture that then limits the space of characters that can be created. Oz remains agnostic with respect to architectures and models. If biologically inspired models end up proving useful in opening up some new space of characters, then they will be used. But modeling for its own sake is eschewed in order to stay focused on the construction of characters.

Finally, artificial life focuses on the concept of emergence. As [described above](#), emergence is at odds with maintaining artistic control over believable agent construction.

Humanoids

Humanoids is the label I'm using for a body of work concerned with building systems that have physical properties (arms, legs, sensory systems) similar to humans. *Virtual* humanoid work is concerned with building realistic, animated humans that live in virtual worlds. Two examples of this work are the [Jack](#) project at the University of Pennsylvania and the work done at [MIRALab](#) (including the famous virtual Marilyn Monroe) at the University of Geneva. In both Jack and the projects in MIRALab, the focus is on building general tools for the animation of human figures, including animating complicated tasks, providing automatic reach and grasp capabilities, and supporting collision detection. MIRALab is currently focusing on the animation of clothes, faces and hair as well as developing architectures to give virtual humanoids autonomy. Though virtual humanoid work started in the graphics and animation communities and was informed by that research agenda, as the humanoid figures have become more sophisticated there has been a natural progression into research concerned with giving these figures autonomous intelligence.

Virtual humanoid work differs from Oz believable agent work in its concern with generality and realism. General toolkits for realistic movement are certainly useful for designing avatars for virtual worlds and perhaps for building background characters (extras). Much of a character's personality, however, is reflected in the unique way a character moves. For building main characters, an author needs detailed control over a character's movement. Similarly, much of the autonomy work associated with virtual humanoids is concerned with providing humanoids with competent action (perhaps to accomplish tasks in virtual worlds) rather than with rich personality.

Japanese robotics researchers are building *physical* humanoids (e.g. [JSK](#), [Waseda Humanoid Project](#)). Examples of this work include a robot that can [swing on a swingset](#), and a [robot with a 3D face](#) that can recognize and produce facial expressions. Such work has focused primarily on the engineering necessary to build and control a complex, jointed humanoid. These robots are not yet capable of sophisticated, autonomous behavior. As such technology becomes more mature, it may open up the possibility of physical believable agents.

Finally, there is a small body of humanoid work concerned with growing intelligence through interaction with the world via a humanoid body ("grow a baby" projects). [Cog](#), at the MIT AI Lab, is probably the best known example of this work. Cog is a robot which has been engineered to have sensory and movement capabilities similar to humans (though its torso is fixed to a pedestal). Cog started with simple motor and sensory reflexes. The hope is that as Cog interacts with the world, it will begin developing intellectual capabilities similar to a human. The guiding hypothesis is that much of human intelligence is the result of sensory-motor interactions with the environment as constrained by human bodies. [Neo](#), at the University of Massachusetts, is a virtual baby living in a simulated world. Neo, like Cog, starts with simple sensory-motor reflexes. As Neo interacts with its world, it learns concepts through a hierarchical sequence of abstractions on streams of sensory-motor data. Both these projects are concerned with growing human-like intelligence (realism) as opposed to building characters.

All the humanoid work shares with Oz the desire to build broad agents which have bodies, sense the world, and take action. Capabilities developed by these projects, either for animating human movement, moving physical humanoid bodies, or physically grounding conceptual thought, may indeed prove useful for opening new levels of sophistication in believable agent behavior. The challenge will be to translate work that seeks to develop general solutions and realistic models into a framework which provides authorial control over the construction of characters.

Embodied characters

Embodied character work is concerned with building physical characters. The physicality of such characters seems to evoke a strong, visceral effect in an audience. While I know of no formal studies of this effect, there is informal evidence.

For example, [Tamagocchi](#), a toy from Bandai corporation, is wildly popular in Japan. It is about the size of a key chain, has a small LCD screen and three buttons. By pushing the buttons to administer positive and negative feedback, provide food and medicine, and clean up feces, the user nurtures a small bird-like creature that lives on the screen. If the creature is not taken care of, it dies. Stores can't keep Tamagocchi in stock; it is being sold for many

times its retail price on the street. Office workers bring Tamagocchi to work and care for it throughout the day. Theft of Tamagocchi is on the rise, especially among teens for whom it is a valued status symbol.

It is unclear how much of this powerful effect is due to social conditions unique to Japan, such as the high cost of pet ownership. However, much of this effect may be due to Tamagocchi's physicality: the fact that it is a small, jewelry-like object (and in fact, teenage girls are wearing Tamagocchi on chains around their necks) that can be incorporated into daily life. Since the character itself is not that complex, the emotional intensity surrounding Tamagocchi may be related to its *ubiquitous* presence.

At Agents 97, Sony demoed a robot dog as an example of their [OpenR standard for household entertainment robots](#). The dog responds to colors, audible tones, and physical touch on its head. The most impressive feature of the dog was its fluid, lifelike movements. As an example, it can smoothly lay down, place its head on its paws, then get back up. In the demo group I was in, everyone responded to this action with "ahhhhh" (cuteness). In this case, I believe the strong response comes from the animal-like movement of a physical object.

[Karl Wurst](#) at the University of Connecticut is building robotic puppets based on the woggles (characters built by the Oz project). While these puppets roll rather than hop (the original woggles hop), they are able to stretch and squish (woggle body language) and communicate with each other via IR sensing. It would be interesting to compare the audience response to these puppets with the response to the behaviorally similar screen-based woggles.

Providing a believable agent with a physical body is an interesting research direction to pursue. The combination of rich behavior, personality, and physicality could produce a powerful audience response.

Chatterbots

Chatterbots are programs that engage in conversation. The original chatterbot is [Eliza](#), a program that uses sentence template matching to simulate the conversation of a non-directive therapist. [Julia](#) is a chatterbot that connects to multi-user dungeons (MUD). Besides engaging in conversation, Julia has a simple memory that remembers what's been said to her and where she's been. She uses this information in her conversations (repeating what someone else said or providing directions). When she is not engaged in conversation, she wanders about the MUD exploring. Erin the bartender from [Extempo](#) is a recent example of a chatterbot. Erin serves drinks and converses with customers (music is a favorite topic of conversation). She has an emotional state (influenced by what you say to her, whether you argue with her, etc.) and forms attitudes about customers. How she responds to any particular utterance is influenced by her current state.

There are several differences between chatterbots and believable agents. First, chatterbots primarily interact in language. Body movement and physical activity play a secondary role; if it is present at all, it is used to provide some background color during lulls in a conversation. The language interaction is primarily reactive; the chatterbot is responding to each utterance without its own goals for the conversation. In the absence of an utterance to respond to, the chatterbot may fall back on some small number of stock phrases that it uses to try and start a conversation. Second, many chatterbots are designed for entry into a restricted form of the

Turing test (the [Loebner Prize](#)). The goal is to fool a human for some short period of time into thinking that they are interacting with another human. Notice that the goal is not to communicate a personality, but rather to briefly fool a user into thinking that they are talking to some generic person during a context-free conversation. Finally, most chatterbots don't have a long-term motivational structure; they don't have goals, attitudes, fears and desires. The conversations they engage in don't go anywhere. A chatterbot's only goal is to engage in open-ended conversation.

In contrast, believable agents express their personalities through their movements and actions, not just through language. Believable agents are designed to strongly express a personality, not fool the viewer into thinking they are human. When watching a play or film, viewers know that the characters are not "real" but that does not detract from being engaged by the character. Finally, believable agents have long-term motivational structures. Their behavior is designed within the context of a particular world. Within this context, the believable agent's behavior is conditioned by desires and attitudes. On the other hand, the lack of a long term motivational structure and the focus on language interaction allows chatterbots to function within open environments (such as chat rooms or MUDs) where they can serve as a social catalyst for the human participants.

Behavioral animation

Behavioral animation has developed in the graphics community as an alternative to hand-animation. In more traditional computer animation, the animator builds a model of the character they wish to animate, defines parameters that move and deform parts of the model, and writes functions that smoothly interpolate the values of parameters given beginning and end values. Now, in order to make the figure do something, the animator must define a set of initial and final values of all the parameters (keyframes) and apply the interpolation functions to generate all the intermediate frames. Even after doing all the upfront work of building the model and defining the functions, the animator still needs to define keyframes in order to make the model move.

Behavioral animation seeks to eliminate the work involved in defining keyframes by pushing more work into the upfront creation of a model. Instead of just defining the geometry, the model also includes code that tells the model how to move in different situations. Given a state of the world, the model moves itself. Some behavioral animation work focuses on general strategies for realistic movement. In this respect, behavioral animation shares some common goals with virtual humanoid work. However, as more internal state is added to the behavioral routines, state which may represent emotions or social attitudes, behavioral animation begins converging on believable agents. Whereas believable agent research begins in AI (the building of minds), and then appropriates and modifies AI technology to the task of building characters (minds and bodies), behavioral animation research begins in graphics (the building of bodies), and adds behaviors to these bodies to build characters (bodies and minds).

A good example of the convergence between behavioral animation and believable agents is [IMPROV](#), a system built by Perlin and Goldberg at NYU. As part of IMPROV, they have developed a scripting language for writing animation behaviors. Behaviors written in this language can be conditional on author-maintained internal state as well as external events. The main mechanism for creating non-deterministic characters is the tuning of probabilities. The author communicates the character's personality and mood by tuning probabilities for

selecting one action over another. Both IMPROV and Oz share an author-centered point of view. However Hap (the Oz believable agent language) provides more support for expressing complex control relationships among behaviors. In addition, [Em](#) provides support for maintaining complex emotional state (something that would have to be done manually using the IMPROV language). On the other hand, the procedural animation portion of the IMPROV scripting language provides more language support for animating control points on the model.

Why study believable agents?

I've described the Oz philosophy regarding the believable agent research program and reviewed related research areas. The reader still may be left with a nagging question: why study believable agents at all? The most obvious answer is that believable agents are necessary if you want to build interactive story worlds. This is the primary motivation behind the Oz research program. There are other reasons to pursue this research, however.

Believable agents may greatly enhance learning in educational settings by providing engagement and motivation for the learner. Research in this area is being pursued by the [IntelliMedia](#) project at North Carolina State University. They have built a constructivist learning environment in which children learn about the biology of plants by building a plant (selecting different kinds of roots, and leaves, etc.). A believable agent serves as a companion and guide for the student.

Believability will be important for building anthropomorphic interface agents. Research by [Nass and Reeves](#) at Stanford University has shown that users interpret the actions of computer systems using the same social rules and conventions used to interpret the actions of people, whether or not the computer system is explicitly anthropomorphic. Since most systems aren't designed with this fact in mind, the resulting social behavior of the system (its personality) is accidental. As designers begin building systems with *designed* personalities, they will need techniques for communicating this personality to the user. This is precisely the research area of believable agents.

The three motivations given above are pragmatic reasons to pursue this research. There is also a more distant, idealistic, yet compelling reason for pursuing this research: the AI Dream. This Dream, to build companions such as Data on Star Trek, has motivated many workers in the field of AI. Woody Bledsoe, a former president of AAAI, captured this dream nicely in his [1985 Presidential Address](#). In describing the dream that motivated his career in AI, he opened:

Twenty-five years ago I had a dream, a *daydream*, if you will. A dream shared with many of you. I dreamed of a special kind of computer, which had eyes and ears and arms and legs, in addition to its "brain." ... my dream was filled with the wild excitement of seeing a machine act like a human being, at least in many ways.

Note that he did not talk about some disembodied mind; this is a complete creature. Later he states:

"My dream computer person *liked* (emphasis added) to walk and play Ping-Pong, especially with me."


Clearly the AI dream is not just about rational competence, but about personality and emotion. As described above, believable agents research is not a subfield of AI, but rather a stance from which AI can be reinterpreted and transformed. The believable agents research program, by directly engaging the issue of building complete agents with rich personality and emotion, provides a new agenda for pursuing the AI Dream.

Interactive story

Drama consists of both characters and story. In *interactive* drama, believable agents are the characters. Now it's time to talk about story.

Interactive story: an oxymoron

Many observers have remarked that the concept of interactive story contains a contradiction. A story is an experience with temporal structure. Interaction is doing what you want, when you want (interaction as control; other models are possible).



Accounts of story structure often describe some form of dramatic arc (first introduced by Aristotle in *The Poetics*). One form of the dramatic arc is shown above. The vertical axis represents tension, or unresolved issues or questions. The horizontal axis represents time. At the beginning of the story, during the exposition, the tension rises slowly as the audience learns the background of the story. An inciting incident then sparks the story. Tension begins rising more rapidly after this incident. Eventually, the amount of tension, the number of unresolved questions, the intertwining between plot elements, reaches a critical state. During this crisis, the tension rises rapidly to the climax. During the climax, questions are answered and tensions resolved. After the climax, the tension falls rapidly as any remaining tensions are resolved. Finally, during the denouement, the world returns to some status quo. The experience of a story is thus structured; events don't happen in some willy-nilly fashion. The experience has a global shape. Interaction, on the other hand, is generally construed as the freedom to do anything at anytime. Story is predestination; interaction is freedom. Thus the conflict.


Some have resolved the conflict by saying that interactive story is impossible. Others have redefined the notion of story to have less structure; whatever emerges from interaction is defined as story. Brenda Laurel, in her [1986 thesis](#), described a hypothetical expert system that causes a structured story to happen in the face of interaction. While the technology is different, the Oz drama manager takes this approach of simultaneously honoring story structure *and* interaction.

Oz drama manager

The [Oz drama manager](#) controls a story at the level of plot points. Plot points are "important moments" in a story. In an hour and a half film, there may be 12-15 of them. Given a particular set of plot points, the space of all possible stories is the set of permutations of all possible plot points. The vast majority of these permutations will be garbage - unsatisfying stories which don't make sense. The author of the story has some particular ordering of the plot points in mind - this is the story she wants to tell. Rather than expressing this preferred sequence via structural constraints on the story world, the author writes an evaluation function that captures her sense of aesthetics for the story. This aesthetic is captured by some set of features the evaluation function looks for in a permutation. Conforming to the shape of some dramatic arc may be one feature in the function. Given a permutation of plot points, the evaluation function rates the permutation. Assuming the author has successfully captured her aesthetic, the original story should be ranked high by the function. So the authorial process is:

1. write some linear (non-interactive) story as a sequence of "important moments"
2. reverse-engineer your own thinking to figure out why you think that particular sequence is a "good" story
3. capture this aesthetic as a set of features (over sequences) in an evaluation function
4. make sure that you have (approximately) captured your aesthetic by comparing the output of the evaluation function with your own evaluation of a set of sequences (of course include the original story - the one your "really" want to tell)

With an evaluation function in hand, you can now do search.



The drama manager watches the state of the world (including the user interaction). While the user is moving around and interacting with characters *within* some particular plot point, the system isn't doing anything but watching. Eventually, some sequence of activities in the world will be recognized as causing a plot transition. The drama manager springs into action. There exists some past history of plot points. At this point in time, the future histories consist of all possible sequences of remaining plot points. Sequences of events that result in a plot transition are abstracted as user moves. The drama manager has a set of operations it can perform to warp the world: these are the system moves. In a manner similar to game playing programs (such as chess programs), the manager examines every possible system move it could perform to warp the world, every possible user move the user could make to cause a plot transition, every possible system move from that new state of the world, etc. until it has played out the possible histories. The past history plus each possible history forms a set of total histories. The evaluation function can now evaluate each total history. The system then

makes a system move (warping the world in some way) that maximizes the probability of generating a highly ranked total history. In this way, a story structure is imposed on the viewer's experience, while still allowing interaction.

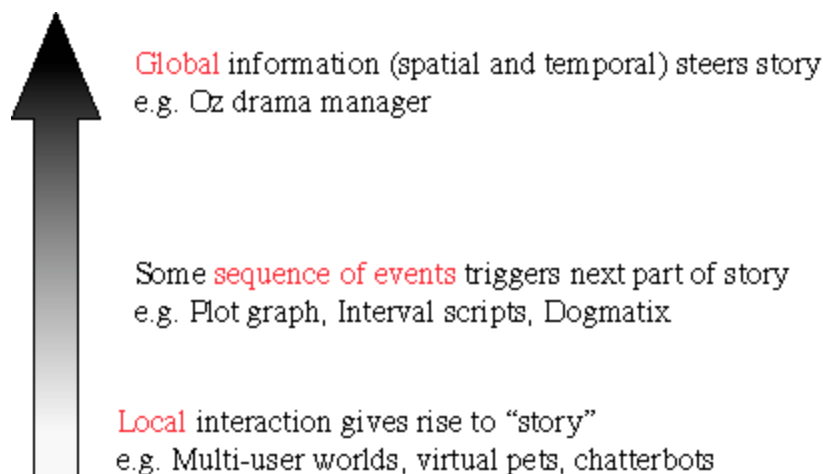
Design dimensions

Having briefly examined the Oz approach to interactive drama, I will now examine related work. The first comment to make is that there is less related work on interactive story than on believable agents. Believable agents work can be construed as the construction of little "people". Even though there is not much work directly concerned with believability, there is a body of work concerned in one way or another with building little people. Interactive story, by comparison, is relatively unexplored. Instead of describing the relationship between various research areas and the Oz approach, as was done for believable agents, I will describe three design dimensions. Each of these dimensions represents a spectrum of choices that can be made with respect to a design question. Various research projects in interactive story can be displayed along these dimensions.

While each dimension has a "low" end and a "high" end, this is not meant to imply that low is bad and high is good. Systems laying on different points along these dimensions have different properties and will be useful for generating different kinds of experiences. The dimensions merely indicate a space of potential to be explored.

Local vs. global control

A drama manager can take smaller or larger blocks of spatio-temporal structure into account when deciding how to control a story. By spatio-temporal structure I mean *action* as it unfolds in the space of the story world across the time of the story. To the extent that a drama manager only looks at the action that has immediately occurred in the area around the audience (user), the control is local. To the extent that the manager takes into account the entire history of the story across the entire space of the story world, the control is global.



At the extreme local end of the spectrum are systems in which interaction with characters is the only mechanism structuring experience. For example, when interacting with other people in a multi-user world, the structure of the experience arises out of these moment-to-moment interactions. As a shared history develops among users, this history will condition future

interactions. Similarly, interaction with artificial characters such as virtual pets and chatterbots share such local structure. Such purely local control doesn't give rise to story in any strong sense of the word; an audience is not carried through some author-defined shaped experience in the course of their interaction with a system.

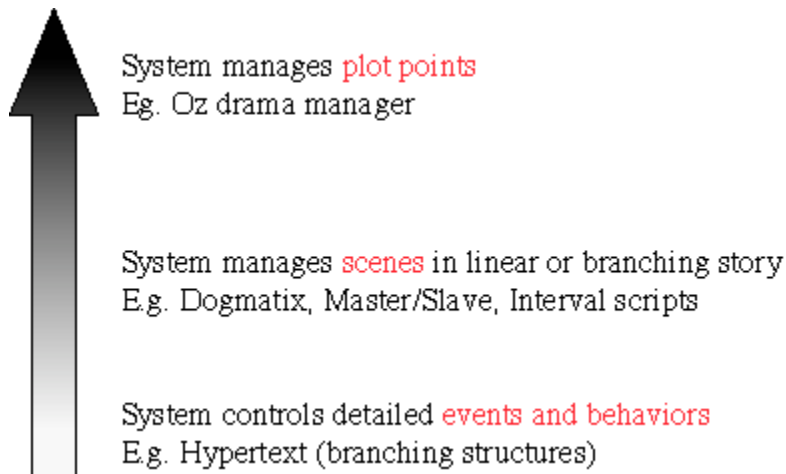
At an intermediate point on the spectrum are systems that control a story by taking into account some history across some physical space of the story world. Such systems can be characterized as *script-and-demon* systems. The script specifies a linear or branching sequence of events. These events can be guarded by demons that won't let the event happen unless some preconditions on the state of the world have been satisfied. [Plot graphs](#), an early approach to drama in the Oz project, are one example of such a system. A plot graph lays out scenes in a directed acyclic graph (DAG). The arcs represent the must-precede relationship. Only after all preceding plot points have happened can the next plot point be entered. Associated with the arcs are hints and obstacles. These are ways that the drama manager can influence the world. Hints make it more likely that the user will move into the next scene; obstacles slow the user down. Demons recognize when a user has completed a scene. Another example, Pinhanetz's [Interval Scripts](#), represents the script by using a temporal calculus to record temporal relationships among intervals. Some of these intervals are connected to sensors (demons) that wait for events to occur in the world; other's are connected to actuators that make events happen in the world. A constraint propagation mechanism is used to determine the state of each interval (*now*, *past*, *future*, or some mixed state). When a sensor has the value *now*, it begins looking for its associated event to happen in the world. When an actuator has the value *now*, it makes its associated event happen in the world. The final script-and-demon system I'll discuss is the plot control mechanism in [Galyean's Dogmatix](#). Galyean makes an analogy between the action selection problem in behavioral agents and the event selection problem in plot control. At each point in time, a behavioral agent must select one (or in general, some small subset) behavior from its pool of possible behaviors. This selection is accomplished as a function of the internal state of the agent and the external state of the world. Analogously, at each point in time a plot selection mechanism must select an event to make happen out of the set of all events it could make happen. In Galyean's system, this selection is a function of story state variables (*history*), sensors (demons watching for events in the world), and *temporal relationships*. The temporal relations *hierarchy*, *before*, *xor*, and *must-happen*, place a partial order on the possible sequences of events chosen by the selection mechanism. At each point in time, the event that has the highest "fitness" is chosen for execution.

In script-and-demon systems, the complexity of the demons is the limiting factor. In order to take more and more global information into account, the firing conditions on the demons must become more and more complex. Perhaps because of this complexity, in practice demons tend to fire on relatively local sequences of events. And regardless of how complex a demon's firing condition becomes, it can only take the past into account. It can not look into the future to see what might happen in the story.

At the global end of the spectrum is the Oz drama manager. Whenever it detects a user move, it considers total story histories by concatenating the entire past history with projected future histories. These total histories are evaluated to determine which events to make happen in the world.

Granularity of story control

A drama manager can seek to control the story at different levels of detail. To the extent that a manager controls precisely the actions of characters (what they do and when they do it), the manager is controlling the story at a small grain size. To the extent that a manager controls the general direction of the story, but does not directly control the activities of particular actors, the manager is controlling the story at a large grain size.



At the extreme small-grain-size end of the spectrum, are systems that directly control the detailed events in the story and behaviors of the characters. In such systems, there isn't really a distinction between the drama manager and the world; the structure of the world *is* the drama manager. Hypertext stories are an example of such a system. The branching structure of the story precisely describes what happens and when it will happen. Within a node of the branching structure, there is no variation. The same fixed events happen at a given node every time it is visited. Some CD-ROM games also use this approach to story; each node in a branching structure completely describes what a user will experience.

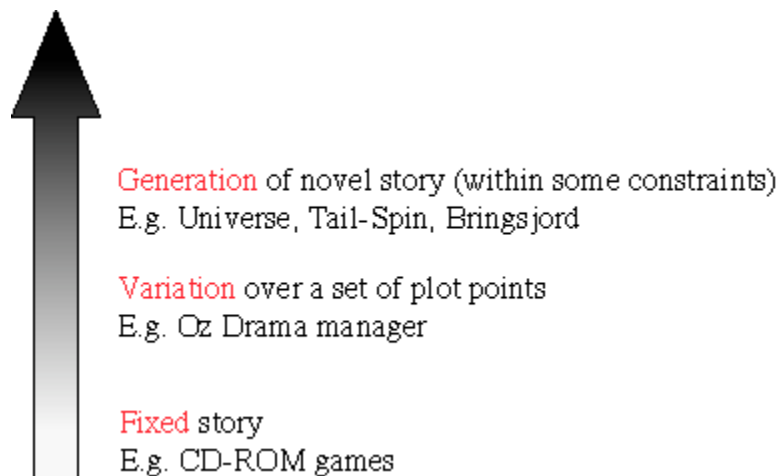
At an intermediate point on the spectrum are systems that manage scenes. In such systems, the progression of scenes is fixed by a linear or branching structure. But what happens within a scene is not completely predetermined; within the scene, there is room for variation in response to user action or non-determinism on the part of the agents. Script-and-demon systems can be used to provide this granularity of control. Two examples are Galyean's event selection system and Pinhanez's interval scripts (described above). [Hayes-Roth's master/servant scenario](#) is another example of scene level control. In this system, which is not interactive (the user doesn't play a character), a master and servant play out a power struggle which can end in the master and servant switching roles. The script issues directives to the characters. The characters engage in action as a function of these directives and their internal state. The script specifies the order in which directives are issued. Demons wait for certain conditions to be met in the world (e.g. "improvise until the master's demeanor is low") before allowing the script to continue.

At the large-grain-size end of the spectrum are systems that decide the order of plot points (which can be scenes); there is no linear or branching structure fixing the order of scenes. For example, the Oz drama manager repeatedly searches the space of scene orderings in order to decide what to do next to influence the story. "Good" orderings are captured *implicitly* in the evaluation function. Each time the user runs through the story, the ordering of scenes can be different.

A single story system may need multiple drama managers at different granularities of story control. A system like the Oz drama manager could select scene orderings. However, within a scene, story control will still be required to handle staging. One approach is to have the individual characters have enough knowledge to not only play their roles but also control the staging. Another approach is to have some sort of script-and-demon system control the staging within scenes.

Generation

A drama manager can be more or less generative while it controls a story. To the extent that a drama manager has a fixed description of a single story (linear) or set of stories (branching), it is not generative. The possible stories that a user can experience while interacting with the system are fixed. To the extent that the manager can create a new story each time a user experiences the system, the story is generative. Another way of thinking about this is capacity for surprise. To the extent that a manager can surprise its author with a novel story, the system is generative.



At the fixed end of the spectrum lie systems like CD-ROM games. The story structure is completely fixed by a branching structure. Such games don't often bear replaying; after having played through the game, there is nothing new to experience.

A bit higher on the spectrum are systems that support variations on a theme. For example, the Oz drama manager can change the order of plot points, or not include plot points, each time the user experiences the story. Though the same story is not experienced each time, it will consist of some sequence of plot points from a fixed pool. The extent to which such a system seems generative will depend on the level of abstraction of the plot points and the complexity of the evaluation function.

Still higher on the spectrum are systems that generate novel stories. Unfortunately, the examples of such systems are not interactive; the systems generate a textual story that is read by the user. [Universe](#) tells a serial soap-opera-like story. Characters are described by sets of attributes. Example attributes are interpersonal relationships (e.g. ex-spouse, div-mom), stereotypes (e.g. party-goer, egomaniac), and goals (e.g. become-famous, associate-right). A library of plot fragments (plans) serves as the raw material for composing stories. Each plot fragment describes the kinds of characters it requires (constraints on the

traits), the goals the plot fragment can be used to satisfy, and the subgoals necessary to accomplish the plot fragment. Stories are told by composing these plot fragments. In addition, the system learns new plot fragments by generalizing old ones. [Tail-spin](#) tells Aesop-fable-like stories. It does not use a library of plot fragments. Instead, stories are generated purely by trying to accomplish the (sometimes conflicting) goals of characters. Both these systems view story telling as a planning problem. [Bringsjord's](#) work is a modern example of non-interactive story generation.

Generation raises the interesting riddle of authorial control. Good authors write good stories - that is, stories which audiences find engaging. If an author takes the trouble to write a good story, you probably want your system to tell that story. At what levels of abstraction can an author still exert authorial control? An author can say "tell this exact love story." Clearly they have control; its a fixed story where interaction basically means moving through at your own pace. An author might say "tell a love story generally similar to this one." Somehow you would have to capture the author's knowledge of what makes a story "similar to this one." This is the aesthetic as captured by the evaluation function in the Oz drama manager. An author might say "make up a love story that sounds like I made it up." What aspects of the author (knowledge, feelings, history) have to be captured in the system to maintain authorial control but allow this kind of flexibility? As you increase the generative power of a system, can you still capture the richness of a particular authorial point of view?

Clip-based work

A non-agent based approach to interactive drama is interactive digital video. The work of [Davenport](#) is characteristic of this approach. I include this work in a separate section, rather than including it under character or story, since interactive video combines aspects of both.

The basic approach is to store and index some large number of video segments. As the user interacts with the system, the system must decide which segment is the appropriate one to play next. The system making this decision may be something very like a drama manager. However, interactive video can also be used primarily as a character, rather than a story technology. The [Entertainment Technology Center](#) at Carnegie Mellon has built several prototypes of such systems. One of their systems, recently demoed at the ACM 50th Anniversary conference, allows the user to have conversation with Einstein. In this system, many shots of an actor playing Einstein are stored and indexed on disk. A user speaks to the system. A speech recognition system converts the user's utterance into text. Based on this text, the most appropriate clip is played. While the simplest version of such a lookup involves comparing the user's utterance against the video index using word-based indexing technology, one can easily imagine building some kind of personality model that maintains state based on the previous course of the conversation. This state could then be used to bias the selection of the video clip. The goal of these interactive interviews is to give the audience the feeling of actually talking with a famous personage.

For both character and story, a clip-based approach is based on selection rather than generation. An appropriate clip must be chosen from some set of clips. In a believable agent approach to character, the behavior is generated in real time. Though the library of behaviors for such an agent is fixed, the granularity is much smaller than a video clip. There is thus much more flexibility in composing these behaviors. In addition, the structures for representing internal state (such as emotion) and for representing behaviors are made out of the same "stuff" - computer code. This allows the state information and the behaviors to

intermingle in complex ways. In interactive video, the code representing the character's state and the data representing the actions (video) are of different kinds. To achieve the same flexibility as the computational representation of behavior, there would have to be a video indexing scheme that captures the detailed action of each clip, as well as a means for changing clips during playback (speeding them up, slowing them down, changing whether the figure in the scene is looking left or right, etc.).

On the other hand, clip-based approaches can immediately take advantage of the skills of actors. Rather than having to generate a facial expression, a movement or an utterance with all the subtlety of a human actor, you can immediately use the skill of the human actor by filming her and including the clip in the database. Also, all the techniques that have been developed by the movie industry for creating engaging video sequences are at your disposal.

Conclusion

Believable agents and interactive drama are two relatively new research fields. Both research areas are pursuing the combination of insights and knowledge from the dramatic arts with computer technology. Bringing rich personalities and story structures to computing promises to open up new realms of human expression and experience.