

**This Is a Publication of
The American Association for Artificial Intelligence**

This electronic document has been retrieved from the
American Association for Artificial Intelligence
445 Burgess Drive
Menlo Park, California 94025
(415) 328-3123
(415) 321-4457
info@aaai.org
<http://www.aaai.org>

*(For membership information,
consult our web page)*

The material herein is copyrighted material. It may not be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from AAAI.

Autonomous Agents as Synthetic Characters

Clark Elliott and Jacek Brzezinski

■ Humans are social creatures. Much of our intelligence derives from our ability to manipulate our environment through collaborative endeavors. Most extant computer programs and interfaces do little to take advantage of such manifestly human talents and interests, leaving broad avenues of human-computer communication unexplored. Although it is still considered controversial, there are many who believe the harnessing of social communication to be rich in possibilities for modern software. In this article, we look at a number of autonomous agent systems that embody their intelligence at least partially through the projection of a believable, engaging, synthetic persona. Among other topics, we touch briefly on samples of research that explore synthetic personality, representations of emotion, societies of fanciful and playful characters, intelligent and engaging automated tutors, and users projected as avatars into virtual worlds.

The idea of machines with humanlike qualities has been with us since the advent of science fiction. In recent years, we have seen an explosive growth in the power of inexpensive multimedia computers that can support interfaces that look and sound very much like humble versions of these fanciful autonomous machine-characters of years past. Software agents that run on these platforms and that make use of both deep reasoning about humanlike qualities and multimedia techniques to manifest them might, some day, make fundamental changes in the way we relate to our computers. Here we look at a number of current approaches to building synthetic characters that embody these techniques.

Agents of Change— The Emerging Power of Personified Interfaces

The central theme of this article is that synthetic agents are attractive to us because they communicate in ways seldom before used in extant computer applications and that the novel portion of this communication is primarily **social** in nature. With most software currently in use, humans, being both smarter and more flexible than computer systems, have shouldered the burden of adapting to the artificial communication protocols established between users and software programs. Preferring speech, we type; preferring conversation, we give commands; preferring negotiation, we get no negotiable feedback. However, people are social animals. They have highly developed social skills that are not typically being used to direct, and gather information from, their computer applications. Progress in reducing the adaptive requirements of human users by increasing the ability of software to communicate within the human social context is a big win. It is this area where agent-based software, especially as embodied in agents with redeeming social qualities, is particularly attractive.

Some Open Areas of Broad Interest in Synthetic Agent Research

Among the most difficult problems will be the development of widely applicable authoring tools that bridge the gap between theoretical foundations laid in the research labs and the ad hoc intentions of system designers. In some larger projects, this goal is made explicit (for example, the JACK Project at the University of Pennsylvania [Badler, Phillips, and Webber 1993], the PPP [personalized plan-based presenter] Project at DFKI [Andre, Rist, and Muller

1998], and the STEVE and ADELE projects at the University of Southern California Information Sciences Institute [USC-ISI] [Johnson and Shaw 1997; Rickel and Johnson 1997]); in virtually all others, it is at least implicit. Even when a strong theoretical foundation gives rise to a system, it is often the artistic detail that makes the system work. Systems such as Cyberlife's CREATURES (Grand, Cliff, and Malhotra 1997) and North Carolina State University's (NCSU) COSMO (Lester, Callaway, et al. 1997; Lester, Converse, et al. 1997b) require long hours in creating personas that are fun, interesting, and informative. Authoring tools that make use of relevant metaknowledge of the architectures, and past artistic design, will go a long way toward reducing the time necessary for building new agents. For example, an authoring tool for PPP will need to understand something about temporal planning constraints, but one for COSMO will need to know about sets of overlapping behaviors. Similarly, authoring tools applicable to agent personalities such as those proposed by Elliott, Reilly, Velasquez, Botelho, Rousseau, and Sloman (Botelho 1997; Rousseau and Hayes-Roth 1997a; Sloman 1997; Velasquez 1997; Reilly 1996; and Elliott 1992) will allow for the development of lay-personality characterizations from the rich formal personality systems developed by these labs.

The automated real-time control of music stands out as a unique opportunity for synthetic agent research. Music is at once both emphatically human in quality yet not at all natural in real-world social interaction. It can be profoundly engaging, highly communicative in a subjective way, and extremely powerful as a mood-manipulation tool. The entertainment industry has long recognized music's power to engage us, yet although there certainly is work under way to understand computer-delivered music's effects, on the whole, its effectiveness as a computationally indexable resource for agents seems woefully underused. Often, it is true that interface efforts focus almost exclusively on graphic representations, but it is not at all clear that this area is where the biggest payoff lies for equivalent effort. Skeptics of the power of music to convey important social information should note how exceedingly rare it is for a Hollywood movie, on which hundreds of millions of dollars can ride, to be successful without a quality musical score.

Although there are extant computational models for some forms of humor, theoretical models of humor for proactive use in agents seem at best ad hoc. Given the story-telling

capabilities of agents, their emerging social intelligence, and their rudimentary emotional intelligence, this seems a natural course to pursue—and one with a wide applicability. Furthermore, it is generative models of humor—certainly easier than humor understanding (for example, see Zrehen and Arbib [1998])—that would be most immediately useful in the synthetic agent paradigm.

The integration of past work in natural language understanding, and discourse understanding, with modern speech-recognition software seems a likely “killer app” for lifelike agent research. Similarly, natural language generation and text-to-speech technology seem made to order for extending the capabilities of many existing agents.

Interactive agents for the web are certainly going to be popular and will undoubtedly be a major contributor to rapid growth in autonomous agents research. It is not clear that assumptions of greatly increased network bandwidth, at least of the sort that will support network-based real-time animations in the near future of agent research, are warranted. Systems whose intelligence can be focused through the existing narrow bandwidth of existing home telephone lines will have a big advantage in accessibility. This constraint suggests strong considerations with respect to local execution within the context of web browsers (for example, as in ISI's ADELE and DFKI's PPP) or through the use of efficient mechanisms for carrying the “socially intelligent” signal (as in DePaul's AFFECTIVE REASONER).

Finally, building lifelike agents, especially those that mimic human emotion and personality, might well improve our ability to model and detect similar states in users. At some level, the illusion of life in agents breaks down without an ability to form concepts about, and respond to, the state of the user (see the discussions on affective user modeling in the AFFECTIVE REASONER and in the work of Blumberg).

After Artificial Life...Artificial Death? Questions to Ponder

Synthetic agent characters work because people see them as social counterparts. When bonds are formed, what happens when these creations fail to act in socially responsible ways? Furthermore, is such a goal even achievable? After all, people will themselves generally fail this test!

We might find it possible for synthetic agents to form synthetic yet highly plausible relationships with users and with each other; indeed, this is indirectly a goal of building believable agents. However, with respect to

users, the concept of a synthetic relationship might well be spurious, and such a relationship is qualitative in nature at the least. As agents are better able to create the illusion of artificial life, the social bond formed between agents, and the humans interacting with them, will grow stronger. New ethical questions arise. Each time we inspire an agent with one or more lifelike qualities, we muddy the distinction between users being amused, or assisted, by an unusual piece of software and users creating an emotional attachment of some kind with the embodied image that the lifeless agent projects.

Especially so for children but also for adults, it is clear that agents that understand social relationships, maintain histories with users, have some knowledge of human emotion, are beginning to understand human speech, can speak themselves, and have control over media channels to deliver morphing faces, music, and theater-quality sound—all responsively and in real time—have tremendous inherent attachment-forming capabilities. Additionally, these agents, by definition, are at least partially autonomous. They might well live on after the user walks away from the terminal and might form relationships with other users. In short, they have their own (albeit impoverished) synthetic lives. What happens to the agent, and how it changes over time, might not be something designed for the user to fully control. When wedded with adventure games, goal-based products such as tutoring systems where a user's well-being can be seen as depending on his/her relationship to the embodied agent or systems that deliver critical or personal information (such as might occur in, for example, a medical patient advocate system [see the related Miksch, Cheng, and Hayes-Roth [1997]], the possibilities for as-yet-undiscovered social phenomena are wondrous (frightening?) to ponder (Foner 1997).

Will You Still Like Me?

As an example of the sorts of social sophistication that agents are beginning to manifest, consider this sample dialog (figure 1) from the Affective Reasoning Project at DePaul University, which took place between a researcher, Elliott, and an affective reasoning agent, SAM (figure 1). The agent was embodied as a morphing schematic face (with a wide range of emotion expression), used speech recognition for input and text-to-speech for dialog output, and played musical selections to enhance the expression of emotion. Although not interesting from a natural language-understanding point of view (the agent has no understanding

about anything except emotion states), it does illustrate the elements of history maintenance, some sophistication in social awareness, and the rudiments of understanding about the social structure of emotional interaction. (The dialog has been cleaned up minimally to filter out speech-recognition miscues.)

Elliott: I am really furious at Joe.

SAM: Boy what a blameworthy thing he must have done! What do you call such an action?

Elliott: Joe was stealing.

SAM: When Joe was stealing, he must have caused you distress. What is the name of the thing you wanted, or wanted to preserve?

Elliott: My sandwich.

SAM: How important was your sandwich to you, Clark?

Elliott: Not very important.

SAM: How surprising was it that Joe was stealing and caused you distress about your sandwich?

Elliott: It was not surprising.

SAM: How bad was it that Joe was stealing?

Elliott: It was not very bad, actually.

SAM: How important is your principle about stealing?

Elliott: Quite important.

SAM: Hmmm. I would think that you are only somewhat annoyed at Joe. I do not understand that you are furious at Joe. After all, what Joe did was not that bad. Your sandwich is not that important to you. It is not very surprising that Joe was stealing. Perhaps your principle about stealing is very important to you. Is this correct?

Elliott: Yes.

SAM: I will remember that about you.

... later ...

SAM: Clark, if I was stealing and it affected your sandwich, would you still be my friend?

In this sequence, SAM “believes” that anger is caused when some entity performs a blameworthy act that affects one’s goals. SAM uses dialog to discover what the blameworthy act was and what the goal was and then asks for information about variables that can affect the intensity of anger. Because these variables do not match a general pattern for fury, SAM reasons that one explanation for this anomaly might be the greater weight of the principle

Especially so for children but also for adults, it is clear that agents that understand social relationships, maintain histories with users, have some knowledge of human emotion, are beginning to understand human speech, can speak themselves, and have control over media channels to deliver morphing faces, music, and theater-quality sound—all responsively and in real time—have tremendous inherent attachment-forming capabilities.




Figure 1. Real-Time Dialog between Elliott and SAM, an Affective Reasoning Agent, Using Speech Recognition and Text to Speech.

about stealing and asks Elliott about this. After confirmation, SAM updates his internal representation of what is important to Elliott. Later, SAM draws on this information to ask Elliott an intelligent question about what might happen if their friendship were at odds with a strongly held principle.

A Sampler of Research Paradigms for Lifelike Synthetic Agents

In many cases, one of the key requirements of useful lifelike agents is that they be engaging. There are a surprising number of disparate ways to effect this, such as inspiring the agents with realistic models of emotion and personality, giving them a sense of humor and social grace, or creating a believable social fabric in interesting virtual worlds. In this section, we introduce a sampler of agents that manifest such engaging qualities in a variety of ways.

Agent-Based Models of Emotion and Personality

The building of computer emotion models has long been viewed with a certain amount of skepticism; after all, even deciding what comprises emotion in humans is rife with controversy. Whispering a phrase such as *basic emotions* in the wrong company is liable to cause verbal fireworks, if not actual fisticuffs. However, it is clear that computer agents are capable of

expressing a rich variety of emotions (Elliott 1997b; Koda 1997) and that these emotions can be effective in conveying much about the intended personality of the agent. Few would argue that some working model of emotion systems will greatly benefit both the believability and the intuitiveness of automated characters.

The first problem in sorting out the field is that work labeled as centering on effect, emotion, or personality tends to be grouped together out of hand regardless of the work's focus and goals, although these vary widely. In some cases, the researchers are attempting to faithfully model emotion subsystems (if not architecture), both human and otherwise; in other cases, they are attempting to support a believable social structure based on a descriptive model of eliciting scenarios and response behaviors. Rosalind Picard of the Massachusetts Institute of Technology (MIT) Media Lab has a good starting list of references on affective computing (starting with her new book, *Affective Computing* [MIT Press, November 1997]), and many of these have relevance to the intelligent, interactive agent paradigm (vismod.www.media.mit.edu/vismod/demos/affect/AC_bibliography.html). Here we consider only three of the many labs that are primarily dedicated to building computer models of emotion and personality appropriate for use in computer agents; also see the work of Blumberg, Koda, and Maes from the Affective Computing Group at MIT, the work of the VIRTUAL THEATER Project at Stanford University's Knowledge Systems Laboratory (KSL), and the work of Sloman and Humphreys at the University of Birmingham (www.cs.bham.ac.uk/~axs/cog_affect/sim_agent.html).

Affective Reasoning Is Effective Reasoning In our own work on the Affective Reasoning Project, we have used a descriptive model of emotion based on the seminal work of Ortony, Clore, and Collins (1988). The AFFECTIVE REASONER has been agent-centric since its inception in 1990. The model is manifested in, albeit entirely independent of, network-efficient multimedia agents that have highly expressive schematic faces, speak with somewhat emotionally inflected voices, listen to users through speech recognition, and use a rich set of musical selections to help reflect their current states. The agents have 26 emotion types along with a rich set of variables for controlling intensity. A key element of the AFFECTIVE REASONER agents, which run on a PC platform, is that they respond in real time to input, or lack thereof.

Recent work in the project has several branches. In one study presented at the

Autonomous Agents '97 Conference, we reported results showing that subjects did significantly better at correctly matching videotapes of computer-generated multimedia AFFECTIVE REASONER presentations with intended emotion scenarios (70 percent) than they did with videotapes of a human actor attempting to convey the same scenarios (53% $\chi^2(1, N = 6507) = 748.55, P < .01$) (Elliott 1997b). Verbal information was ambiguous, such as, "I saw Butler in the news again today," and users had to match intended meanings with as many as 12 different possibilities, such as, "Wanda is angry because Butler, one of her subordinates, is again saying damaging things about her in the news" and "Wanda is fearful because Butler, the district attorney who is prosecuting her, is in the news again today."

In another branch, the AFFECTIVE REASONER agents were shown able to generate, and present, a large number of stories based on a single external plot sequence. The stories varied according to the dispositions and attendant appraisals of the agent-actors presenting the individual stories, so that the themes and characters in each differed significantly. In short, what happened stayed much the same, but how the characters felt about it and why they felt this way varied from story to story, under the control of the program. In a pilot study, it was shown that stories fabricated by the AFFECTIVE REASONER from plot templates and then presented by AFFECTIVE REASONER multimedia agent actors were considered to be highly plausible by subjects (Elliott et al. 1998).

Work is also under way to incorporate the AFFECTIVE REASONER's model of emotion and personality into agent-based tutoring systems with the idea of making the automated tutoring personalities more engaging, more motivating, and more expressive. Preliminary designs have been outlined for integrating a subset of the AFFECTIVE REASONER's models into the STEVE Project at USC-ISI and the COSMO Project at Intellimedia-NCSU (Elliott 1997a; Elliott, Rickel, and Lester 1997) (c.f. the work of Rickel and Lester later).

One of the lines of research in the Affective Reasoning Project that has yet to be developed to any depth but that we feel is promising is the idea of *affective user modeling*. In this approach, the hard problems of general user modeling are left alone, with the focus being placed not on what the user knows but, rather, on how the user feels (Elliott, Rickel, and Lester 1997). Because AFFECTIVE REASONER agents and other emotionally intelligent systems necessarily keep at least implicit internal models of how others see the world—for how else can

one, for example, feel sorry for someone if not by making presumptions about them is being sad?—it is not a big step to then keep a model of a user's presumed emotion state. In the AFFECTIVE REASONER agents, this internal model of others' presumed emotion states is explicit (Elliott and Ortony 1992), and it is only a minor theoretical leap to use this as a basis for input relevant to tutoring, and other, goals. Bolstering this approach is something that we have observed informally in the relationship between users and AFFECTIVE REASONER agents but that is also commonsensical: People are socially motivated to express their emotion states (for example, I am frustrated, I am angry, I admire the way you...) even to a computer agent, as long as the agent has some way to respond appropriately. This research is in contrast to, but can work in concert with, those developing real-time sensing mechanisms for detecting human emotion, such as those in the emotion-sensing subgroup of the Affective Computing Project at MIT (c.f. vismod.www.media.mit.edu/vismod/demos/affect/AC_research/sensing.html).

Just as Fogg and Nass (1997) have shown that users appreciate flattery (at least in the short term) even when they know it for what it is, it might not be so far-fetched to find that users can accept a computer agent that says, for example, "I am just a simple computer program. Still, I consider you to be my friend. I believe that you are unhappy about [fill in the blank]. In my own small way, I am sorry that it happened." For further information on the Affective Reasoning Project, see www.depaul.edu/~elliott.

The Ebb and Flow of Affective Life Juan Velásquez at the MIT Artificial Intelligence Lab has developed a computational model called CATHEXIS for generating emotions in autonomous agents. It makes use of a set of about six emotion families, such as anger and fear, drawing from the work of Ekman, Izard, and others, as well nine *motivational drives*, such as hunger, fatigue, and curiosity, based on control system theory. Each of these has *releasers* that regulate the duration and intensity of states over time. Additionally, the emotion states have elicitors that fall in the categories *neural*, *sensorimotor*, *motivational*, and *cognitive*.

Of particular interest in the CATHEXIS system is that it has a fairly robust model of the ebb and flow of affective states over time. To understand how difficult a problem this is, consider the aforementioned state of fear. Certainly it is true that the more important the goals, or the more likely that they will be blocked, the more intense the fear. Suppose

The building of computer emotion models has long been viewed with a certain amount of skepticism; after all, even deciding what comprises emotion in humans is rife with controversy.

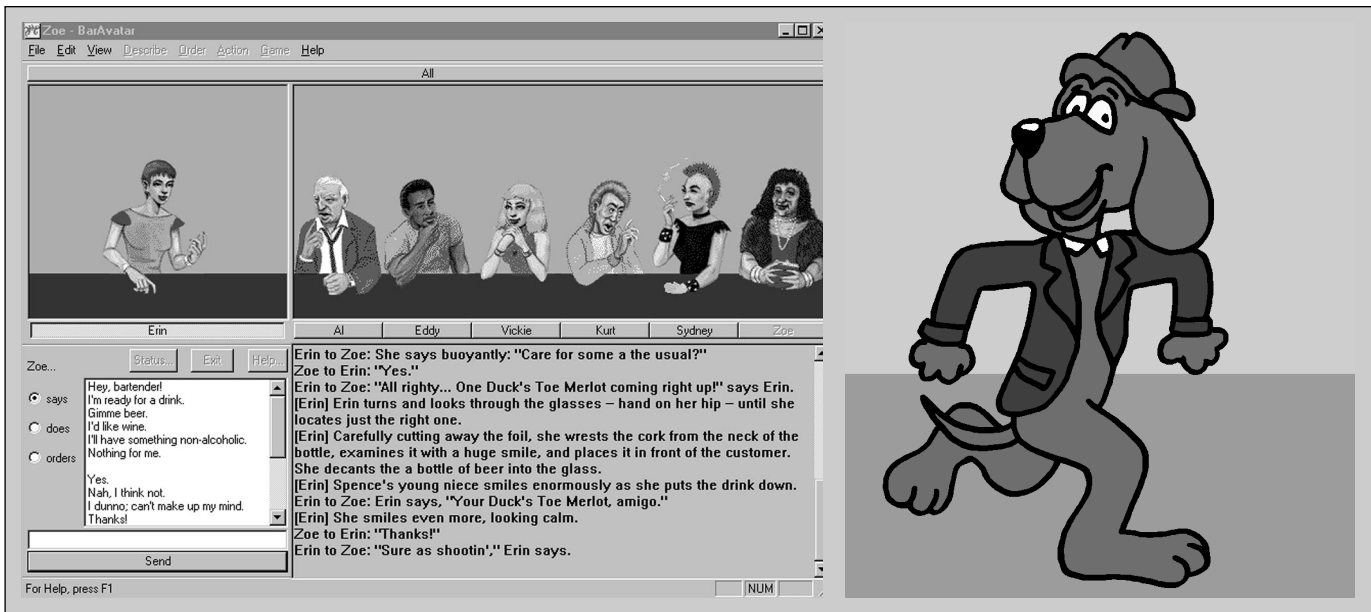


Figure 2. The Agent ERIN at Extempo's Award-Winning SPENCE'S BAR and Max, a Web Tour Guide. (Reproduced with permission.)

that one is fearful about a menacing person believed to be immediately threatening one's social status. Now consider one's dispositional fear of death. In contrast, this fear is likely to be more intense, on the one hand, because of the extreme nature of death but, on the other hand, less intense because it is less immediate. Now we have to ask is this long-term, elemental fear (1) more intense on average (that is, a stronger influence on current behavior); (2) more intense for short, repeating bursts; (3) less intense at each moment (and, of course, what is a moment in an AI system?) but more intense in sum; (4) and so on? Is intensity strictly quantitative, or is it qualitative? (That is, is being mildly pleased closer in nature to being slightly angry or to being intensely rapturous?) The Velásquez model goes farther in addressing these difficult issues than many with its computational mechanisms of elicitors and releasers.

The integrated behavior systems, such as some of those discussed by researchers in Aaron Sloman's (1997) group, operate somewhat autonomously, competing in an inhibition network for chances to manifest themselves. Other behaviors, as well as initial effectors, can regulate a behavior's form and eligibility to fire. In this model, a winner-take-all strategy, which is somewhat controversial in its pure form, is used.

The model is comprehensive and designed for use in both software agents and robots. It has partially been implemented in synthetic agents such as SIMON THE TODDLER and a robot

MUTANT (Fujita and Kageyama 1997). Only time will tell if the sophisticated checks and balances used to create the dynamic processes in the system prove generally applicable in practice, but the group is to be commended on the push toward implementation and testing in such varied environments.

Personalities for Cyberspace As project leader for the Adaptive Intelligent Systems Group at Stanford University's KSL, Barbara Hayes-Roth has long been one of the main proponents of character-rich agents. Hayes-Roth was program chair for the inaugural Autonomous Agents '97 Conference in Marina del Rey, California, and has been tireless in her efforts to legitimize interactive agent research. Most closely related to this article is the Hayes-Roth group's work on the VIRTUAL THEATER Project and a commercial offshoot, Extempo Systems. The latest VIRTUAL THEATER communications can be found at ksl-web.stanford.edu/projects/CAIT/index.html. The latest Extempo developments can be viewed at www.extempo.com/.

One of the research lines for this group is in developing user-guided characters and fully autonomous characters for use in shared virtual environments. These environments, such as KSL's CYBERCAFE and Extempo's award-winning SPENCE'S BAR, allow users to meet and engage in directed social interaction. The user-controlled avatars and autonomous characters build interactive stories together in an online chat room (figure 2).

At KSL, Daniel Rosseau and Hayes-Roth

(1997a) have been designing schemes for a high-level descriptive representation of personality appropriate for such interactive characters. The goal of the research is to allow authors to build distinct and recognizable personality types suitable for use in the interactive virtual environments. The models are based partly on *trait theory* (wherein a multidimensional space of characteristics such as sociability and extraversion is seen as determining how one will act in society) and social learning theories (wherein behavior is partly determined by context characteristics and the individual's past experiences in similar situations).

The personality-profiling model allows for the specification of traits, such as self-confidence, activity, and friendliness, which can be varied along a numeric continuum. These, in part, determine how an agent reacts to situations in the virtual environment. The traits, in turn, can depend on values of agent states, such as happiness-sadness (self-oriented affective states), gratitude-anger (other-oriented affective states, for example, grateful to someone), and liking-hatred (attraction-oriented affective states). Characteristics such as these are used to create the dispositional, and dynamically variable, personalities of agents used in the interactive environment.

Using such controls over agent behavior, one is able to define personalities that reflect the intended high-level characteristics of labeled lay-personality types (c.f., Elliott [1993]). For example, one might create agents with general types of nasty, friendly, shy, lazy, choleric, and selective (friendly with some, nasty with others) (Rousseau and Hayes-Roth 1997b).

Early exploratory studies using this approach have shown promise in that users are able to recognize the intended personality characterizations and that they respond to them in socially coherent ways.

The efforts of the VIRTUAL THEATER Project and Extempo Systems share many common elements with the groups that follow in the next subsection, in that they support the creation of virtual human-computer communities through the use of agents (Rousseau and Hayes-Roth 1997b).

Let Me Virtually Introduce You to a Friend

In the work of Carnegie Mellon University (CMU) Zoesis, Pfmagic, Cyberlife, and Extempo, we find research groups and companies that have built marketable models of virtual friends (and antagonists!). These groups have in common the creation of likable, attractive

agents that interact with both one another and users in creating virtual communities.

Follow the Yellow Brick Road...to Carnegie Mellon The OZ Group at CMU, led by Joseph Bates (one of the original area chairs for Autonomous Agents '97), has made some major contributions to the development of computer agents as characters. Not the least of these has been the long-term efforts of the group to legitimize the very hard, and sometimes very mainstream AI-ish, problems to be solved in creating interesting, engaging, believable, interactive characters. A postulate of the OZ Project's work is that the body of artistic knowledge developed for film animation, acting, and fiction writing not only can be transferred from the original sources into interactive computing but, ultimately, must be. The OZ Group sought input from artists on the projection of personality into characters and suggested early on that progress in these areas would generalize to products for entertainment, commercials, computer games, and corporate training.

Driving this research is the idea that much like musical conception, the artistic spark needed to inspire a truly great character might first be envisioned in an unencumbered, purely artistic world, but it ultimately must be manifested in the real world—here in sequences of computer code and graphic bit mappings. Along the way, great attention to detail is required so that the *there* originally there is not lost in the shuffle of timing, protocols, AI planning, knowledge representation, and the like.

One difficult problem the group has faced in bringing this work into the mainstream of AI is that the empirical testing of progress is difficult. The goals of the project are wedded to an artistic sensibility, albeit in the AI paradigm. Testing whether a character works as a plausible, interesting entity is a far cry from measuring execution time, polygons, or even the quasirigorous educational effectiveness of a program: There is no *big oh* of plausibility.

One lesson that seems to have fallen from the corpus of their research is that creative work on interesting interactive believable characters, just like characters for other mediums, requires that one must each time create, from scratch, a substantial amount of what is intuitively novel and intriguing about a new character. One way to consider this is that although tools and frameworks can be extrapolated, original art, once reused, is no longer original. The group has additionally worked on usable models of emotion (Reilly 1996); interactive fiction (Peter Weyhrauch [Wey-




Figure 3. *The Playful Woggles.* (Reproduced with permission.)

hrauch and Bates 1995)]; natural language for believable agents (Loyall and Bates 1995); and, most recently, what we might call *action discourse* (Sengers 1998), wherein actions based on disparate motives are constrained to be presented in realistic sequences (and, hence, related to the work of Lester et al. and others on diectic believability [see *Agents That Educate in an Entertaining Way*]). The tools generated by this group have been used by a number of

others (for example, Hayes-Roth et al. in the animated puppets work in the VIRTUAL THEATER Project at KSL) and Wurst and McCartney at the University of Connecticut).

Like many agent researchers, they are probably best known by the look of some of their characters, that is, as the “Woggles guys and gals from CMU.” The woggles might be described as cute, squishy, minimalist, highly expressive, playful, social balls of graphic fur. In their classic implementation, they are, above all, playful (figure 3).

Most recently, the group has spawned a commercial enterprise, Zoesis, which finally settled in Boston, Massachusetts, this past year. Necessarily, with the migration to building a commercial product, the group now has to be more protective of its current ideas. Zoesis still maintains strong ties to CMU, however, and its members remain active at conferences. See 128.2.242.152/afs/cs.cmu.edu/project/oz/web/ for further information.

The Mad Scientists at Work At Autonomous Agents '97, Stephen Grand, at the time not well known to others at the conference, gave a presentation on a system, CREATURES, that was in production and continued development at CyberLife Technology Ltd. (Grand, Cliff, and Malhotra 1997). Grand showed pictures of the relatively detailed and dense 2-1/2-dimensional (2-1/2 D) (with flat but layered graphics) CREATURES world, filled with cute imaginary animals, elevators that rise and fall, toys, food, and so on, that he presented as objects with scripts that describe how they interact with other objects (figure 4). The talk continued with details of how the biped creatures respond to positive and negative reinforcement; mature over time as they grow; and, most importantly, are influenced by their (symbolically represented and effected) genetic makeup. By the time Grand began discussions of the genetic rules, user control of genetics, adaptation between generations, the creatures' ability to learn a simple verb-object language, and a surprisingly robust (for a game!) neural net “brain” model of a quasibiological self-modifying response to stimulus, the audience reaction was probably one of, “We have another mad scientist here, albeit one who certainly knows all the right buzzwords. Too bad such systems are so deadly dull when one looks at the code and sees them run.”

To the delight of many (all?), Grand's presentation ultimately did convince—and decidedly so. Creatures turned out to be one of those rare instances where a diverse set of content theories, after being filtered into similar levels of granularity, were synthesized into a

real-time, temporally robust system that works. Grand and his colleagues mix knowledge of genetics, biochemistry, object-oriented programming, engineering, artificial life, artistic craft, and a great sense of both humor and fun into an inspired piece of small-scale software that not only holds up as entertainment but is also fascinating to ponder from an intellectual point of view. Further papers and general information on this work can be found at www.cyberlife.co.uk/.

PETZ—Agentz for the Massez With virtual *PETZ* (both cats and dogs) from PF Magic, users can invite autonomous agents to play on their PC desktop. A user might interact with her *PETZ* by petting them, introducing them to one another at different ages, feeding (or not feeding) them, and so forth. Unlike *CREATURES*, users do not have direct access to any sort of slider-bar tweaking of internals. In this way, a user's *PETZ* are truly autonomous. The *PETZ* agents use real-time animation and the layering display of multiple simultaneous behaviors to create the illusion of lifelike continuous motion (figure 5).

PETZ graphic behaviors mimic those simple behaviors we expect of real pets: *DOGZ*, for example, wag their tails; have perky ears and expressive eyes; have noses that follow "scents"; have spots in which they like to be petted; and have tongues that can be used to express being tired, excited, thirsty, and hungry.

PETZ have persistent personalities, but these can change over time if, for example, a pet is not fed appropriately (pets come with virtual food and virtual toys). Additionally, users can modify the behavior of their *PETZ* by squirting them with a virtual spray bottle.

The newest version of this product allows *PETZ* to interact with one another without direct input from the user. Some relationships are biased (for example, *DOGZ* and *CATZ* will tend not to get along well, older *Petz* nurture younger *PETZ*), but these can indirectly be modified by the user over time. These autonomous relations reflect a general trend in the computer game industry wherein *non-player characters* (NPCs) such as those in Origin's *ULTIMA ONLINE* (www.owo.com) make increased use of intelligent decision making, state maintenance, and autonomous action.¹

From an agent-sociological perspective, it will likely be interesting to study the nature of the long-term relationships that develop between users and their *PETZ* and the virtual communities that have sprung up on the web for discussing and sharing information on *PETZ*.




Figure 4. *Creatures—Sophisticated Imaginary Animals from CyberLife Technology Ltd. (Reproduced with permission.)*

Greasing the Wheels: Tempering Foundational Theories of Believable Interaction

Autonomous agents are necessarily complex, and different labs focus on different components of these sophisticated devices. In some cases, simple believability is the primary goal. In other cases, there is an effort to distill foundational principles that are both faithful to the




Figure 5. Pfmagic's PETZ Interacting with One Another Autonomously. (Reproduced with permission.)

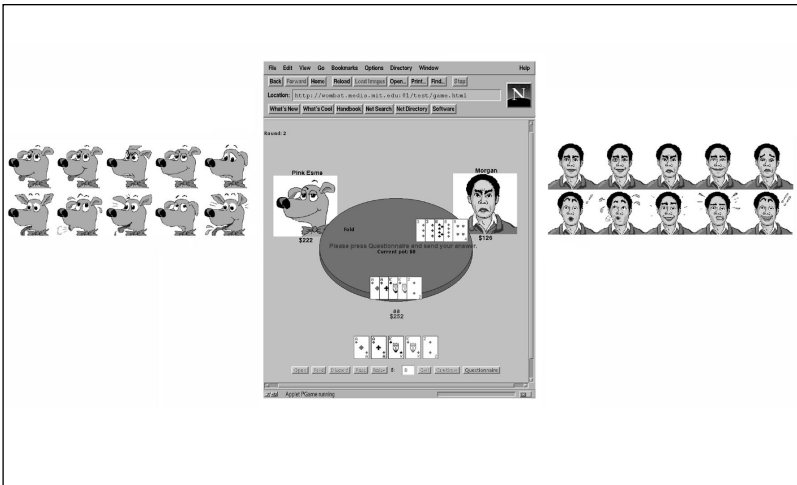


Figure 6. Some of Tomoko Koda's E-POKER Agents at the Massachusetts Institute of Technology. (Reproduced with permission.)

real-world examples from which we derive our models and also broadly applicable to large classes of agents. Labs in this latter class would include, for example, those associated with Aaron Sloman (University of Birmingham), Norman Badler (University of Pennsylvania), and Justine Cassel (MIT). Here, we present a few novel approaches to systems that have this more foundational approach.

You Know You're Alive When You Have...Synthetic Emotions A number of ongoing projects are relevant to lifelike agents at MIT, many under the influence of Pattie Maes. In one such project, Tomoko Koda's E-

POKER, users were studied in the way they interacted with computer poker players who displayed emotions using a variety of characters, each of which had a number of different expressions. As part of her master's work, "Agents with Faces: A Study on the Effects of Personification of Software Agents," Koda built agents capable of playing poker and designed several user studies. Koda's emotion model was an augmented subset of that originally specified by Ortony, Clore, and Collins (1988). Among her findings were that personified interfaces helped users become engaged in tasks, faces were seen differently in isolation than they were when experienced as part of a task, and perceived intelligence of a face has more to do with the underlying competence of an agent than its graphic presentation (figure 6) (tomoko.www.media.mit.edu/people/tomoko/).

In the Media Lab's ALIVE (ARTIFICIAL LIFE INTERACTIVE VIDEO ENVIRONMENT) (Blumberg 1995; Maes et al. 1994) work (lcs.www.media.mit.edu/projects/alive/), users are able to interact with virtual creatures without being constrained by the usual trappings of virtual reality systems. Rather, as long as users stay within a large rectangle of floor space, their image appears in the virtual world on a large-screen television image using the wonders of video tracking. Autonomous animated characters, such as Bruce Blumberg's well-known SILAS T. DOG, cohabit the space with the user image, and users are able to interact with them (Blumberg 1996; Blumberg, Todd, and Maes 1996) (figure 7). Rather than rely on author input to tweak the behavior of the agents, these rather phenomenal creatures operate truly autonomously, using a set of motivations and goals




Figure 7. Bruce Blumberg's *SILAS T. DOG* from the Massachusetts Institute of Technology. (Reproduced with permission.)

[Patrick Doyle] ... is addressing the problem of generalizing the relationship between an agent and the virtual environment in which the agent exists. The central thesis of this work is that knowledge for agenthood, wherein the agent has its own personality and its own history of interaction with users is internal to the agent, but that domain knowledge is internal to the environment.

based originally on a number of ideas from animal ethology. SILAS, for example, renders scenes from its viewpoint as it navigates through the virtual space, then uses the processed image as input that affects its behavior. At conference demonstrations, the creatures would exhibit engaging and humorous high-level behavior as users waved their arms, embraced them, and walked around the space.

Most recently, in leading the Synthesizing Emotions Group at the MIT Media Lab, Blumberg is pursuing research that focuses not only on the ability of machines to reason about which emotions are appropriate in a given situation but also on the building of machines capable of having emotions, whereby the emotion processes form an integral part of the decision-making processes (bruce.www.media.mit.edu/people/bruce/).

Merlin Uploads Virtual Worlds Patrick Doyle's work at Stanford applies independently developed techniques, both old and new, in a novel application. Doyle's system is a work in progress, but one that, if successful, will benefit a community likely to expand significantly in the near future. Doyle is attempting to build what we might call an infrastructure for intelligent virtual environments (Doyle and Hayes-Roth 1998). That he is working in multiuser dungeons (MUDs) (virtual interactive playgrounds in which users can participate [c.f., Curtis (1992)]) is less important than that he is addressing the problem of generalizing the relationship between an agent and the virtual environment in which the agent exists. The central thesis of this work is that knowledge for *agenthood*, wherein the agent has its own personality and, more importantly, its own history of interaction with users (for example, as a mentor-guide for children), is internal to the agent, but that domain knowledge (for example, how to play chess in one "room" of an MUD and how to buy bread in another) is internal to the environment. Environment knowledge is annotated so that agents, through an established protocol based minimally on metalevel understanding of a domain, can supplement their own beliefs and actions.

To old hands at AI and human-computer interaction (HCI), Doyle's efforts might sound like a rehash of various ideas that have come before, and to an extent, this is true. The work traces its roots back to J. J. Gibson's (1977) affordances, Don Norman's (1990) knowledge in the world, and even Roger Schank's scripts (Schank and Abelson 1977). What makes the work appealing as a modern research paradigm is that it integrates the previous AI and HCI

work with that of creating believable agents, giving it what might be a highly appropriate, and workable, application.

Virtual worlds of various types, with believable agents of all sorts populating them, are here to stay. Any leverage to be gained in making agent intelligence portable from one virtual world to another is well worth seeking, and this is a serious, novel attempt at such an effort.

Computers as Social Actors At Stanford's Social Responses to Communication Technologies Research Group (www-leland.stanford.edu/group/commdept/), Nass, Reeves, et al., formalized the intuition that people apply social rules to many aspects of HCI (for example, Nass, Moon, and Fogg [1995] and Nass, Steuer, and Tauber [1994]). The often-cited studies of this group have been used to counter arguments that the attempt to build social intelligence into computer programs is frivolous. This work supports a counterargument that can roughly be stated as "designing software as a social interface is not something we can avoid because it happens whether we plan for it or not; we have no choice in doing it but only in doing it right."

The studies have illustrated that even when computers were not given explicitly anthropomorphic interfaces, users tended to see them in this light anyway and showed preferences relative to artificial personalities. Although these studies were not intended to suggest that the computer programs used were autonomous agents as such, they did serve to illustrate that the association of a persona with certain types of program was relatively easy to establish and, in some cases perhaps, hard to avoid. In addition to the Stanford group, Youngme Moon continues similar work in MIT's relatively new Social Intelligence Research Group.

Agents That Educate in an Entertaining Way

Outside computer game applications, the use of anthropomorphic interfaces probably has the longest history in tutoring systems. Here we look at five current, intriguing applications from large labs that are devoting significant resources to this research.

HERMAN and His High-Flying Big Brother James Lester and other members of the IntelliMedia Initiative at the Multimedia Laboratory at NCSU have two agents that are of interest to this discussion. The first is HERMAN THE BUG, an impish buglike creature that teaches children about biology, and the second is COSMO, a frenetic adviser on internet protocols




Figure 8. HERMAN THE BUG and COSMO Form the Multimedia Laboratory of North Carolina State University. (Reproduced with permission.)

(figure 8). Both projects share tremendous attention to detail in graphics, gesture logic, sequencing, and theoretical underpinnings. HERMAN THE BUG has undergone rigorous testing, and COSMO is headed for the same. The character of the two agents is different, with HERMAN a likable clown and COSMO a hip and hyped internet cosmonaut.

The thrust of these efforts is rooted in both AI and education. The work seeks to clarify the issues raised when animated interface agents use gesture, speech, and guided locomotion to communicate important information (for example, to focus a student's attention on an object in the world) but must also maintain believability in the characters' presence and actions. These two goals are not necessarily linked. Lester gives the examples, on the one hand, of a humorous, lifelike, joke-cracking character that ultimately impedes problem solving through its distracting presence and, on the other, of a dull assistant that always operates appropriately but fails to engage the student. When communications from an agent must be coordinated to be both engaging and purposeful, issues in timing and the multilayering of actions arise. The Lester group is pursuing diectic believability, wherein the behavior planner for animated agents allows them to move through space and refer to (sometimes dynamic) objects in its dialog with the user in a way that is both natural and unambiguous.

An important aspect of the Lester group's research is that it seeks to carefully document and test its findings about precisely what kinds of contribution lifelike agent capabilities make to learning. That the systems have the look and feel of computer games is a distant second

to the core academic principles from which derive even the smallest gestures and action sequences in the systems. Intellimedia uses rigorous empirical testing to assess the contributions of agents to problem solving, higher-order learning, and the effective impact on students (for example, see Lester, Converse, et al. [1997a]). Several extant studies discuss what might best be referred to as existence proofs that the intended influences on education are manifested by systems using these agents (Lester, Callaway, et al. 1997; Lester et al. 1997).

A number of other labs are also working on issues related to the diectic believability work at NCSU and, more generally, on developing components of semantically realistic agent interfaces. Among them are USC-ISI (see later discussion), the University of Pennsylvania Center for Human Modeling and Simulation (Badler, Phillips, and Webber 1993) (www.cis.upenn.edu/~hms/badler/pap/pap.html),² New York University's Media Research Lab (Perlin and Goldberg 1996) (www.mrl.nyu.edu/improv/), and the MIT Media Lab's Gesture and Narrative Language Group (Vilhjalmsson and Cassell 1998; Thorisson 1997; Cassell et al. 1994).

Don't Blow Up That H-Pack! Lewis Johnson, in his capacity as conference chair, was a major force behind the success of the Autonomous Agents '97 Conference hosted by USC-ISI (www.isi.edu/isd/Agents97/info.html). Johnson's Educational Technology Group at ISI has a pair of agents under development that we discuss here. The newest addition, ADELE (AGENT FOR DISTANCE LEARNING ENVIRONMENTS) is a 2D-3D pedagogical agent for presenting web-based course materials (Johnson and Shaw



Figure 9. The Agents from the Educational Technology Group at the University of Southern California Information Sciences Institute.

Top: ADELE. Bottom: STEVE. (Reproduced with permission.)

1997) (www.isi.edu/isd/ADE/ade.html) (figure 9a). A copy of ADELE runs locally, monitors student actions, and reports end-of-session information to a central server. The project is tackling both case-based courses and the more difficult problem-based learning environments wherein teachers act as guides for student discovery of relevant information. Authoring tools are also being developed for ADELE's troika of instructional narrative, problem-solving simulation, and reference materials. Students are exposed to the instructional narrative, take part in discovery through interactive problem solving, and can refer to the online instructional materials. Additionally, online discussion with other students and instructors can take place. ADELE plays the part of the tutor-guide during problem solving and in the after-exercise summary. An interesting component of ADELE is that its thrust is on *client-side* intelligence, where the user and agent operate as a somewhat autonomous social unit.

A longer-running project, embodied as the pedagogical agent STEVE (SOAR TRAINING EXPERT

FOR VIRTUAL ENVIRONMENTS) is coordinated under Johnson at ISI by Jeff Rickel. STEVE, embodied in various hand-head incarnations (Rickel and Johnson 1997) lives in an immersive virtual 3D environment and helps students learn physical procedural tasks such as operating or repairing complex equipment (for example, a naval H-pack compressor). As STEVE observes the dynamic state of the world and a student's interactions with it, it can choose to intervene by giving demonstrations, coaching the student, and manipulating objects in the world ad hoc. STEVE "talks" with students using text-to-speech software and uses gestures to make indications in the virtual world. STEVE helps students in various ways; it can demonstrate tasks, answer questions about the rationale behind task steps; and monitor students while they practice tasks, providing help when requested (figure 9b).

Additionally, there are plans to extend STEVE to fill in the role of missing team members during instruction in team training. This type of interactive simulation-based training is not possible without an embodied tutor that takes part in exercises within the virtual environment.

The system has three parts: (1) STEVE, the *pedagogical agent*; (2) the *virtual reality software* that handles the interface between students and the virtual world, updating a head-mounted display and detecting user interactions with the objects in the virtual world; and (3) a symbolic *world simulator* that maintains the state of the virtual landscape. STEVE gets messages from the virtual reality software about user actions as well as from the simulator about, for example, the resulting new state of the world.

The two projects have roots in, and are responsible to, both the AI community and the education and training communities. One theme they share is in looking at how an embodied agent sharing a virtual 2D or 3D space permits the use of gesture and other nonverbal actions to convey critical tutoring information (for example, by giving virtual demonstrations or otherwise visually manipulating objects in the virtual world).

Some of the hard problems this group faces have to do with maintaining diegetic believability in gesture and reference (see the Lester work discussed earlier), building robust personalities (see the AFFECTIVE REASONER discussion earlier [Elliott, Rickel, and Lester 1997]), and developing authoring tools for use with the graphic and speech subsystems. An introduction to the work can be found at www.isi.edu/isd/etg.html. For insight into the related work of Billingham and Savage (1996)

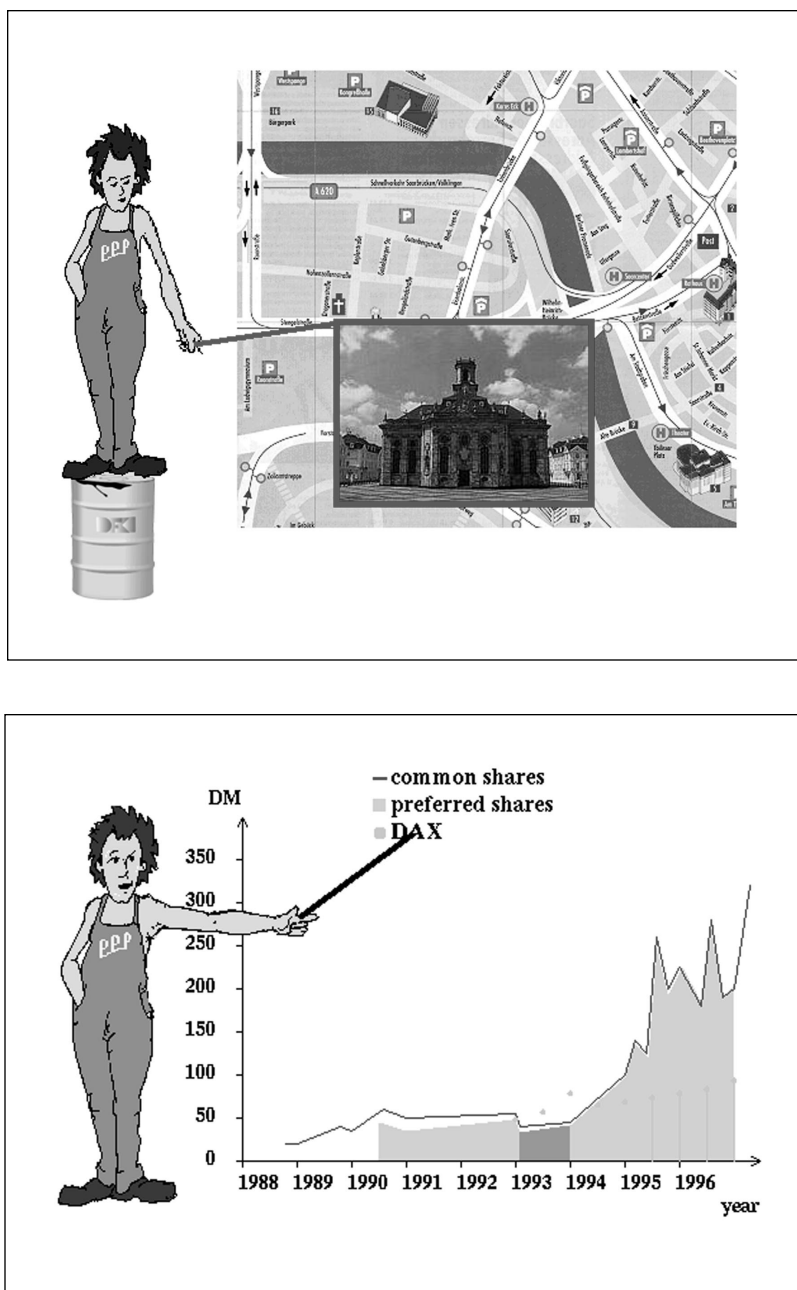


Figure 10. PPP System of DFKI. (Reproduced with permission.)

at the University of Washington's Human Interface Technology Lab, who are also working on complex human interface environments, go to www.hitl.washington.edu/information/index.html.

Just Ask the Playful Pointing Person for Directions

The work of Elisabeth Andre,³ Thomas Rist, Jochen Muller, and others at DFKI is manifested in one incarnation of the PPP system (although one author always thinks of it as the playful pointing person) as a likable, overall-clad, virtual persona that helpfully points out aspects of commercial presentations and computer-based instruction (Andre, Rist, and Muller 1998; Andre and Rist 1997). It has a wide variety of expressions and gestures that it uses to guide the user through tasks. The agent is embedded in web pages and carries out presentations such as *idle-time actions*, such as breathing and tapping a foot; *reactive behaviors*, such as keeping current when an object that the agent is pointing to is moved by the user; and *navigation acts*, such as walking to another screen position (figure 10).

The interface model has a carefully crafted appeal, designed to be used in extant products, such as those for illustrating product design or showing real estate on the web. At the same time, the theoretical underpinnings of the system look carefully at the specifics of user focus and attention and are manifested in a robust hierarchical planning model, after earlier work by the same researchers, and temporal reasoning models (after Allen [1983]). Furthermore, unlike for some commercial production systems, there is a scholarly record of both the agent development and the results of research studies.

Presentations come as sequences of *presentation acts* that are developed in concert with a human author. Although the completed "performances" are generally script based, the user can interrupt the scripted sequence of events, causing the agent to initiate replanning of the presentation under the new constraints. This replanning occurs under program control, wherein new script material is imported into the current environment so that the web context remains constant (that is, the user is not transported to another linked portion of the web). Careful coding for efficiency allows PPP to run acceptably, and background resources are devoted to planning and recompiling on the fly.

DFKI is actively working on authoring applications to be used with PPP. Additionally, although it uses a relatively simple hierarchical rhetorical structure, through its on-the-fly compiling, context-sensitive reactivity, and

innovative use of planning and temporal constraint technology, it supports an impressive set of interactions with the user. Like the Cyberlife and Extempo work noted previously, the PPP Project is a commendable instance of commercial development that maintains a strong academic record. An introduction to the demonstrations can be seen at www.dfki.uni-sb.de/jmueller/ppp/persona.

And Behind Door Number Three...

Only time will tell whether the current interest in software systems that embody lifelike qualities will continue apace. There are many signs that interest will continue to increase and that system intelligence manifested through synthetic agent emissaries will become, if not commonplace, at least highly accepted in the mainstream. The combination of a new breed of web-software consumers likely to be interested in nontraditional applications and interfaces, the increase in local computing power allowing for pragmatic approaches to speech recognition and sophisticated real-time expressive graphics, and the emergence of "killer apps" for many AI-related technologies suggests that this area will grow significantly. After all, our own world is infinitely richer and more engaging than that of computers. In the end, it might well make sense to have programming interfaces join our own complex social fabric now that we have the skills to inspire them with the illusion of life.

Notes

1. Additionally, it seems that multiplayer online games, where NPCs thrive, are a growth industry. For example, a recent check of the internet PC game charts showed five interactive action games in the Top 40 (www.worldcharts.com).
2. This group also maintains an extensive set of pointers to about 70 other virtual humans around the world (www.cis.upenn.edu/hms/badler/vhlist.html).
3. Andre recently chaired the Animated Interface Agents Workshop at the Fifteenth International Joint Conference on Artificial Intelligence.

References

- Allen, J. F. 1983. Maintaining Knowledge about Temporal Intervals. *Communications of the ACM* 11(26): 832-843.
- Andre, E., and Rist, T. 1997. WEBPERSONA: A Lifelike Presentation Agent for the World Wide Web. Paper presented at the Fifteenth International Joint Conference on Artificial Intelligence Workshop on Animated Interface Agents: Making Them Intelligent, 23-29 August, Nagoya, Japan.
- Andre, E.; Rist, T.; and Muller, J. 1998. Integrating

- Reactive and Scripted Behaviors in a Lifelike Presentation Agent. In *Proceedings of the Second International Conference on Autonomous Agents*. New York: Association of Computing Machinery. Forthcoming.
- Badler, N. I.; Phillips, C. B.; and Webber, B. L. 1993. *Simulating Humans: Computer Graphics, Animation, and Control*. New York: Oxford University Press.
- Billinghurst, M., and Savage, J. 1996. Adding Intelligence to the Interface. In *Proceedings of the IEEE 1996 Virtual Reality Annual International Symposium*, 168–176. Washington, D.C.: Institute of Electrical and Electronics Engineers.
- Blumberg, B. M. 1996. *Old Tricks, New Dogs: Ethology and Interactive Creatures*. Ph.D. dissertation, Media Lab, Massachusetts Institute of Technology.
- Blumberg, B., and Galyean, T. A. 1995. Multilevel Direction of Autonomous Creatures for Real-Time Virtual Environments Computer Graphics. In *SIGGRAPH '95 Proceedings*, 295–304. New York: Association of Computing Machinery.
- Blumberg, B. M.; Todd, P. M.; and Maes, P. 1996. No Bad Dogs: Ethological Lessons for Learning in Hamsterdam. In *From Animals to Animats: Proceedings of the Fourth International Conference on the Simulation of Adaptive Behavior*. Cambridge, Mass: MIT Press.
- Botelho, L. M. 1997. *Building Intelligent Agents: An Approach Based on the Cognitive Models of Decision Making*. Ph.D. dissertation, Department of Information Sciences and Technologies, Instituto Superior de Ciências do Trabalho e da Empresa.
- Cassell, J.; Pelachaud, C.; Badler, N.; Steedman, M.; Achorn, B.; Becket, T.; Douville, B.; Prevost, S.; and Stone, M. 1994. Animated Conversation: Rule-Based Generation of Facial Expression, Gesture, and Spoken Intonation for Multiple Conversational Agents. In *Proceedings of SIGGRAPH '94*, 413–420. New York: Association of Computing Machinery.
- Curtis, P. 1992. Mudding: Social Phenomena in Text-Based Virtual Realities. Paper presented at the 1992 Conference on Directions and Implications of Advanced Computing, 2–3 May, Berkeley, California.
- Doyle, P., and Hayes-Roth, B. 1998. Agents in Annotated Worlds. In *Proceedings of the Second International Conference on Autonomous Agents*. New York: Association of Computing Machinery. Forthcoming.
- Elliott, C. 1997a. AFFECTIVE REASONER Personality Models for Automated Tutoring Systems. Paper presented at the Workshop on Pedagogical Agents, 19 August, Kobe, Japan.
- Elliott, C. 1997b. I Picked Up Catapia and Other Stories: A Multimodal Approach to Expressivity for “Emotionally Intelligent” Agents. In *Proceedings of the First International Conference on Autonomous Agents*, 451–457. New York: Association of Computing Machinery.
- Elliott, C. 1993. Using the AFFECTIVE REASONER to Support Social Simulations. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 194–200. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Elliott, C. 1992. The AFFECTIVE REASONER: A Process Model of Emotions in a Multiagent System. Ph.D. dissertation, Technical report, 32, The Institute for the Learning Sciences, Northwestern University.
- Elliott, C., and Ortony, A. 1992. Point of View: Reasoning about the Concerns of Others. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 809–814. Hillsdale, N.J.: Lawrence Erlbaum.
- Elliott, C.; Rickel, J.; and Lester, J. 1997. Integrating Affective Computing into Animated Tutoring Agents. Paper presented at the IJCAI-97 Workshop, Animated Interface Agents: Making Them Intelligent, 24–29 August, Yokohama, Japan.
- Elliott, C.; Brzezinski, J.; Sheth, S.; and Salvatoriello, R. 1998. Story-Morphing in the Affective Reasoning Paradigm: Generating Stories Automatically for Use with “Emotionally Intelligent” Multimedia Agents. In *Proceedings of the Second International Conference on Autonomous Agents*. New York: Association of Computing Machinery. Forthcoming.
- Fogg, B. J., and Nass, C. 1997. Silicon Sycophants: The Effects of Computers That Flatter. *International Journal of Human-Computer Studies* 46(5): 551–561.
- Foner, L. 1997. Entertaining Agents: A Sociological Case Study. In *Proceedings of the First International Conference on Autonomous Agents*, 122–129. New York: Association of Computing Machinery.
- Fujita, M., and Kageyama, K. 1997. An Open Architecture for Robot Entertainment. In *Proceedings of the First International Conference on Autonomous Agents*, 435–442. New York: Association of Computing Machinery.
- Gibson, J. J. 1977. The Theory of Affordances. In *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*, eds. R. Shaw and J. Bransford. Hillsdale, N.J.: Lawrence Erlbaum.
- Grand, S.; Cliff, D.; and Malhotra, A. 1997. CREATURES: Artificial Life Autonomous Software Agents for Home Entertainment. In *Proceedings of the First International Conference on Autonomous Agents*, 22–29. New York: Association of Computing Machinery.
- Johnson, W. L., and Shaw, E. 1997. Using Agents to Overcome Deficiencies in Web-Based Courseware. Paper presented at the Workshop on Pedagogical Agents, 19 August, Kobe, Japan.
- Koda, T. 1997. Agents with Faces: A Study on the Effects of Personification of Software Agents. Master's thesis, Media Lab, Massachusetts Institute of Technology.
- Lester, J.; Callaway, C.; Stone, B.; and Towns, S. 1997. Mixed-Initiative Problem Solving with Animated Pedagogical Agents. Paper presented at the Workshop on Pedagogical Agents, 19 August, Kobe, Japan.
- Lester, J.; Converse, S.; Stone, B.; Kahler, S.; and Barlow, T. 1997a. Animated Pedagogical Agents and Problem-Solving Effectiveness: A Large-Scale Empirical Evaluation. In *Proceedings of the Eighth World Conference on Artificial Intelligence in Education*, 23–30. Amsterdam, The Netherlands: IOS Press.
- Lester, J.; Converse, S.; Kahler, S.; Barlow, T.; Stone, B.; and Bhogal, R. 1997b. The Persona Effect: Affec-

tive Impact of Animated Pedagogical Agents. Paper presented at CHI '97, 22–27 March, Atlanta, Georgia.

Loyall, A. B., and Bates, J. 1995. Personality-Rich Believable Agents That Use Language, Technical report, CMU-CS-95-139, Computer Science Department, Carnegie Mellon University.

Maes, P.; Darrel, T.; Blumberg, B.; and Pentland, S. 1994. ALIVE: Artificial Life Interactive Video Environment. In Proceedings of the Twelfth National Conference on Artificial Intelligence, 1506. Menlo Park, Calif.: American Association for Artificial Intelligence.

Miksch, S.; Cheng, K.; and Hayes-Roth, B. 1997. An Intelligent Assistant for Patient Health Care. In Proceedings of the First International Conference on Autonomous Agents, 458–465. New York: Association of Computing Machinery.

Nass, C.; Moon, Y.; and Fogg, B. J. 1995. Can Computer Personalities Be Human Personalities? *International Journal of Human-Computer Studies* 43:223–239.

Nass, C.; Steuer, J.; and Tauber, E. 1994. Computers Are Social Actors. Paper presented at the CHI Conference, 24–28 April, Boston, Massachusetts.

Norman, D. 1990. *The Design of Everyday Things*. New York: Doubleday.

Ortony, A.; Clore, G. L.; and Collins, A. 1988. *The Cognitive Structure of Emotions*. New York: Cambridge University Press.

Perlin, K., and Goldberg, A. 1996. A System for Scripting Interactive Actors in Virtual Worlds. In SIGGRAPH '96 Proceedings, 205–216. New York: Association of Computing Machinery.

Reilly, W. S. 1996. Believable Social and Emotional Agents. Ph.D. dissertation, Computer Science Department, Carnegie Mellon University.

Rickel, J., and Johnson, L. 1997. Integrating Pedagogical Capabilities in a Virtual Environment Agent. In Proceedings of the First International Conference on Autonomous Agents, 30–88. New York: Association of Computing Machinery.

Rousseau, D., and Hayes-Roth, B. 1997a. A Social-Psychological Model for Synthetic Actors, Technical report, KSL-97-06, Knowledge Systems Laboratory, Stanford University.

Rousseau, D., and Hayes-Roth, B. 1997b. Interacting with Personality-Rich Characters, Technical report, KSL-97-06, Knowledge Systems Laboratory, Stanford University.

Schank, R. C., and Abelson, R. 1977. *Scripts, Plans, Goals, and Understanding*. Hillsdale, N.J.: Lawrence Erlbaum.

Sengers, P. 1998. Do the Thing Right: An Architecture for Action Expression. In Proceedings of the Second International Conference on Autonomous Agents. New York: Association of Computing Machinery.

Slovan, A. 1997. What Sort of Architecture Is Required for a Humanlike Agent? In *Foundations of Rational Agency*, eds. M. Wooldridge and A. Rao. New York: Kluwer Academic.

Thorisson, K. 1997. GANDALF: An Embodied Humanoid Capable of Real-Time Multimodal Dia-

logue with People. In Proceedings of the First International Conference on Autonomous Agents, 536–537. New York: Association of Computing Machinery.

Velásquez, J. D. 1997. Modeling Emotions and Other Motivations in Synthetic Agents. In Proceedings of the Fourteenth National Conference on Artificial Intelligence, 10–16. Menlo Park, Calif.: American Association for Artificial Intelligence.

Vilhjalmsson, H. H., and Cassell, J. 1998. Understanding Jokes: A Neural Approach to Content-Based Information Retrieval. In Proceedings of the Second International Conference on Autonomous Agents. New York: Association of Computing Machinery. Forthcoming.

Weyrauch, P., and Bates, J. 1995. Guiding Interactive Drama. Paper presented at the AAAI Spring Symposium on Interactive Story Systems, 27–29 March, Stanford, California.

Zrehen, S., and Arbib, M. 1998. Understanding Jokes: A Neural Approach to Content-Based Information Retrieval. In Proceedings of the Second International Conference on Autonomous Agents. New York: Association of Computing Machinery. Forthcoming.



Clark Elliott is an associate professor at DePaul University's School of Computer Science, where he heads the Affective Reasoning Project and is associate director of the Distributed Systems Division. He has served as an area chair for Autonomous Agents '98 and on the program committees for the American Association for Artificial Intelligence and the International Joint Conferences on Artificial Intelligence. He received his Ph.D. in computer science in 1992 from Northwestern University. His research emphasis is AI representations of emotion and personality.



Jacek Brzezinski is a Ph.D. student in the School of Computer Science, Telecommunication, and Information Systems. His research interests include AI, particularly machine learning.