

Putting Intelligent Characters to Work

Barbara Hayes-Roth

■ *Extempo Systems, Inc., was founded in 1995 to commercialize intelligent characters. Our team built innovative software and novel applications for several markets. We had some early-adopting customers during the Internet boom, but the company could not survive the significant downturn in corporate IT spending when the bubble burst. In 2004, Extempo ceased operations and was formally liquidated. Although our commercial venture failed, we advanced the technology for intelligent characters and learned a lot about how (not) to take them to market. Now we are trying again with a new company. To others who aspire to commercialize AI technology, I say: Take a chance!*

Extempo was founded to commercialize intelligent characters. Our technology grew out of research on intelligent agents, begun at the Rand Corporation and pursued over a decade at Stanford University. At its core, the “dynamic control architecture” models the orderly yet opportunistic nature of human cognition (Hayes-Roth, 1985). It provides a uniform mechanism by which an agent triggers and chooses among context-appropriate actions for both task-level and meta-level actions. A series of challenging applications were built to evaluate the architecture and inspire extensions for new functionality.

In early applications, cognitive agents performed heuristic design tasks, such as protein structure modeling and site layout, where dynamic control produced efficiency in complex reasoning computations. The architecture was extended with reusable design ontologies, languages, and strategies, which also supported semantic explanation of the design process and explicit reasoning traces for learning. In later applications, situated agents performed monitoring and control tasks, for example office robots, semiconductor manufacturing, and patient monitoring in the intensive care unit. Here, the architecture was extended to support continuous operation in complex real-time environments; coordination of concurrent perception, cognition, and action; asynchronous performance of loosely coupled tasks, such as situation monitoring, fault detection, diagnosis, and planning; con-

text-sensitive selection among alternative methods for particular instances of each task; and closed-loop control (Hayes-Roth 1995). Addressing practical matters of software development, our next-generation architecture also provided a component-based plug-and-play approach to supporting these capabilities for a broad class of adaptive intelligent agents (Hayes-Roth and Larsson 1997).

In the early 1990s, we were ready for new challenges. At the 1994 AAI Spring Symposium on believable agents (Stanford, California), Joe Bates demonstrated his delightful Woggles—spherical creatures that evinced social-emotional dynamics while playing together in a colorful virtual world (Bates 1994). I was hooked! Bates’s evangelizing of believable agents reignited my original professional interest in psychology (on hiatus since arriving at Stanford) and the potential for embodied agents that manifest personality, emotion, and social dynamics. It also led to the performing arts literature, notably Keith Johnstone’s (1987) brilliant, deeply psychological theory of theatrical improvisation. It was intriguing to realize that our agent architecture offered a felicitous foundation for a new class of “improvisational characters” and new classes of application genres related to learning, play, and the performing arts.

Our first new application was a toy that turned out to be a learning toy. We created two “Improv Puppets” by using our architecture to build improvisational character “minds” and integrating them with

two of Bates's Woggle "bodies" (Hayes-Roth and van Gent 1997). My own children helped to create the puppets by generating and recording hundreds of lines of dialogue for specified situations (for example, introduce yourself, invite the other puppet to play, accept or decline an invitation, sing a song), with variations along three mood dimensions (happy-sad, peppy-tired, friendly-shy). Although the puppets had no mouths, speaking in real children's words and voices imbued them with adorable personalities and emotions. They also had actions and behavior patterns, with variations for different moods and situations (for example, hop on pedestals, jump in the chute, hide, play follow the leader). The puppets had internal logic that modulated their moods and inclinations to perform different behaviors or speak different dialogues, based on the immediate situation and its changes over time. Sitting at networked computers, two children could play with (or "be") the puppets in a shared virtual world, much as they would with traditional puppets in a shared physical world—but our puppets were smart and semiautonomous. Instead of manipulating a puppet physically and speaking for it, each child used soft buttons to tell a puppet what to say or do (for example, Greet the other puppet, Play alone) and used soft sliders to modulate its moods (for example, to be a little peppy). The puppets followed the children's directions—more or less. After all, they had minds of their own! Perceiving one another's behavior and their own dynamic internal states, the puppets improvised what they were willing to do and how they were willing to do it. The result was four-way collaborative play: two smart children playing with two smart puppets. We also discovered that playing with Improv Puppets helped children develop key cognitive skills for literacy and social empathy. This was fun—and useful.

Determined to pursue intelligent character technology and applications, but unsure how sponsors would respond, I submitted several proposals for research funding. My usual courses on intelligent agents morphed into courses on intelligent characters and interactive story, several cotaught with Professor Larry Friedlander and cross-listed in the computer science and English departments. At the same time, here in Silicon Valley, the possibility of a "startup" is always in the air. Although no market yet existed for intelligent characters, this technology would enable all kinds of new products and services. So I also submitted a proposal to the NIST Advanced Technology Program (ATP), which funded high-risk, high-promise technology research and development in the private sector. The plan was to continue basic research if the Stanford funding came in or, if the ATP funding came in, to start a company to do applied R&D and commercialization.

As it turned out, all of the funding came in. DARPA, NSF, and Intel funded the Virtual Theater Project at Stanford, allowing my research group to pursue theoretical ideas (for example, personality and social dynamics, the autonomy continuum, mind-body connections, and multicultural issues) and to experiment with new application genres (for example, Tigrito, a virtual stuffed animal; master-servant improvisational sketch actors; Cyber Café, a virtual world with a virtual waiter and smart masks for visitors; Merlin, a children's online art museum guide; and Patient Advocate, a wellness coach for maternity patients). Meanwhile, NIST ATP funding allowed the Extempo team to begin applied R&D to commercialize intelligent characters.

Commercial Intelligent Character Technology

The Extempo team built the Imp Character Technology for creating, deploying, monitoring, analyzing, and reporting on intelligent characters. The patented approach underlying it significantly extends and specializes the architecture conceived and developed at Rand and Stanford. Each Imp character "mind" integrates a plug-and-play "persona" with at least one "role." A persona structures a character's back-story, emotional dynamics, abstract body language, social exchange, and topical chat. A role controls the character's mission-oriented behavior in a specified domain (for example, guiding visitors at a particular website). It organizes behavior with context-sensitive agendas, where each item is a proactive or reactive behavior, which may itself have multiple steps and interactive logic based on speech acts (for example, ask and accept a reply to a yes-no question, ask and accept a reply to an open-ended question, make a statement, reply to a question or statement, tell a story). The architecture supports normal variability of behavior, mood-sensitive behavior, natural language user input, run-time variable binding and instantiation, short- and long-term memory for events and facts, and electronic actions (for example, speak, gesture, control user's browser, read/write to a database, display data, access the Internet). It writes specified variables to a database and logs complete, annotated interaction transcripts.

A typical Imp character projects a vivid personality and plays a useful, easily recognized role. It is goal directed, but responds to the user. It is highly context sensitive. It reflects internal emotional dynamics and may recognize and respond to the user's emotions or other states. It recalls facts and conversations from earlier in a visit or from prior visits and uses these memories to enhance its services and its relationship with the user. It may use

sophisticated interaction techniques, such as motivational interviewing, to encourage behavior change in a user, or social exchange to build a reciprocal relationship with the user. As a result, users interact easily and naturally with Imp characters.

The Imp Mind-Body API allows alternative body technologies (for example, graphics, animation, speech). In a typical application, a user has a mixed-initiative natural language conversation with an expressive animated character. The user freely types dialogue to the character. The character types its dialogue, speaks aloud, and displays its abstract body language through its body technology. However, different applications may use different combinations of communication channels and modalities.

Three dialogues below illustrate the behavior of two Imp characters, Coach Harmony and Patient Nina, in STAR Workshop for training medical and nursing students to do brief intervention for alcohol abuse.

In the first two dialogues, Joan, a medical student, interviews Nina, a virtual patient. Nina has come to primary care for a sprained ankle, but also has a “red flag” for alcohol abuse: 4–5 glasses of wine, 6–7 days a week. Joan must use the flexible brief intervention protocol summarized in table 1 (Miller and Rollnick 2002) to persuade Nina to make a positive change related to her alcohol consumption. Virtual patients like Nina are designed to provide authentic practice (Druckman and Bjork, 1994.) She accepts Joan’s typed natural language inputs and responds appropriately in printed and spoken remarks, complemented by gestures. In particular, when Joan correctly instantiates the intervention protocol, Nina responds compliantly and continues the dialogue. When Joan departs from protocol, Nina resists. In some situations, Nina takes the conversational initiative out of turn. For example, in the first practice, Joan pursues the sensitive topic of alcohol in a vague way. Nina complies, but she is uncomfortable. When Joan pauses too long, Nina gets impatient and terminates the interaction. In the second practice, Joan relates the topic to Nina’s specific health risks and takes her conversational turns promptly. Again, Nina complies, but this time she is more comfortable and allows Joan to continue. Because students must learn to manage patients’ emotional responses, Nina’s body technology uses high-affect photos and voice recordings made by an actress. After each session, Nina records the student’s scores and annotated transcripts for use by Coach Harmony and her human supervisor.

Joan’s First Practice with Virtual Patient Nina

Nina: It sure is a relief to know my ankle is OK.

Joan: On another matter, Nina, can we discuss the relationship between alcohol and your health?

Nina: OK, I guess so. What did you want to talk about?

Joan: There is quite a bit of evidence that drinking alcohol can be a risk factor for certain health problems.

Nina: Um, OK. I didn’t know that.

[Joan pauses too long before continuing.]

Nina: Well, if that’s all, I guess I’ll be going. Thanks for your help with the ankle. Bye-bye.

Joan’s Second Practice with Virtual Patient Nina

Nina: Thanks for checking my ankle. I’m glad it’s not serious.

Joan: Nina, before you go, I would like to talk about how your drinking might be affecting your health.

Nina: Um, all right.

Joan: You know, alcohol might be contributing to your high blood pressure.

Nina: Really? I’ve never heard that before. Everyone in my family has high blood pressure.

Joan: That may be, Nina, but I think excessive drinking is contributing to your high blood pressure.

Nina: What do you mean excessive drinking? I don’t drink that much, compared to a lot of people.

Joan: From a health perspective, 4–5 glasses of wine 6–7 days a week is considered excessive.

Nina: Look, I just have a sprained ankle. I didn’t come here for a lecture. I have to go.

In the third dialogue, Coach Harmony reviews Joan’s second practice with Nina. Designed to provide effective coaching (Druckman and Bjork 1994), Harmony begins by making a positive assessment, summarizing Joan’s scores, and noting Joan’s improvement. She then reviews Joan’s performance in detail. For each protocol step, Harmony reminds Joan of the requirements, gives specific feedback on strengths and weaknesses of Joan’s performance, quotes Joan’s exchanges with Nina, and highlights effects on Nina’s compliance, resistance, and emotions. Harmony calls out and praises Joan’s specific improvement over prior practices. She gives constructive advice and offers targeted coaching before each new practice. Because Harmony must construct her behavior and dialogue to incorporate unanticipated run-time information, her body technology uses animation and text-to-speech synthesis. She also records complete annotated transcripts of her interactions with students.

Harmony Gives Joan Feedback on Second Practice with Nina

Joan, that was a good second practice with Nina. You scored 4/12 points, up from 3/12 last time. Let’s review your practice session, one step at a time.

As in your first practice session, you started off well,

Joan's Scores on Her Two Practices with Nina	P1	P2
1. Introduce the topic in a neutral health context.	2/2	2/2
2. Inform the patient of her specific health risks.	1/2	2/2
3. Invite the patient to express her concerns.	0/2	0/2
4. Acknowledge the patient's stated concerns.	0/2	0/2
5. Invite the patient to make an appropriate change.	0/2	0/2
6. Ask the patient to commit to the positive change	0/2	0/2
Total	3/12	4/12

Table 1. Joan's Scores in Two Practice Sessions.

by introducing the topic of Nina's drinking in a neutral health context. You said, "Nina, before you go, I would like to discuss how drinking might be affecting your health." As a result, Nina agreed to discuss the topic. She said, "Um, all right." Good job, Joan.

As in the first session, you also correctly informed Nina of health risks associated with alcohol. But this time, you mentioned her specific risks. You said, "You know, alcohol might be contributing to your high blood pressure." That got Nina's attention and drew her into the conversation. She said, "Really? I've never heard that before. Everyone in my family has high blood pressure." Good improvement, Joan.

However, you let Nina's comment draw you into an argument. You replied, "That may be, Nina, but I think excessive drinking is contributing to your high blood pressure." That was a reasonable medical assessment, but it made Nina feel defensive. She argued, "What do you mean excessive drinking? I don't drink that much, compared to a lot of people." Next time, Joan, keep the intervention on track by sticking to your protocol. Just accept Nina's comments and move on to step 3, inviting her to express her concerns.

Unfortunately, Nina's comeback snared you again. You made another valid point, "From a health perspective, 4–5 glasses of wine 6–7 days a week is considered excessive." But it only made Nina feel more defensive. She got upset and terminated the interview. She said, "Look, I just have a sprained ankle. I didn't come here for a lecture. I have to go."

Let's sum up. Joan. You mastered the first two protocol steps with Nina. Good job. But you got drawn into an argument with her. That happens to a lot of students. I'd like to coach you on how to avoid arguments and keep the intervention on track. Then you can try another practice session with Nina. How does that sound?

In STAR Workshop, Coach Harmony guides each student to mastery of the target intervention proto-

col, defined as error-free performance on a representative sample of virtual patients—a higher standard of competency than what is practical in conventional training. To achieve mastery, Harmony gives the student as many practice sessions as needed with each patient, providing detailed feedback and targeted coaching after each one. Virtual patients give authentic practice in repeat sessions by varying incidental features of their dialogue and behavior, while remaining in character. For most students, errors decrease on successive practice sessions with a patient, but students vary in specific error patterns. After a student achieves mastery with a given patient, Harmony introduces a new patient and iterates the same practice-feedback-coaching process. Errors generally decrease on successive patients, but again, students vary in specific patterns. Also, later patients challenge students in new ways and may elicit new errors. By tailoring coaching and practice to performance, Harmony individually optimizes the learning path, enabling each student to achieve mastery in a relatively short amount of time.

As illustrated in the examples, Imp characters engage users in mixed-initiative conversation. But they do not perform deep natural language processing. Instead, they use context-sensitive recognition and instantiation of parameterized linguistic patterns. This allows characters to handle a wide range of conversational content in real time. It also enables writers to create engaging dialogue and distinctive linguistic styles for individual characters.

Imp Authoring Tools are designed to enable designers and writers to create characters with vivid personalities and role-specific expertise, without significant programming. Figure 1 illustrates authoring for a small part of Nina's behavior. In a prospective run-time context, the student communicates that alcohol can affect a particular health risk. Nina has already agreed to discuss the role of alcohol in her health, but prior events have put her in a defensive mood. If this context occurs, Nina will speak one of the three alternative dialogues and tell her body technology to perform a defensive argue gesture. If the context occurs in multiple practice sessions, she will use alternate dialogues and gestures. Each time, her mood will grow more defensive than it is at context instantiation and she will instantiate the log cue in the student's transcript. If the student's next comment is protocol-compliant and Nina's mood has not grown too defensive, she will continue the interaction and her mood will improve. However, given the "assertion" action type, she will wait only 3–5 seconds for the student to reply before taking the initiative herself.

Designed for commercial deployment, the Imp Engine is lean, efficient, and robust. In conservative benchmarks, it supports multiple applications

Context	Student said: #alcohol #affects health-risk
Mood	defensive
Precondition	topic-ok = yes
Precondition	#my-specific-health-risks includes health-risk
Dialogues	1. Really? I've never heard that before. Everyone in my family has health-risk. 2. I've had health-risk all my life, long before I started drinking. 3. Alcohol and health-risk? I don't see that.
Gesture	argue
Action Type	assertion
New Mood	more defensive
Log Cue	argued-about-relation-to health -risk

Figure 1. Sample of Dialogue Authoring for Nina

and hundreds of concurrent users per CPU; excess users increase response times. Capacity scales linearly in number of CPUs and degrades gracefully on hardware failure. Imp Admin Tools aid system management and online application modification. Imp Log Analysis Tools produce specifiable reports.

Bootstrapping Early Successes

With technology development proceeding, the Extempo team applied for and won an award for business planning from the California Technology Innovation Program. We used it to explore applications, markets, and value propositions for commercial uses of intelligent characters. Instead of the enterprise software model, which was predominant at the time, Extempo offered a software services model. Exploiting technical and conceptual modularity, we offered use of customers used the Imp Engine to run applications on our hosted site, with transparent deployment for users on their own sites and monthly fees based on usage. Some customers used Imp Authoring and Log Analysis Tools, as well as simpler, more specialized tools, with fees set to motivate development and deployment of applications. Extempo also offered custom application development, management, and reporting services, with time and materials fees. Working with early adopters, we created a variety of demonstrations and applications. For example:

Virtual Mr. Clean offered household cleaning tips,

satisfied customer curiosity, built customer relations, and collected marketing data for Procter & Gamble.

Jack, a friendly pup, entertained Petopia.com visitors, taught them about pets, and collected marketing data.

Max, a wisecracking hound, greeted visitors at PeopleWeb.com and assisted portal novices.

Web master and conspiracy theorist Lenny Pochnik dropped clues and taunted game players for La Fong Interactive, Microsoft, and Time-Warner.

Erin, a sassy Texas bartender, mixed virtual drinks, chatted up players, and offered game clues for AOL.

Students in an experiential learning course at the Naval Postgraduate School interviewed Robert and James, former members of fictional terror groups.

Coach Harmony and several virtual employees trained Sun managers in mandated protocols for performance review.

Coach Harmony and several virtual patients trained Stanford University, University of San Francisco, and San Jose State University medical and nursing students in evidence-based protocols for brief intervention in alcohol abuse.

Imp characters achieved their application goals. Users complied with suggestions and requests, gave useful information, answered questions, made repeat visits, stayed to chat, and offered praise. Characters like Mr. Clean and Jack increased website stickiness and persuaded users to provide a variety of useful inbound marketing data. In con-

trolled studies of alcohol intervention and management training, Coach Harmony and virtual role-players were more effective than alternative training methods, including human coaches and role-players, standard web training, or no training. For example, in final assessments with human test patients, eight weeks after training, students trained with Imp characters in STAR Workshop achieved 94 percent accuracy, compared to 50 percent for web training or no training ($p < .01$).

Our customers were satisfied with early results and interested in continuing and expanding their applications. This was critical for Extempo. Since we had not raised significant venture capital, we needed revenue to support operations.

In fact, Extempo never had enough resources to achieve commercial success. Starting in 1996, we spent two years developing technology, doing a few small commercial projects, and writing a business plan for the emerging e-commerce market. We did not raise seed capital, hire professional sales staff, or start to grow the company until 1998. Despite successes with early visionary customers, including large corporations and dot-com newcomers, we scrambled to catch up with competitors who had raised significant venture capital. These companies had pursued a different strategy: selling a large percentage of equity to investors early, growing their companies aggressively before earning substantial revenues, investing in expensive marketing campaigns to gain exposure, eschewing sophisticated functionality, and giving away simple products and services to rapidly growing lists of customers.

The Bubble Bursts—The Market Evaporates

The eventual bursting of the Internet bubble and market downturn affected many companies, including Extempo. When customer prospects and investment capital dried up, we could not afford to maintain our 30-person team. We lost 25 great people overnight.

Extempo almost died in 2000, but remained in business for two reasons. First, we had not yet sold enough equity to lose control of the company. Only people who were passionate about the technology would persist under such daunting circumstances. Second, we had already won a second NIST ATP grant to develop our technology for online learning.

So we started over. Again, we developed technology and worked with visionary customers. With investment capital no longer available, we applied for and won funding from the NIH SBIR program. We were inspired by the potential of online learning and excited by the success of our early applications. However, it became increasing-

ly difficult to sustain even our small team in a slow economy. When two customers simultaneously defaulted on their financial commitments, we had to put Extempo to rest.

Lessons Learned

Although we can only speculate about paths not taken, I believe that Extempo would have benefited from a more aggressive business development strategy. With our technology-push culture and funding, we spent too much time exploring the “possibility space” before going to market. We were visionaries, planning (hoping) to staff the web, introduce innovative toys and game genres, provide authoring tools to creative teachers, and populate the first generation of virtual worlds with entertaining denizens. By contrast, a successful serial entrepreneur advised me, “Go out and sell something. Then come back and build it.”

There is a middle path. Extempo could have gone to market early with a simple product that satisfied a recognized market need. With a lower price point and shorter sales cycle, this strategy would have generated feedback for developers, some good reference customers, and an installed base. A robust early revenue stream might have financed the company’s survival.

On the other hand, Extempo had unavoidable bad timing. Along with half a dozen large and small competitors, we rode a tide of optimism. We built new technology, offered new products and services, operated on new infrastructure, and promoted new business models. Then the market crashed. Most of the small competitors are gone and the larger ones appear to have abandoned the endeavor. It is worth noting that most of these competitors avoided Extempo’s greatest handicap, insufficient capital, either by raising more venture capital or using internal R&D funds. Nonetheless, they did not succeed.

Some technologies require several generations of ventures before they find success in the market. I’m guessing that this will be true for intelligent characters. Their day will come.

Postscript: Lifelike Solutions, Inc.

Lifelike Solutions, Inc. was founded on October 1, 2007, with a core team including Rami Saker, Satoru Isaka, Shirley Tessler, and myself. Still passionate about the Imp technology (purchased during Extempo’s formal liquidation), but following the admittedly speculative “lessons learned,” we are focusing on applications serving recognized market needs in the health industry. Building on earlier successes, Lifelike will offer two lines of STAR Workshops, one training health-care providers in clinical interviewing skills and the

How Inappropriately Heavyweight AI Solutions Dragged Down a Startup

(and Made Me Realize that Industrial Salaries Are High for a Good Reason)

Michael Wooldridge

The basic idea was to build distributed information systems, with intelligent nodes at each repository, able to reflect upon and reason about repository content in order to route queries appropriately. We came up with a heavyweight agent architecture, using ideas from AI planning and robotics. These sorts of architectures were very much in vogue at the time, and the company wanted its own, proprietary technology. We started thinking about programming languages for the agents and the kinds of knowledge representation and reasoning that would be required. We spent a lot of time and money flying from London to the U.S. West Coast, talking to patent lawyers.

It transpired that the architecture, its decision-making and action models, were completely inappropriate for the problem at hand. By the time we realized we should have been focusing on basic software engineering, quality assurance, and end-user requirements, the company had burned out much of the goodwill—and most of the funds—of our investors. I jumped ship after less than a year, back to academic life (on a much reduced salary). The company folded about six months later, in the spring of 1998 (about the same time that Page and Brin were launching Google).

The experience was traumatic, but I learned a great deal about software engineering, the real priorities of companies, and the kinds of things that do and do not make money.

Michael Wooldridge has been a professor of computer science at the University of Liverpool since 1 January 2000. His research is primarily in the theoretical foundations of multiagent systems; he was the recipient of the ACM SIGART Autonomous Agents Research Award in 2006 and is coeditor-in-chief of the journal *Autonomous Agents and Multi-Agent Systems*.

other training consumers in health-related interpersonal skills. In a new product line, Lifelike will offer CoachBots, providing consumers long-term coaching in health-related behaviors. For example, CoachBot Harmony (moonlighting outside her STAR Workshop role) coaches adults to start and stick with a Basic Good Health program for eating, exercise, and play. Here is an excerpt from her first session with Margo, a new client:

Harmony: In the Basic Good Health program, we can focus on improving your weight, fitness, or disease risk. Which one of these is most important to you right now, Margo?"

Margo: I want to reduce my disease risk.

Harmony: Disease risk. Me too. You know Margo,

cancer, diabetes, and heart disease are the top causes of death for women. But here is the good news. You can reduce your risk for all of them with a healthy lifestyle. I have heart disease in my family, so I want to reduce my risk for that. How about you, Margo? Are you most concerned with diabetes, cancer, or heart disease?

Margo: Heart disease.

Harmony: Heart disease. OK. We'll start working on that today. All three parts of the Basic Good Health program will help reduce your heart disease risk. I'll also set up your log to track your risk factors—blood pressure, heart rate, cholesterol, and weight. Then I'll figure out which lifestyle activities are most effective at controlling these factors for you so I can tailor your program. Sound good?

Please Join Us in
Washington, DC for the

AAAI-08 Fall Symposium Series

7–9 November 2008

Westing Arlington Gateway
Arlington, Virginia

Details:

www.aaai.org/Symposia/Fall/fss08.php



Margo: Sure does.

Harmony: I aim to please, Margo.

Margo: Am I too fat to exercise?

Harmony: In a word, Margo—no. It's true that your BMI is in the top range. You might want to talk to your doctor about that. But we are going to start you off slowly, with a nice, safe walking program. We'll build up your exercise gradually, based on your progress. OK?

Margo: Yeah.

Harmony: Excellent!

Extempo was founded in 1995 and restarted in 2000. So I consider Lifelike Solutions to be our third try. Third time's the charm!

Acknowledgements

Many colleagues contributed to the work at Extempo. Perry W. Thorndyke was a founding director. Robert van Gent and M. Vaughan Johnson led engineering, which included Louise Crow, Tom Sephton, Dmitri Bobrovnikoff, Allyn Treshansky, Ian Flanigan, Don Hayler, Danny Hui, and Mitch Ferguson. Character authors included Karen Amano, Hallie Kushner, Winter Mead, Rem Reynolds, and John Calhoun. Michaela Schlocker and Ruth Duran Huard led analytics. Project managers included Keith Wescourt, Corey Fugmen, Satoru Isaka, and Rami Saker. The business team included Clyde Hinshelwood, Kimberly Thomas, Mark Chantland, James Regan, Jean-François Noubel, and Dave Whipple. I regret any omissions and wish to thank everyone for their contributions. I also wish to thank sponsors of the Virtual Theater Project at Stanford (DARPA, NSF, Intel, and Stanford) and Extempo's

sponsors and investors (NIST ATP, NIH NIAAA, CalTIP, Sipl MacDonal Ventures, Angel Investors, and other individuals).

References

- Bates, J. 1994. The Role of Emotion in Believable Agents. *Communications of the ACM* 37(7): 122–125.
- Druckman, D., and Bjork, R. A. 1994. *Learning, Remembering, Believing: Enhancing Human Performance*. Washington, DC: National Academy Press.
- Hayes-Roth, B. 1985. A Blackboard Architecture for Control. *Artificial Intelligence* 26(3): 251–321.
- Hayes-Roth, B. 1995. An Architecture for Adaptive Intelligent Systems. *Artificial Intelligence* 72(1–2): 329–365.
- Hayes-Roth, B., and Larsson, J. E. 1997. Domain-Specific Architectures for Software Reuse. Unpublished paper.
- Hayes-Roth, B., and van Gent, R. 1997. Story Making with Improvisational Puppets. *Proceedings of the International Conference on Autonomous Agents*, Los Angeles. New York: Association for Computing Machinery.
- Johnstone, K. 1987. *Impro: Improvisation and the Theatre*. Toronto, Ontario: Theatre Arts Books.
- Miller, W. R., and Rollnick, S. 2002. *Motivational Interviewing: Preparing People for Change*. London: Guilford Press.

Barbara Hayes-Roth holds a Ph.D. in psychology from the University of Michigan. She is a Fellow of the Association for the Advancement of Artificial Intelligence, and holds two patents on intelligent characters. She founded two companies, Extempo Systems, Inc. and Lifelike Solutions, Inc. She also worked in the Computer Science Department at Stanford University, at Rand Corporation, and at Bell Laboratories Murray Hill. Hayes-Roth can be contacted at bhr@lifelikesolutions.com or bella@hayes-roth.net.