Multi-Agent Collaboration in Directed Improvisation

Barbara Hayes-Roth and Lee Brownston

Stanford University
October, 1994

Directed improvisation is a new paradigm for multi-agent interaction. One or more human users direct one or more computer characters with scripted or interactive directions. The characters work together to improvise a course of behavior that follows the directions, expresses their distinctive individual styles, reflects social principles, and meets other objectives. The resulting "performance" reflects the collaboration among all of the human and computer agents. Directed improvisation has several properties that make it a promising paradigm for multi-agent human-computer interaction. It also presents distinctive agent requirements that make it a useful addition to the domain of multi-agent paradigms. In this paper, we describe the directed improvisation paradigm, illustrate it in a "computer-animated improvisational theater," present an architecture for improvisational agents, and discuss the agent requirements that distinguish directed improvisation from other multi-agent paradigms.

1. Directed Improvisation

Directed improvisation is the simultaneous invention and performance of a new "work" under the constraints of user-specified directions. Besides its role in theater [7, 19, 21], directed improvisation mediates other diverse human activities: jazz music [9], writing [10, 25], scientific investigation [23], planning [13, 14, 16, 17], reactive behavior [1, 14], conversation [20], human-machine communication [24], children's planning and playcrafting [2, 3, 22], and life management [6]. Indeed, most human behavior and interaction appears to incorporate some degree of directed improvisation.

We are studying directed-improvisation as a new paradigm for multi-agent human-computer interaction. Here, the "new work" is a course of behavior enacted by computer characters. Users give characters abstract directions, either interactively or in preconceived scripts. The characters improvise a collaborative course of behavior that follows the directions, expresses their distinctive individual styles, reflects social principles, and achieves other performance objectives. Thus, the characters obey their users, but also surprise and engage them with artfully improvised behavior.

Directed improvisation has several attractive properties. (a) It provides a *collaboration framework* in which users and computer characters work together to achieve objectives [dai]. (b) It explicitly allows *run-time flexibility* in how characters meet objectives [1, 12, 14, 17]. (c) It provides a *natural, familiar style of interaction*. that is variable, idiosyncratic, unpredictable, and life-like [4, 5]. (d) It introduces a potentially *amusing, engaging, and delightful* quality to interactive experience through characters who combine task-relevant obedience with task-compatible improvisation [5]. (e) It extends the domain of multi-agent systems. In contrast to previously studied paradigms (e.g., cooperative problem solving, distributed work, discretionary cooperation [8]), directed improvisation requires: a greater emphasis on opportunistic behavior; a more complete collaboration among agents; and more process-oriented evaluation criteria.

2. Illustrative Application: A Computer-Animated Improvisational Theater

For illustration, we use CAIT, a computer-animated improvisational theater, in which children direct animated characters to perform stories and plays. Eventually, CAIT will support several creative, playful, and educational activities [18]. Here we focus on its use of directed improvisation.

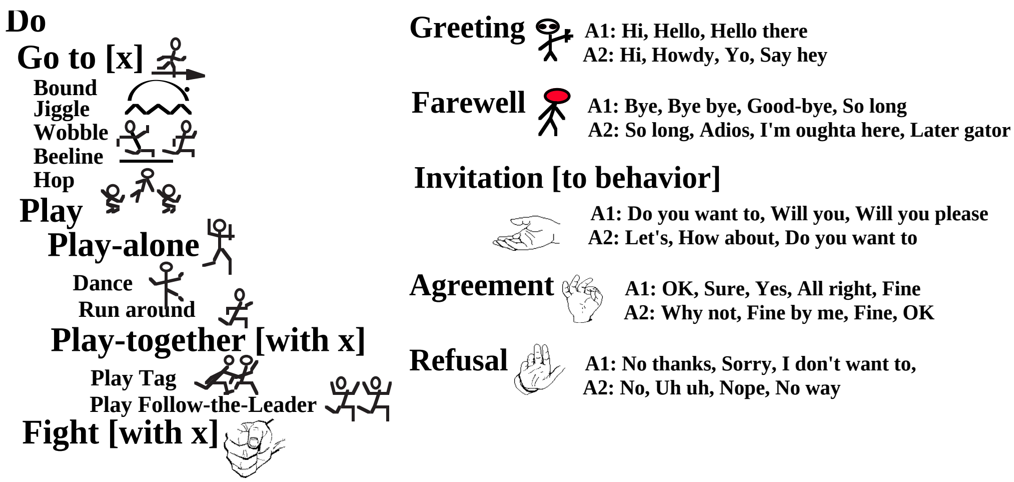

Figure 1. Excerpt from physical and verbal behavior repertoires for two agents, A1 and A2. Different physical characteristics affect their performance of common physical behaviors. Different personalities are reflected in different verbal behaviors.

Each character in the CAIT "company" has a repertoire of physical and verbal behaviors that serve as building blocks for improvisation (Figure 1). Instances of a class permit characters to express moods, individual differences, and general variability in nominally equivalent behaviors . For example, a character may go to a destination by "beeline" or "hop," depending on whether it "feels" determined or playful. CAIT currently has a single set and two characters (Figure 2), which we adapted from the "woggles" system developed by Joseph Bates at CMU. (See [18].) Each character has about 10 classes of physical behaviors and 20 classes of verbal behaviors, with 1-5 instances of each class. Their verbal behaviors—and thus, to some degree, their personalities—were conceived and recorded by Aaron Hayes-Roth, age 13, and Nora Hayes-Roth, age 10. To create stories, children direct characters to perform sequences of behavior classes. The characters improvise within those constraints. We have conceived several different interaction modes, but focus here on two modes.

In *animated-puppets mode*, children direct characters interactively using "situated behavior menus." Each menu displays a few abstract behaviors that its character "is considering now." In Figure 2, the large character on the left feels cheerful and energetic. It considers going somewhere, playing alone, or inviting the small character to play. The small character feels OK, but tired. It considers going to the rest area, playing alone, or speaking. Children direct characters by selecting menu options. In Figure 2, one child directs the large character to invite the small one to play. It must improvise an appropriate sequence of executable behaviors, for example: approach the small character, make a greeting, and say "Do you want to play follow the leader?" A character's improvisational choices reflect its personality, its relationships, and its recent interactions with other characters. Executed behaviors produce changes in the shared situation and, therefore, changes in which behaviors the characters subsequently consider. Either child may direct his or her character by choosing menu options at any time. The characters improvise both directed and undirected behaviors (e.g., returning one another's greetings even without being directed to do so).

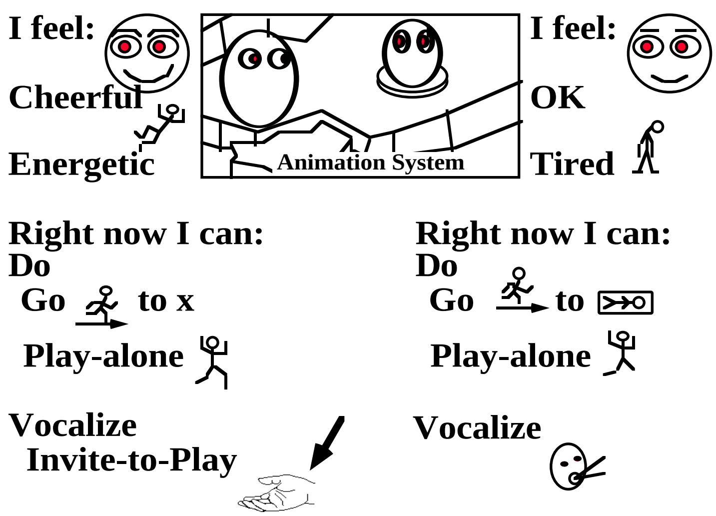

Figure 2. Snapshot from situated behavior menus for a two-character play.

In *animated-actors mode*, children direct characters with synchronized "behavior scripts." Characters work through a script, improvising both scripted and unscripted behavior sequences by choosing among dynamically triggered situated behaviors. During the play-alone scene in Figure 3, each character could dance, run around, hop on the pedestal, etc. They might also greet one another in a chance encounter. The character A2 determines when "awhile" has passed and what game to suggest (e.g., hide and seek, tag, follow the leader). A2 agrees and the two characters collaborate on playing the game. Again, the characters' behavioral choices are influenced by their personalities and interpersonal histories. Children can override the characters' autonomous choices at any time.

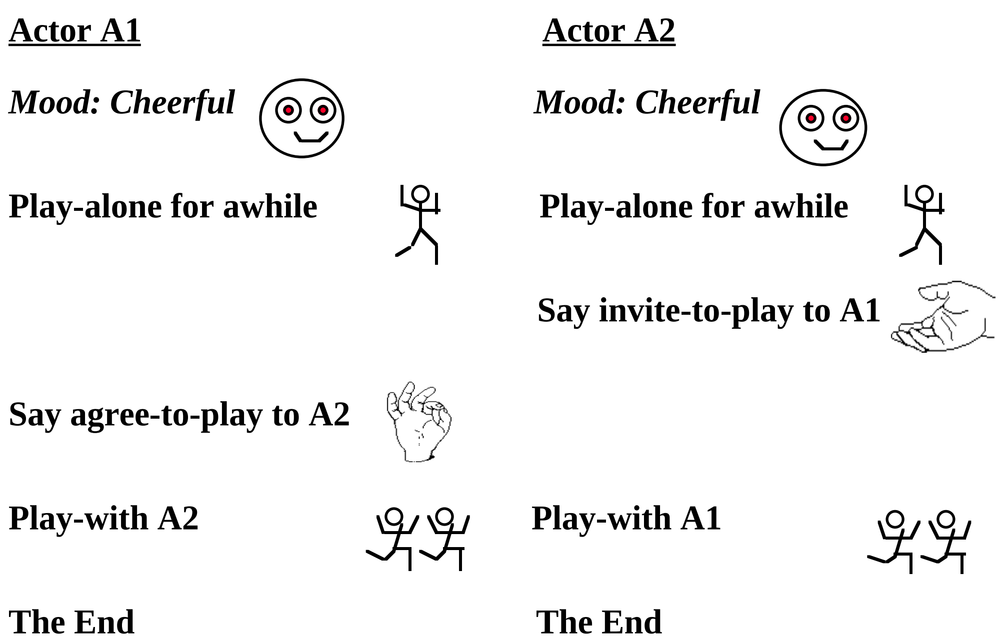

Figure 3. Illustrative Behavior Script.

CAIT illustrates multi-agent human-computer collaboration in the directed improvisation paradigm. Children work together to create the narrative structures of their stories. Children direct the behavior of characters within the narrative structure. Characters collaborate with one another to improvise the details of their interactions. The balance of creative work shifts among children and characters depending on the specificity of the children's scripts and directions. If they wish, the children can determine very precisely each successive behavior their characters perform. Here, however, we focus on the case in which characters have considerable freedom to improvise.

3. An Architecture for Improvisational Actors

Each animated character in a production embodies an intelligent agent that mediates all of its interactions with users, other characters, and the environment (Figure 4). Full and restricted versions of the "dynamic control model" serve as cognitive and physical controllers in a two-level agent architecture [11, 12, 13, 14, 15, 17]. An agent's cognitive controller receives perceptual inputs from its physical controller, constructs an evolving model of its situation and its interactions with other characters, plans its physical and verbal behaviors, and sends control plans to its physical controller. An agent's physical controller interprets and filters inputs from its user interface and from sensors on its animated embodiment and sends those perceptions to its cognitive controller. It receives control plans from its cognitive controller and enacts those plans with physical behaviors, sending their outputs to to effectors on its animated embodiment and to its user interface.

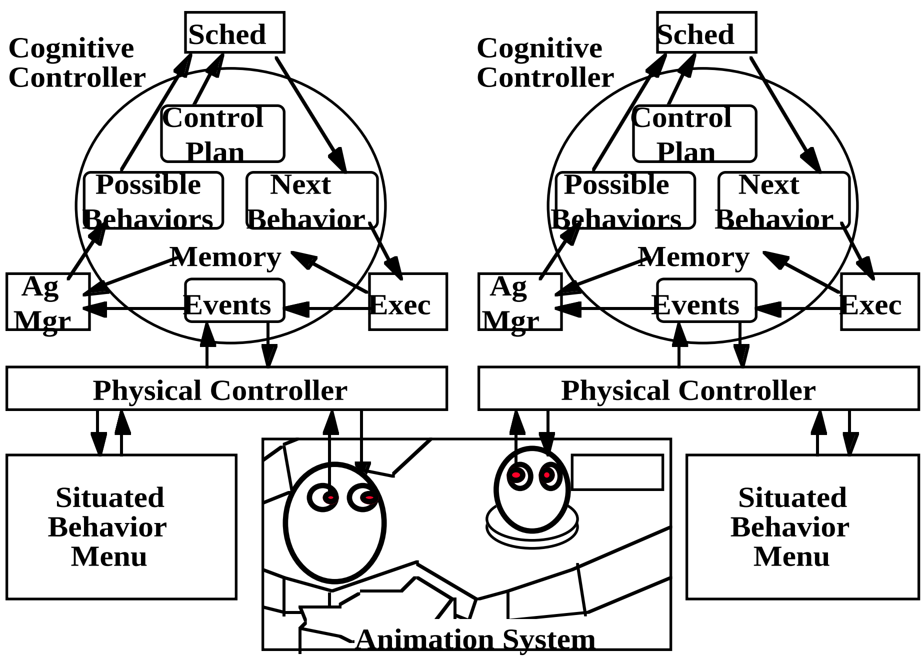

Figure 4. Framework for a Two-Character Production.

Each of an agent's controllers iterates its own three-step cycle: (a) its agenda manager triggers behaviors for the current situation; (b) its scheduler chooses the best triggered behavior based on the current control plan; and (c) its executor executes the chosen behavior. Behaviors at each level are represented declaratively as illustrated in Figure 5. Because control plans are abstract, they typically allow execution of alternative behavior sequences. For example, a character could satisfy a cognitive control plan to "interpret partner's movements" with either deductive or case-based reasoning behaviors. It could satisfy a physical control plan to "go to destination-d" with different physical gaits. Because control plans are data structures, a character can modify them at run time. For example, observing a change in a partner's mood, a character might change its cognitive control plan from "interpret partner's movements" to "plan a response to partner's mood."

Cognitive Behaviors
- Interpret (Other-Agent's-Behavior)
  - Deductively-Interpret
  - CBR-Interpret
- Plan-response-to (Other-Agent's Mood, Other-Agent's Gesture)

Physical Behaviors
- Go-to (Destination)
  - Hop-to
  - Beeline-to
  - Wobble-to

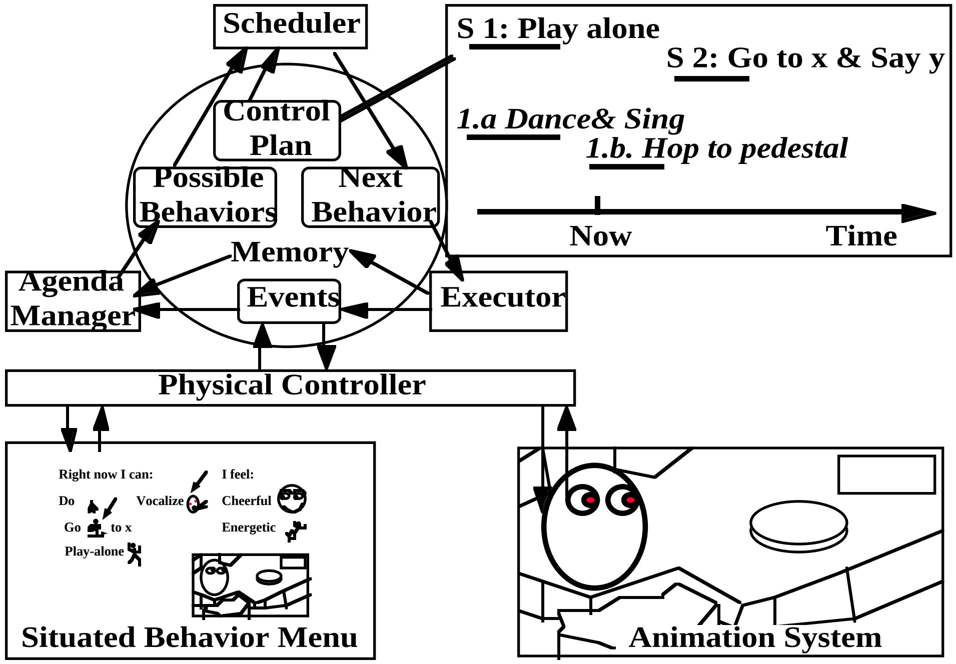

Figure 5. Dynamic Control Plans Integrate Users' Scripted and Interactive Directions with Characters' Improvisational Decisions (in *italics*).

Our architecture is a natural framework for directed improvisation. The agenda manager generates situated behaviors—the raw materials out of which a character creates meaningful behavior patterns. Dynamic control plans, which the scheduler uses to make context-dependent choices of behaviors to execute, provide an ideal representation for users' scripted and interactive directions. They can represent detailed directions, which effectively specify all physical and verbal behaviors and severely restrict improvisational oportunity, or abstract directions, which minimally constrain behavior and permit more improvisation. Dynamic control planning enables characters to improvise the details of users' directions and to augment them with unscripted behaviors on the fly. In Figure 5, a character follows an abstract direction to "play alone" by improvising a sequence of solo games. Table 1 describes improvisational forms that vary in behavioral complexity and the involvement of partners. We are developing cognitive behaviors to instantiate these and other improvisational forms as dynamic control plans.

Table 1. Improvisational Forms.

Solo Improvisation	*Collaborative Improvisation*
One-Step Improvisation	
Choose among alternative logically equivalent behaviors Direction: Go to pedestal Improvise: Hop to pedestal-3	Respond to a partner's behaviors Partner: Greeting Improvise: Return greeting
Sequential Improvisation	
Construct a coherent path to a dramatic moment Direction: Play alone, Rest Improvise: Play alone, Get tired, Rest	Recognize and coordinate with a partner's behavior sequence Partner: Going toward pedestal Improvise: Meet at pedestal
Patterned Improvisation	
Instantiate an improvisational schema Direction: Dance: Improvise: Iterate(Hop, twirl)	Recognize and participate in a partner's schema Partner: Play hide and seek? Improvise: I count to 10, etc.

To support multi-agent collaborative improvisation, we are replicating the architectural mechanisms agents use to represent, plan, monitor, and control their own behavior so that they can apply them to their partners' behavior. Figure 6 illustrates undirected collaborative improvisation. The large character is on the set and has made a plan to look for something. The small character enters, observes its partner standing still, and infers nothing about its mood or plans. The small character plans to "start something" by playing alone for awhile and then hiding. Meanwhile, the large character observes its partner enter and start playing alone and infers that it is portraying a "shy" character. The large character plans an appropriate interaction: approach, greet, and invite the shy character to play. At this point, the two characters have unknowingly constructed conflicting plans. One of them must change its plan and probably its model of its partner. However, it doesn't matter which character changes, as long as their observable interactions appear plausible. Thus, if the large character succeeds in inviting the small character to play, the small character must abandon its plan and agree to play. If the small character notices the large character approaching and rushes over to tell it to hide, the large character must abandon its plan and hide. The architecture enables characters to generate and adapt their models of one another and their plans for their own behavior as their interaction unfolds.

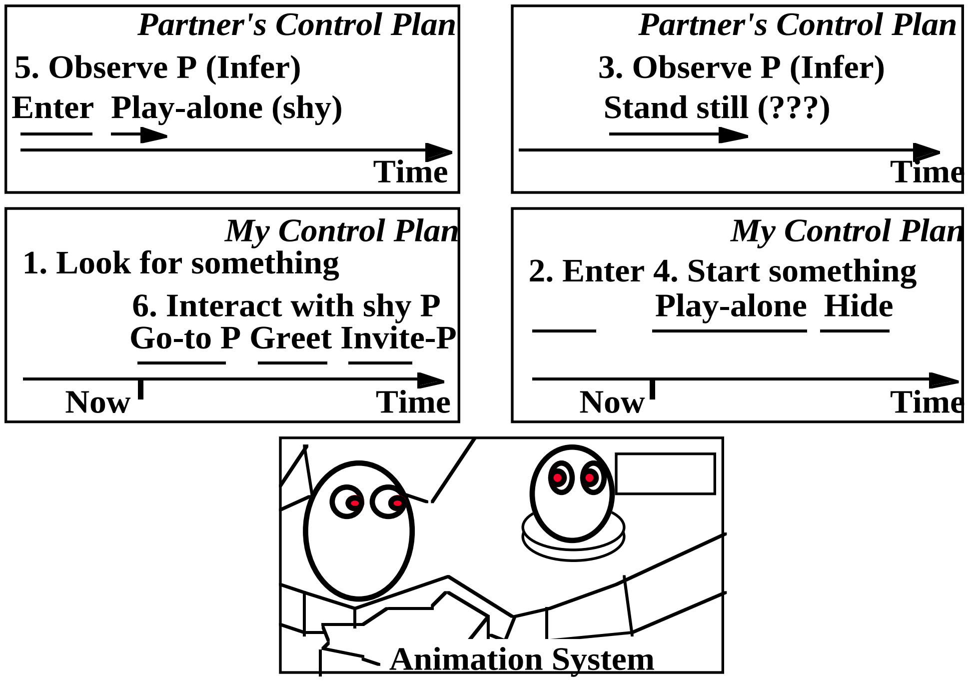

Figure 6. Simple Illustration of Collaborative Improvisation.

4. Issues in Multi-Agent Interaction

Directed improvisaation is a new interaction paradigm for multi-agent systems. Although it shares important requirements and solution techniques with previously studied paradigms (e.g., cooperative problem solving, distributed work, discretionary cooperation [8]), it presents three distinctive requirements (discussed below) that make it a useful addition to the domain of inquiry. We view these requirements as design objectives in our own efforts to elaborate and instantiate our proposed agent architecture.

Directed improvisation emphasizes behavior that is situated, spontaneous, and opportunistic in the service of abstract and weakly constrained goals.

Most multi-agent paradigms assume that agents' individual and collective activities are predominantly and specifically goal-directed. In cooperative problem solving, the team's goal is to solve a shared problem. In distributed work, the team's goal is to get all of the work done. In discretionary cooperation, individual agents have their own goals. Techniques for these paradigms focus on goal-directed reasoning in support of individual and group planning. Unplanned behavior occurs only when it is necessary to compensate for plan failure of when it happens to advance planned efforts to achieve established goals.

The directed improvisation paradigm also assumes that agents have goals, but they typically are less well defined and less constraining than in other paradigms. Agents working in a joint production share a general goal to produce a successful joint performance that meets the constraints of the users' directions. Although directions can vary in specificity, effective application of the paradigm typically involves relatively abstract directions that only weakly constrain the performance and give characters plenty of freedom to improvise. The working assumption is that producing the directed behaviors is easy; the art lies in the improvisation.The performing arts literature [7, 19, 21] reflects this assumption in its prescription of two underlying cognitive heuristics for good improvisation: (a) *welcome possibilities* (Just let the words flow. Do not fear mistakes. Turn off the censor. Look for relationships. Do not plan too far ahead.); and (b) *pursue promising possibilities* (Relate present actions to past actions. Keep the action on stage. Make the natural response. Listen and respond to your partner. Take it to the extreme. Accept (don't block) offers). In contrast to the foward-looking, goal-directed reasoning and planning of agents in traditional paradigms, effective improvisers engage in backward-looking efforts to *reincorporate* incidental themes and behavioral qualities that they or their partners happen to have generated previously. Professor Patricia Ryan, who teachers improvisation at Stanford, describes the improviser at work as "a man walking backwards, trying to make sense of where he's been" (personal communication). In sum, the individual agent's behavior is firmly situated in the dynamic context, spontaneous in its short-term etiology, and opportunistic in its thematic relationships to other aspects of the performance. Ensemble behavior builds incrementally out of individual agents' actions and reactions. Achievement of the "goal" is not the specific product of a deliberate, provably correct process, but an emergent and uncertain epiphenomenon of the agents' real-time interactions.

Directed improvisation demands intimate collaboration and shared control among agents engaged in closely intertwined and interdependent behaviors.

Most multi-agent paradigms assume that agents have limited interest in working together and limited interactions during the actual performance of their collective activities. In cooperative problem solving and distributed work, agents work together because they are committed to solving a shared problem or to completing a shared job. The prevailing interaction model is divide-and-conquer the joint task so that individual agents can work more or less independently. Multi-agent planning is used to allocate and coordinate tasks among agents, with run-time communication used primarily to exchange selected results of individual activities and, in some cases, to reallocate responsibilities. In discretionary cooperation, individual agents have their own goals and may or may not be willing to work together at all, depending on what it costs them. Competition for resources or conflict among goals may impede cooperation or even motivate agents to thwart one another's efforts, raising issues of trust and deception. Communications focus mainly on determining agents' willingness to cooperate. In all of these paradigms, research concentrates on the key question of how agents should decide in advance who is able and willing to do what. Techniques have been developed for: problem decomposition, communication of assumptions, beliefs, decisions, and commitments; negotiation and persuasion; conflict resolution and consensus building; organization of effective chains of command; and establishing mutually beneficial social laws.

By contast, the directed improvisation paradigm assumes that agents are 100% committed to collaboration on their joint performance and that they will work together every step of the way. Their prospective work is so intertwined (and unpredictable) that agents do not even think of dividing it up ahead of time. Instead, effective improvisers rely on one another to do what is necessary to generate the work jointly and interactively in real time. In fact, none of the participating agents can make individual progress without collaborative interaction with the others. As illustrated above for CAIT, the work-in-progress is intrinsically collaborative. Individual agents may introduce plot devices or instantiate improvisational routines. They may build small individual plans of activity in particular situations. However, all such structures, routines, and plans are tentative and dispensable under the fundamental rule of collaborative improvisation: *Accept all offers* [19]. An "offer" is any explicit or implicit assertion. No matter how an individual agent feels about a partner's offer, no matter where the individual had been planning to go in the performance prior to the offer, no matter how the partner's offer redirects the individual's behvior, there can be only one response: "Yes." An effective improviser always embraces a partner's offers and tries to advance them with constructive, collaborative behavior. Since all of the participants are allowed to make offers and since not all offers are detected or understood by partners, each improviser must be willing to both lead and follow, with no preconception of when, how often, or how long. The dynamism and mutual adaptation of good improvisers reflects this underlying willingness to share control. Unlike control regimes based on organization, negotiation, turn-taking, or other explicit arrangements, shared control is an implicit arrangement in which the participants readily and immediately adopt and contribute to one another's goals, strategies, and tactics so that they can move forward together.

Directed improvisation succeeds when it produces a joint performance that follows the script and directions in an engaging manner.

Most multi-agent paradigms are "product-oriented"—they evaluate the objective consequences of agents' behavior against high standards: achievement of agents' individual and collective goals, optimal use of resources, and acceptability of side effects. These criteria for evaluating group performance imply related criteria for evaluating individual performance. To succeed as a member of the group, an agent must: reason correctly in performing its own tasks, model its partners' knowledge and behavior correctly, make rationale decisions about commitments to cooperative behavior, keep its commitments or at least inform its partners of changes, etc. Features of agents' behavior that do not affect the objective consequences are not valued, may be distracting, and, in the worst case, may carry unacceptable costs.

By contrast, the directed improvisation paradigm is "process-oriented"—the joint course of behavior enacted by computer characters <u>is</u> their product.

Other than meeting the constraints of users' directions (which is usually assumed to be easy), there is no "correct" or "incorrect" performance. On the contrary, there can be many alternative, equally "successful" performances of a given script—that is, performances that follow the directions in an engaging manner. In fact, in domains like our computer theater, children (and adults) may find it especially charming to observe or participate in repeated performances of a favorite story or one they have created themselves just to see how the characters will improvise anew. By implication, instead of behavior that is correct, rational, and reliable, effective improvisers produce behavior that appears appropriate in context, varies in different performances, and, in the best case, is endearingly idiosyncratic. The individual qualities that agents bring to a production are not costly distractions, but powerful sources of texture and depth in their contributions to the joint performance. In contrast to the all-business mentality of the ideal agent in traditional multi-agent paradigms, effective improvisers bring believable characters to life.

1. Agre, P.E., and Chapman, D. What are plans for? Robotics and Autonomous Systems, 6, 17-34, 1990.

2. Applebee, A.N. The User's Concept of Story. Chicago: University of Chicago Press, 1978.

3. 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. New York: Harvester Wheatsheaf, 1992.

4. Bates, J. The role of emotion in believable agents. Communications of the ACM, 1994, 37, 122-125.

5. Bates, J., Hayes-Roth, B., Laurel, B., and Nilsson, N. Working notes of the AAAI Symposium on believable agents, Menlo Park, CA: AAAI, 1994.

6. Bateson, M.C. Composing a Life. New York: The Atlantic Monthly Press, 1989.

7. Belt, L., and Stockley, R. Improvisation through Theater Sports. Seattle: Thespis Productions, 1991.

8. Charib-draa, B., Mandiau, R., and Millot, P. Distributed artificial intelligence: An annotated bibliography. SigArt Bulletin, 3,1992, 20-37.

9. Dean, R. Creative Improvisation: Jazz, Contemporary Music and Beyond. Philadelphia: Open University Press, 1989.

10. 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.

11. Hayes-Roth, B. A blackboard architecture for control. Artificial Intelligence, 1985, 26, 251-321.

12. Hayes-Roth, B. An architecture for adaptive intelligent systems. Artificial Intelligence, in press, 1994.

13. Hayes-Roth, B. Architectural foundations for real-time performance in intelligent agents. Real-Time Systems, 2, 99-125, 1990.

14. Hayes-Roth, B. Opportunistic control of action. IEEE Trans. on Systems, Man, and Cybernetics, 12, 1575-1587 1993.

15. Hayes-Roth, B., and Collinot, A. A satisficing cycle for real-time reasoning in intelligent agents. *Expert Systems with Applications*, 7, in press, 1993.

16. Hayes-Roth, B., and Hayes-Roth, F. A cognitive model of planning. Cognitive Science, 1979, 3, 275-310.

17. Hayes-Roth, B., Pfleger, K., Lalanda, P., Morignot, P., and Balabanovic, M. A domain-specific software architecture for adaptive intelligent systems. IEEE Trans. on Software Engineering, in press, 1994.

18. Hayes-Roth, B. Sincoff, E., Brownston, L., Huard, R., and Lent, B. Directed improvisation. Stanford University: Knowledge Systems Laboratory Technical Report KSL-94-61, 1994.

19. Johnstone, K. IMPRO: Improvisation in the Theatre. New York: Penguin Books, 1987.

20. Moore, J. Invited presentation at the National Conference on Artificial Intelligence, 1994.

21. Nachmanovitch, S. Free Play: Improvisation in life and art. Los Angeles: Jeremy P. Tarcher, Inc. 1990.

22. Rogoff, B., Gauvain, M., and Gardner, W. The development of children's skills in adjusting plans to circumstances. In S. Friedman, E. Scholnick, and R. Cocking (eds.), Blueprints for Thinking: The role of planning in psychological development. New York: Cambridge University Press, 1987, 303-20.

23. Stefik, M. Planning and meta-planning (MOLGEN: part 2) Artificial Intelligence, 16, 1981, 141-169.

24. Suchman, L. A. Plans and Situated Actions: The Problem of Machine/Human Communication. Cambridge, England: Cambridge University Press, 1987.

25. Winokur, J. Writers on Writing. Philadelphia: Running Press, 1990.