

Modeling Interactive Agents in ALIVE

Pattie Maes, Bruce Blumberg, Trevor Darrell, Alex Pentland, Alan Wexelblat

MIT Media Laboratory

20 Ames Street

Cambridge, MA 02139

pattie/bruce/trevor/sandy/wex@media.mit.edu

In this video we discuss the design and implementation of a novel system which allows wireless full-body interaction between a human participant and a graphical world inhabited by autonomous agents. The system is called "ALIVE," an acronym for Artificial Life Interactive Video Environment [5]. One of the goals of the ALIVE project is to demonstrate that virtual environments can offer a more emotional and evocative experience by allowing the participant to interact with animated characters which have complex behaviors and which react to the user and the user's actions in the virtual world.

The ALIVE system has been demonstrated and tested in several public forums. It was demonstrated for 5 days at the SIGGRAPH-93 Tomorrow's Realities show in Anaheim, California and for 3 days at the AAAI-94 Art Show in Seattle, Washington. The system is installed permanently at the MIT Media Laboratory in Cambridge, Massachusetts. It will feature in the Ars Electronica Museum, currently under construction in Linz, Austria and the ArcTec electronic arts bienale in Tokyo, Japan, May 1995.

In the style of Myron Krueger's Videoplace system, the ALIVE system offers an unencumbered, full-body interface to a virtual world [4]. The ALIVE user moves around in a space of approximately 16 by 16 feet. A video camera captures the user's image and removes the background environment; thus, no blue-screens or other special walls are needed. The separated outline is then composited into a 3D graphical world. The resulting scene is projected onto a large (approximately 10'x16') screen which faces the user and acts as a "magic mirror:" the user sees him/herself in the environment, surrounded by objects and agents. No goggles, gloves, or tethering wires are needed for interaction with the virtual world.

Computer vision techniques are used to extract information about the person, such as where in the space the person stands and the position of various body parts. A pattern-matching technique called dynamic time-warping is used to recognize simple gestures as they are performed. ALIVE combines active vision and domain knowledge to achieve robust real-time performance [5].

The user's position as well as hand and body gestures are used as input to affect the behavior of agents in the virtual world. Agents have vision sensors which allow them to "see" the user and react to gestures such as

pointing or throwing a virtual ball. The user receives visual (on the big screen) and auditory (prerecorded sound) feedback about the agents' internal state and reactions. Agents have a set of needs and motivations, a set of sensors to perceive their environment, a repertoire of activities which they can perform and a physically-based motor system that allows them to move in and act on the environment. The behavior system decides in real-time which activity the agents should engage in so as to meet their internal needs and to take advantage of opportunities presented by the current state of the environment.

The system allows a direct-manipulation style of interaction in which users interact directly with the environment – such as pushing buttons or moving objects – and also an indirect style of interaction in which users give agents commands and the agents carry out the commands based on the user's input, the current environment, and their internal state. For example, the meaning of a gesture is interpreted by the agents based on the situation the agent and user find themselves in. When the user points away from herself, and thereby "gives the command" to send a character away, the character responding to the command will go to a different place in the virtual environment depending on where the user is standing and which direction she is pointing. In this manner, a relatively small set of gestures can be employed to mean many different things in many different situations.

The ALIVE system incorporates a tool, called "Hamsterdam" [2] [3], for modeling semi-intelligent autonomous agents that can interact with one another and with the user. Hamsterdam produces agents that respond with a relevant activity on every time step, given their internal needs and motivations, their past history and the environment they perceive with its attendant opportunities, challenges and changes. Moreover, the pattern and rhythm of the chosen activities is such that the agents neither dither between multiple activities, nor persist too long in a single activity. They are capable of interrupting a given activity if a more pressing need or an unforeseen opportunity arises.

The Hamsterdam activity model is based on elements taken from animal behavior models initially proposed by ethologists. In particular, several ethological concepts such as behavior hierarchies, releasers, fatigue, and so on have proven to be crucial in guaranteeing the robust and flexible behavior required by autonomous

interacting agents [2] [3].

When using Hamsterdam to build a creature, the designer specifies the sensors of the agent, its motivations or internal needs, and its activities (behaviors) and actions (motor system movements necessary to fulfill a behavior). Given that information, the Hamsterdam software automatically infers which of the activities is most relevant to the agent at a particular moment in time according to the state of the agent, the situation it finds itself in, relevant input from the environment, and its recent behavior history. The observed actions of the agent are the final result of numerous potentially executable behaviors competing for control of the agent. The activities compete on the basis of the value of a given activity to the agent at that instant, given the above factors. The details of the behavior model and a discussion of its features are reported in [2] [3].

The most sophisticated creature built so far is a dog called Silas. Silas's behavioral repertoire currently includes following the user, sitting when asked by the user, going away when ordered by the user to do so, and performing other tricks such as standing on his hind legs, fetching a ball, lying down and shaking paws. Silas also will chase the Hamster if the latter creature is introduced into the same virtual environment as the dog. Along with visual sensors and feedback, the ALIVE environment also uses sound. Silas provides auditory output in the form of a variety of prerecorded samples.

The ALIVE system demonstrates that entertainment and the effort to model believable creatures in simple virtual environments can be a challenging and interesting application area for autonomous agents research. ALIVE provides a novel environment for studying architectures for intelligent autonomous agents. As a testbed for agent architectures, it avoids the problems associated with physical hardware agents or robots, but at the same time forces us to face non-trivial problems, such as noisy sensors and an unpredictable, fast-changing environment. It makes possible our study of agents with higher levels of cognition, without oversimplifying the world in which these agents live.

ALIVE represents only the beginning of a whole range of novel applications that could be explored with this kind of system. We are currently investigating ALIVE for interactive storytelling applications in which the user plays one of the characters in the story and all other characters are artificial agents which collaborate to make the story move forwards (for more exposition on this topic, see the three short papers on ALIVE in [1]). Another obvious entertainment application of ALIVE is video games. We have hooked up the ALIVE vision-based interface to existing video game software, so as to let the user control a game with his full body. In addition, we are investigating how autonomous video game characters can learn and improve their competence over time, so as to keep challenging a video game player. Finally, we are modeling animated characters that teach a user a physical skill in a personalized way. The agent is modeled as a personal trainer that demonstrates to the user how to perform an action and provides personalized and timely feedback to the user, on the basis of the sensory information about the user's gestures and body positions. The ALIVE system shows that animated characters that are based on Ar-

tificial Life models can not only look convincing – that is allow suspension of disbelief on viewing – but can act and interact in a realistic enough manner to maintain this suspension of disbelief during unpredictable real-time interaction with users.

Acknowledgements

Alan Wexelblat spent many long hours shooting and editing the video. Thanks Alan!

References

- [1] Joseph Bates, Barbara Hayes-Roth and Pattie Maes, Workshop notes of the AAAI Spring Symposium on Interactive Story Systems: Plot and Character, AAAI, March 1995.
- [2] Bruce Blumberg, "Action Selection in Hamsterdam: Lessons from Ethology", Proceedings of the 3rd International Conference on the Simulation of Adaptive Behavior, Brighton, August 1994, MIT-Press.
- [3] Bruce Blumberg and Tinsley Gallyean, "Multi-Level Direction of Autonomous Creatures for Real-Time Virtual Environments", Proceedings of the Siggraph 1995 conference, Los Angeles, CA, August 1995.
- [4] Krueger M.W., Artificial Reality II, Addison Wesley, 1990.
- [5] Pattie Maes, Trevor Darrell, Bruce Blumberg and Alex Pentland, "The ALIVE system: Full-body Interaction with Autonomous Agents", Proceedings of the Computer Animation '95 Conference, Geneva, Switzerland, IEEE-Press, April 1995.