

# Story-Making with Improvisational Puppets

Barbara Hayes-Roth and Robert van Gent

Department of Computer Science  
Stanford University  
Stanford, CA 94305

Extempo Systems, Inc.  
Santa Clara, CA 95054

{bhr, vangent}@cs.stanford.edu

## Abstract

"Improv Puppets" is an interaction paradigm modeled after traditional puppet play. With physical puppets, children move their puppets' bodies and speak their puppets' words in real time. Thus, the puppets serve as passive vehicles for the children's collaboration on improvisational story making. Similarly, with Improv Puppets, players direct the physical and verbal behavior of their puppets in real time. Now, however, the puppets are synthetic agents. They are animated and smart. The puppets tell the children which high-level directions make sense in a given situation. Given the children's choices, the puppets improvise a joint course of behavior. Thus, improvisational story making becomes a collaboration among children and puppets. This paper describes the design and behavior of an implemented system that allows two children to direct the improvisations of two puppets in real time.

## 1. Directed Improvisation

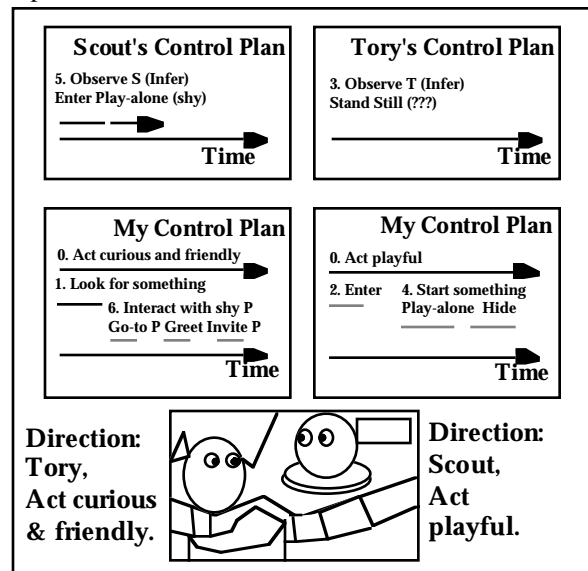
Improvisational actors create performances in real time, without planning, often working under constraints from the audience. There spontaneous performers do not reach for the artistic heights of conventional theater. Nonetheless, audiences find them entertaining. A skilled improv troupe produces stories that are pleasing in form and content. Even when they are only moderately successful, audiences enjoy watching the actors work to create a story while meeting the constraints. They take a special pleasure in knowing that each performance is unique.

We are exploring the possibility of creating synthetic actors that can be embodied as animated characters, perform in a manner loosely resembling that of human improvisers, and tailor their performances to abstract directions from users or other system components.

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

Agents '97 Marina del Rey CA USA  
(c) 1997 ACM

To illustrate the basic idea, Figure 1 presents a hypothetical episode involving two agents embodied as animated characters: a large character (Tory) and a small character (Scout). Tory has been directed to act curious and friendly. Scout has been directed to act playful. These directions are quite general; neither character has been directed to do anything in particular, only to behave in ways that reflect the specified moods. Each character incorporates the directions into the "control plan" that will guide his or her own behavior (Plan 0) for the duration of the episode.



**Figure 1.** Two actors perform directed improvisation. The box above each one shows its "control plan" for its own behavior. The upper box shows the control plan it infers from its partner's behavior. Elements in each control plan, numbered 0-6, indicates "intended" kinds of behavior. (The actors shown are embodied as "Woggles," after the Bates animation [Loyall & Bates 1993])

As the episode unfolds, each character improvises its own behavior, while interpreting and responding to its partner's behavior and satisfying the constraints of its directions. At first, Tory is alone in the world but has many possible behaviors. He can't act friendly without a partner so, following his direction to act curious, Tory decides to look around for something (1). Scout enters (2), observes Tory standing still, tries to interpret his behavior, but infers

nothing (3). Following her direction to act playful, Scout decides to start something: to play alone and then hide(4). Observing Scout enter and play alone, Tory interprets her behavior and infers that she is shy (5). Specializing his act-friendly direction for a shy character, Tory decides to approach Scout, greet her, and invite her to play (6).

Although they don't know it, Tory and Scout now have incorrect models of one another and conflicting plans for interaction. One of them must change, but it doesn't matter which one. If Tory invites Scout to play first, Scout will drop her own plan and accept his invitation. If Scout hides first, Tory will drop his plan and join her game. Following the most basic rule of improvisation, the characters *accept all offers* [Johnstone 1987]. Each one actively seeks to interpret the other's behavior and readily changes his or her own plans to accommodate the other's apparent intentions.

We envision the following capabilities for synthetic actors. Each agent, embodied as an animated character, would accept directions from one or more exogenous sources, either in real time or in advance of a performance. The directions would constrain, but not completely specify the agent's behavior, for example assign the agent a role, endow it with personality features, change its mood, or instruct it to perform a kind of behavior. A sequence of directions might shape the narrative structure of the agent's individual behavior and its interactions with other agents. Each agent would construct its own detailed course of behavior, following directions, interacting appropriately with other agents, and filling in unspecified elements along the way. Its behavior would make sense in the real-time situation, reflect its role, personality, and mood, and manifest life-like qualities, such as normal variability and idiosyncrasies in behavior. In some cases, the agent's improvisations might contribute to the narrative form and content of the story. If a given sequence of directions were repeated on different occasions, the agent might improvise different performances, surprising the audience and perhaps even its director.

We believe that improvisational actors would be useful components in diverse applications in commercial, education, and especially entertainment, the arts, and children's learning toys [Hayes-Roth 1995a]. Alternative interaction paradigms are possible. In this paper, we focus on a paradigm called "Improv Puppets."

## 2. Technical Approach

### 2.1 System Organization

An agent has a "body," a "mind," and a "mind-body interface." The logical system organization (Figure 2; see also [Hayes-Roth et al. 1994]) permits alternative physical organizations.

An agent's body is a computer program controlling a graphical manifestation in a virtual world, which it may cohabit with other embodied agents. It has sensors and

effectors to perceive events and execute behaviors. It also may have an interface for communicating with users.

For Improv Puppets' bodies, we adapted the "Woggles" animation developed by Bates [Loyall & Bates 1993]. Augmenting their basic physical behaviors (e.g., leap to x, y, stretch, turn, blink), we created a few new gaits, so that puppets can introduce normal variability and mood-related expressiveness into their movements. For voices, we gave each puppet 250 lines, conceived and recorded by Aaron and Nora Hayes-Roth, ages 14 and 11, in response to elicitation cues comprising classes (e.g., greeting, invitation to play) and qualitative values of three moods: happy-sad, peppy-tired, friendly-shy. (A more recent version of the system incorporates the voice of Nathan

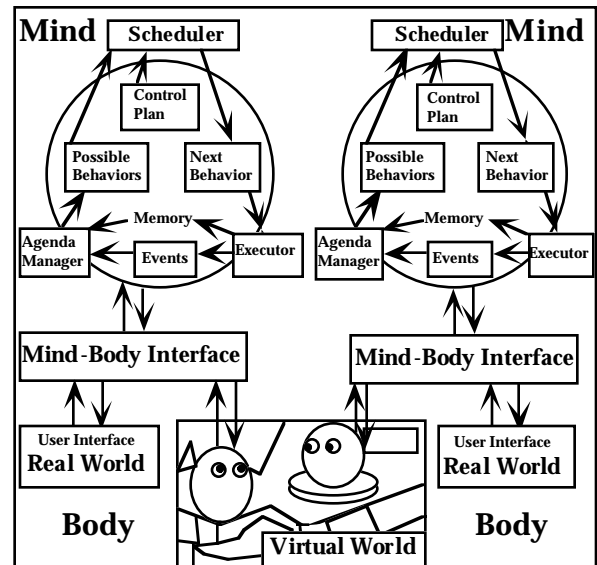


Figure 2. Logical system organization for two agents.

Hayes-Roth, age 5). Thus, the puppets engage in a variety of activities and interactions, introducing normal variability and mood-related expressiveness into their verbal behavior. The children's voices and words endow them with personalities. Each puppet's user interface has buttons, sliders, and a clickable world, as discussed below.

An agent's mind is a computer program that performs three functions. It integrates perceptual inputs with knowledge and inferences to assess the agent's dynamic situation. It instantiates and decides when to execute particular behaviors. It performs all processing intervening between situation assessment and behavior. The puppets' mind architecture and knowledge are discussed below.

An agent's "mind-body interface" mediates interactions between its mind and body in a tight real-time control loop. It classifies patterns of sensor data as meaningful perceptions, which it relays to the mind. It translates motor commands from the mind into sequences of executable instructions, which it relays to effectors. Thus, the mind-body interface allows an agent's mind to be transferred among different bodies with minimal modification and vice versa, so long as the agent is given an appropriate set

of mappings between sensor data and perceptions and between motor commands and effector instructions.

In the current implementation of Improv Puppets, system components are distributed among processes, platforms, and locations. Each puppet's mind, mind-body interface, and user interface run as separate processes on a single Sun SPARCStation 10 in our laboratory. Multiple copies of the virtual world, each one controlling all of the puppets' bodies and displaying one puppet's user interface (see Figure 3), run on different SGIs in other buildings. (SGIs are required for the animation; other system components run in any C++ environment.) Puppets' mind-body interfaces broadcast instructions to their bodies in all copies of the virtual world. Players direct their own puppets and view the joint performance at individual SGI's.

## 2.2 Architecture of the Mind

Each agent's mind instantiates and elaborates a *dynamic control architecture*, which has been applied in other intelligent agent applications [Hayes-Roth 1995b, Hayes-Roth 1993, Hayes-Roth et al. 1994] and is available as the BBI software system. As illustrated in Figure 2, the architecture iterates a three-step execution cycle. First, an *agenda manager* uses recent perceptual and cognitive events, states information, and knowledge to instantiate behaviors that are *possible and relevant* in the present situation. Second, a *scheduler* selects relevant behaviors that satisfy the constraints in the current *control plan*. Third, an *executor* sends the associated behavior commands to the mind-body interface. It also executes associated cognitive behaviors, which produce cognitive events and may cause changes in state, including the contents of control plans. The puppets' minds introduce two additional steps into the BBI control loop:

(a) Whenever the mind either identifies a new set of relevant behaviors or changes its mood, it sends corresponding new information to the user interface.

(b) Whenever the mind receives a new direction from the user interface, it improvises a course of behavior consistent with that direction.

The agent architecture can exploit many different kinds of knowledge and information in various data structures.

Potential behaviors are organized in class hierarchies. Each behavior has a set of relevance conditions that can be satisfied by the occurrence of certain classes of events. When a behavior's relevance conditions are satisfied, the agenda manager instantiates the behavior for the current situation. The class hierarchy allows the agenda manager to search efficiently for relevant behaviors. Each of our puppets has class hierarchies for 15 physical behaviors and about 250 verbal behaviors. For example, "hop to destination" is a sub-class of "go to destination;" "greet friend" is a sub-class of "speak to friend."

Each behavior has annotations specifying "typical" values on a set of mood variables (e.g., happiness). A behavior is considered more appropriate when the agent's actual mood approximates the behavior's typical mood

values. The agenda manager uses mood values and random variability to prune the current set of relevant behaviors. The scheduler uses mood values and other information to choose which relevant behaviors to execute.

Each leaf in the behavior hierarchy has an execution script, which the executor instantiates and executes when that behavior is scheduled. Scripts are represented in a restricted programming language that permits variable instantiation and conditional logic terminating in commands to the puppet's body. Simple scripts (e.g., hop to destination) directly command particular instruction sequences to the animated bodies. Complex scripts (e.g., Play alone for awhile) recursively invoke sequences of instantiated behaviors from the hierarchy. Script parameters permit different improvisations.

Although the puppets' class hierarchies are identical, their executable behaviors are different. Physical behaviors are different because of differences in animation parameters. Verbal behaviors are different because of differences in the children's voices, intonations, and wordings.

Each puppet has three bipolar continuous moods: emotional (happy-sad), physiological (peppy-tired), and social (friendly-shy). Players can direct puppets' moods, but their moods also change in response to events. For example, a puppet will be happier if its invitation to play is accepted and sadder if its invitation is declined. A puppet will tire if it plays hard and pep up if it sleeps.

Each puppet maintains several other kinds of state. It uses state machines to track its interactions with its partner (e.g., greetings, game-playing). It represents the properties and locations of 30 destinations in the virtual world, as well as its own and its partner's dynamic locations. Directions are represented as constraints.

## 3. Story-Making with Improv Puppets

Players and Improv Puppets collaborate to improvise stories in real time. Players collaborate with one another as they do when playing with physical puppets--but now they direct their puppets' moods, actions, and words. Puppets collaborate with one another by perceiving and responding to one another's behavior. Each puppet's behavior might elicit an automatic response from the other, change its mood, or cause it to identify new relevant behaviors. Puppets and players collaborate with one another to determine the puppets' moods and behavior, by means of a mixed-initiative interaction through a real-time graphical user interface (see Figure 3): puppets offer directorial options; players choose among the available direction; and puppets improvise under the chosen directions.

Each puppet offers its player a dynamic set of directorial options corresponding to a small context-relevant subset of its full behavioral repertoire. Treating its user interface as a "window" into its mind, a puppet reveals how it feels--emotionally, physically, and socially--in the positions of its

sliders. It reveals what it is thinking about doing and saying as active blue and red buttons. It reveals where it is thinking about going as active areas of the virtual world. All of these "thoughts" and the corresponding interface displays change with the puppet's situation. For example, a puppet often thinks about going somewhere--but where and how it wants to go depend on its mood. A tired puppet will consider only nearby destinations and calm gaits, while a peppy puppet will consider distant destinations and energetic gaits. On the other hand, traveling to a distant destination with a peppy gait will tire a puppet more than traveling to a nearby destination with a calm gait. As a second example, a puppet that is playing alone ordinarily will want to continue, offering the player only one directorial option: to stop playing. However, if the other puppet invites it to play, the puppet will offer its player new directorial options to accept or decline.

Each player directs his or her puppet by choosing among current directorial options. He or she can move the sliders to change the puppet's mood, select a location in the virtual world to direct the puppet toward that destination, or select a blue or red button to direct the puppet to do or say something.

Each puppet immediately obeys its player's directions by improvising an appropriate course of behavior, coloring its improvisations with life-like qualities: normal variability, idiosyncrasies, mood-related modulations of behavior, event-based changes in mood, and adherence to social conventions.

For example, given a direction to move toward a specified destination, a puppet will move in the appropriate direction with a mood-appropriate gait. If the puppet feels peppy or the destination is close, it will follow a direct path to the destination. On different occasions, the puppet's choices of path and gait between the same origin and destination may vary within the ranges allowed by its mood. On the other hand, if the puppet is tired and the destination is far, it may stop before reaching the destination. A determined player can repeatedly select the same destination, directing the puppet to improvise a complete journey to the destination.

Taking a second example, a puppet will obey a direction to play alone by improvising a sequence of mood-appropriate physical and verbal behaviors. If the puppet's mood should change, autonomously or through direction, the quality of its improvisation also will change. The small puppet's sad play might involve quietly wobbling along the upper steps. Directed to quit, she might say, "All done," in a sad tone of voice. Her happy play might involve high hopping on the pedestals, leaping into the chute, and singing cheerful songs. Directed to quit, she might call out "Finishio!"

With the behaviors and moods currently implemented for Improv Puppets, players and puppets can improvise simple vignettes. More generally, the system illustrates the kind of improvisational story-making supported by the Improv Puppets paradigm: an intimate real-time

collaboration among all of the participating players and puppets.

## 4. Design Studies

We explored a few design alternatives for Improv Puppets.

An early user interface gave children access to puppets' complete repertoires of behaviors, including all executable behaviors, in hierarchically nested menus. This design allowed the children to control the puppets' specific behavior, instead of allowing the puppets to improvise their own course of behavior within the constraints of more abstract directions. We found two weaknesses in this design. First, puppet play should be immediate; it took too long to work through the menus. Second, the puppets' improvisations are cute and bring them to life; controlling them so tightly is not as much fun.

Another early interface was on a separate screen, with the two children viewing the puppets in their world on a third, shared screen between them. This design seemed to enhance the children's collaboration with one another, but to hinder their collaborations with their puppets. It is an interesting tradeoff and both designs seem to have attractive properties. We prefer the current design because it supports geographically distributed puppet play.

We also implemented a version of the puppets in which they have some behavioral autonomy, for example spontaneously returning greetings. The young children who user-tested the system seemed to be confused by the combination of behavioral autonomy and directability. Therefore, we eliminated behavioral autonomy in the current system.

## 5. Evaluation Studies

In ongoing research related to these systems, we are using the Improv Puppets system as a prototype environment for children's learning through creative play [Baker-Sennett, Matusov & Rogoff 1992, Flower & Hayes 1980, Malone & Lepper 1987, McCaslin 1987, Papert 1980].

In a small, informal study, we observed children playing with the puppets in pairs with peers, older siblings, or mothers [Huard & Hayes-Roth 1996a]. We found that mothers and older siblings offered advice to their preschool partners, whereas peers did not. Nonetheless, children in all pairings worked easily with their puppets and their partners to improvise stories without the need for "meta-level" verbal exchanges. That is, the kids collaborated through their puppets.

In a new study, we are investigating the Improv Puppets system's potential to facilitate children's development of a "theory of mind" [Huard & Hayes-Roth 1996b]. Children direct their puppets' moods and then observe the impact of mood on the puppets' directed improvisation. We hypothesize that this activity will lead children to form models of their puppets' minds and, by transfer, to form

and use models of other characters' minds. In particular, we hypothesize that they will bring this knowledge to bear in both the interpretation of behavior by characters in new stories and in the creation of richer characters in their own original stories.



## 6. Related Research

Although the Improv Puppets system was designed for children, many adults also find it engaging. An earlier version of the system was presented "live" at: the ACM Conference on Human-Computer Interaction (CHI) in May, 1995, where 250 people (two at a time) played with it [Hayes-Roth, Sincoff, Brownston, Huard & Lent 1995]; at the Stanford Arts and Technology Exhibit (SATI) in June, 1995, where about 50 people played with them; and in many private showings at Stanford.

Based on informal feedback from these exhibits, we speculate that adults may enjoy similar sorts of Improv Puppets applications adapted for suitably adult-oriented content. In fact, we have implemented a prototype system called "Cyber Cafe," in which player-directed Improv Puppets are customers in an on-line cafe where an autonomous Improv Actor serves as the waiter [Hayes-Roth & van Gent 1996]. Our prototype incorporates graphical embodiments of all characters using animations developed by Perlin [Perlin 1995]. We are extending the prototype with a Java-based user interface to create an on-line gathering place for visitors to our web site (<http://www-ksl.stanford.edu/projects/cait>).

We developed another system, called "While the Master is Away," in which the system organization and architecture of the mind described in this paper were used as the foundation for a pair of autonomous actors. Instead of taking direction interactively at run time, the actors accept directions organized in advance as a scenario. They can improvise alternative performances of a given scenario. This system is described in [Hayes-roth & van Gent 1996]. Other potential applications of Improv Puppets include: multi-player games [Hayes-Roth 1995c], avatars for on-line environments, character-based computer toys, interactive stories [Bates, Hayes-Roth & Maes 1995] high-level authoring environments for animated films [Parisi 1995], actors in simulated training environments, and characters for various kinds of on-line commercial functions.

## 7. Conclusions

Improv Puppets serves both as a testbed for basic research on synthetic agents and as a prototype for a new form of interactive systems.

As a research testbed, the system allows us to investigate fundamental issues in the design of synthetic agents: architecture and knowledge representation, real-time control of behavior, representation and mechanisms of mood, and life-like qualities of behavior. It also supports research on the design of user interfaces and their support for different forms of multi-agent human-computer interaction.

As a prototype, Improv Puppets illustrates a new interactive form. Its push buttons and sliders allow children of all ages to engage in easy social interaction, cooperative story making, character-based games, or free-form fantasy play with on-line companions. Compared to more conventional interactive systems for these kinds of interactions, Improv Puppets produces more interactive content for less user effort and less communication bandwidth. It allows participation by non-typists and non-readers. It produces characters whose behavior is life-like and engaging. These advantages are directly enabled by the system's incorporation of synthetic agents and their capabilities for performing directed improvisation.

## Acknowledgments

The work reported in this paper was supported by: ARPA Contract N66001-95-D-8642-Subcontract #137-1 through Teknowledge, Inc., NSF Contract IRI-9424261, a gift from Intel, and Seed Grants from Stanford's Center for the Study of Language and Information and its Office of Technology Licensing. The work has benefited from discussions with present and past members of the Virtual Theater research group, including: Lee Brownston, Erik Sincoff, Ruth Huard, Todd Feldman, Jack Jen-Jay Yu, Katherine Isbister, Daniel Rousseau, and Brian Lent. It also has benefited from discussions with Stanford Professors Larry Friedlander (English Literature) and Patricia Ryan (Drama). We thank Joseph Bates of CMU and Ken Perlin of NYU for allowing and assisting us in the use of their animation applications. We thank Aaron and Nora Hayes-Roth for creating the voices, lines, and personalities of two animated puppets.

A patent application has been filed for the "Method and System of Directed Improvisation by Computer Characters."

## References

1. Baker-Sennett, J. Matusov, E., and Rogoff, B. Sociocultural processes of creative planning in children's playcrafting. In P. Light and G. Butterworth (eds.), *Context and Cognition: Ways of Learning and Knowing*. NY: Harvester Wheatsheaf, 1992.
2. Bates, J., Hayes-Roth, B., and Maes, P. Working notes of the AAAI Symposium on Interactive Story. Menlo Park, CA: AAAI Press, 1995.
3. Flower, L., and Hayes, J.R. The dynamics of composing: Making plans and juggling constraints. In L.W. Gregg and E.R. Steinberg (eds.), *Cognitive Processes in Writing*. Hillsdale, NJ: Erlbaum, 1980.
4. Hayes-Roth, B. Agents on stage: Advancing the state of the art of AI. Proc. of the Int. Joint Conference on Artificial Intelligence, Montreal, 1995a.

5. Hayes-Roth, B. An architecture for adaptive intelligent systems. *Artificial Intelligence*, 329-365, 72, 1995b.
6. Hayes-Roth, B. Directed improvisation: A new paradigm for computer games. Proc. of the Computer Game Developers' conference, Santa Clara, CA: 1995c.
7. Hayes-Roth, B. Opportunistic control of action. *IEEE Transactions on Systems, Man, and Cybernetics*, 12, 1575-1587, 1993.
8. Hayes-Roth, B., and van Gent, R. Improvisational puppets, actors, and avatars. Proc. of the Computer Game Developers' Conference, Santa Clara, CA, March, 1996.
9. Hayes-Roth, B., Brownston, L. and van Gent, R. Multi-agent collaboration in directed improvisation. Proc. of the First Int. Conference on Multi-Agent Systems. San Francisco, 1995.
10. Hayes-Roth, B., Pflieger, K., Lalanda, P., Morignot, P., and Balabanovic, M. A domain-specific software architecture for adaptive intelligent systems. *IEEE Transactions on Software Engineering*, in press, 1994.
11. Hayes-Roth, B., Sincoff, E., Brownston, L., Huard, R., and Lent, B. Directed improvisation with animated puppets. Proc. of CHI '95 Conf. on Human-Computer Interaction. Denver, 1995.
12. Huard, R., and Hayes-Roth, B. Children's collaborative playcrafting. Stanford University: Technical Report, 1996a.
13. Huard, R., and Hayes-Roth, B. Development of children's theory of mind. Stanford University: Technical Report, 1996b. (NEW 3)
14. Johnstone, K. *IMPRO: Improvisation in the Theatre*. NY: Penguin Books, 1987.
15. Loyall, B., and Bates, J. Real-time control of animated broad agents. Proc. of the Cognitive Science Society, Boulder, CO, June, 1993.
16. Malone, T.W. and Lepper, M.R. Making learning fun: A taxonomy of intrinsic motivations for learning. In R.E. Snow and M.J. Farr (eds.), *Aptitude, learning and instruction*. Hillsdale, N.J.: Erlbaum, 1987.
17. McCaslin, N. *Creative Drama in the Intermediate Grades*. Studio City, CA: Players Press, Inc. 1987.
18. Papert, S. *Mindstorms: Children, computers, and powerful ideas*. New York: Basic Books. 1980.
19. Parisi, P. The New Hollywood: Silicon Stars. *Wired*, December 1995.
20. Perlin, K. Real-time responsive animation with personality. *IEEE Transactions on Visualization and Computer Graphics*, 1, 1995.