

A Believable Agent for First-Person Shooter Games

Dongkyu Choi, Tolga Konik, Negin Nejati, Chunki Park, and Pat Langley

Computational Learning Laboratory

Center for the Study of Language and Information

Stanford University, Stanford, California 94305

{dongkyuc, konik, negin, lovelive}@stanford.edu, langley@csl.stanford.edu

Abstract

In this paper, we present a principled approach to constructing believable game players that relies on a cognitive architecture. The resulting agent is capable of playing the game *Urban Combat* in a plausible manner when faced with similar situations as its human counterparts. We discuss how architectural features like goal-directed but reactive execution and incremental learning can produce more believable synthetic characters.

Introduction

The popularity of first-person shooter games has motivated developers to make them more realistic than ever before, which requires more believable computer-controlled agents. Until recently, however, they have been either too restricted or too informed to be believable.

There have been some efforts to construct more believable agents, with TacAir-Soar (Jones et al., 1999) being an excellent example in the domain of simulated aerial combat. Its developers proposed three principles for generating human-like behavior, which include the same basic processing structure, the same sensory and motor systems, and the same basic knowledge as humans. They used the Soar cognitive architecture (Laird et al., 1987) to develop their agent. Our work employs a different architecture, ICARUS, that shares some features with Soar but that also introduces novel ideas. The ICARUS agent is given sensors with limitations similar to those of human perceptions, along with actions much like those available to human players.

However, we believe that a human-like agent requires additional features. These include a hierarchical organization of knowledge with multiple levels of abstraction that provides a rich, human-like vocabulary and strategies. The agent should also have goal-directed but reactive behaviors that let it maintain achieved subgoals while working on others. Finally, the system should learn incrementally, so that it responds situations differently as its experience accumulates. Combined with appropriate domain knowledge, these features should produce agents that are more plausible and believable to human players.

Copyright © 2007, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

In the following sections, we review our testbed, describe the architecture and the agent implemented within it, and consider the ways in which our architectural assumptions support more human-like, believable behavior. We also explain how the agent performs and learns in a goal-oriented manner, then report some experiences with the system. Finally, we discuss related work and draw some conclusions.

First-Person Shooter Games

One popular genre – first-person shooter games – are characterized by displaying the screen view in the main character’s point of view and by the operation of handheld weapons. They are extremely realistic when compared to more traditional research domains like puzzles. Also, their flexible structure facilitates creation of different missions by modifying maps, items, and obstacles for the environment.

For our research, we used as a testbed *Urban Combat*, which is a stand-alone modification of *Quake 3 Arena*.¹ The ICARUS architecture interfaces with the testbed through shared memory to receive perceptions and apply actions in the simulated world. *Urban Combat* provides several challenging scenarios that involve different levels of task difficulty. Among many possible missions, we focused on ones that required the agent to find and defuse Improvised Explosive Devices (IEDs) at various locations.

An Architecture for Intelligent Agents

Our architecture for intelligent agents, ICARUS, shares many assumptions with previous frameworks like Soar (Laird et al., 1987) and ACT-R (Anderson, 1993). It distinguishes short-term and long-term memories, populates these memories with symbolic list structures, and operates in recognize-act cycles. Also, ICARUS’ learning is incremental and interleaved with performance. Many of these ideas have their basis on psychological theories of human intelligence.

Despite these shared assumptions, the architecture also has many novel features. Separate concept and skill memories support different aspects of knowledge, and another memory contains a prioritized list of goals the agent should attempt to achieve. Long-term memories are organized in hierarchies that describe knowledge at different levels of abstraction. ICARUS indexes procedural skills by goals they

¹A first-person shooter game by Silicon Ice Development.

achieve, which plays a key role in its retrieval mechanisms and its incremental learning processes. Furthermore, goal-directed but reactive behavior makes the framework well suited for physical domains like *Urban Combat*.

The architecture operates in cycles, each of which invokes a number of distinct cognitive processes. As Figure 1 shows, on each cycle, ICARUS first perceives the environment and makes inferences about its situation. If the system finds a path through the skill hierarchy that matches against these beliefs and would make progress toward the current goal, it executes the associated actions. This changes the environment and the agent continues to the next cycle. As the architecture carries out this process, it learns new skills that produce more efficient on future problems.

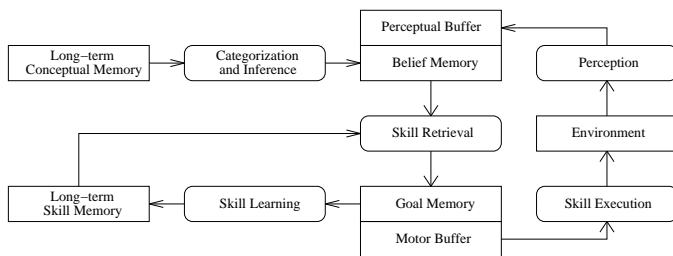


Figure 1: Cycles of ICARUS architecture

We believe that the architectural features above allow more human-like behaviors in embodied agents, while also making it easier to provide background knowledge equivalent to what humans have when they first start playing a game. In general, cognitive architectures have great potential to aid the construction of believable synthetic agents, and they provide a uniform infrastructure for game developers that facilitates applications to many different games.

An Automated Urban Combat Player

Architectural support from the ICARUS framework let us program an automated player for *Urban Combat* with only limited domain-specific engineering. Like human players, the resulting agent can recognize its immediate surroundings, move toward objects, explore the unknown environment systematically, and learn from its experiences. In this section, we describe each of these capabilities and compare them to those of human players.

Representing and Using Spatial Knowledge

Ideally, an agent in a first-person shooter game should recognize objects and their spatial relations from their depiction on the screen, but this would require three-dimensional image understanding. Although an important subject, this type of processing is beyond the scope of our research. Instead, we divide the game environment *region maps* that correspond to distinct but interconnected places, and we provide the agent with perceptual content about objects that reside in the same place. There is strong evidence that humans organize their spatial knowledge into places that have detailed metric information about local objects. This strategy also lets the agent focus attention on the immediate surroundings.

Human spatial knowledge is not limited to places, but also includes routes. Although the region map provides information on the current region, the agent must also know its topological relation of neighboring regions in order to traverse them to reach distant places. To this end, the region map includes *gateway* information that lets the agent can perceive the locations and types of the boundaries to neighboring places. Along with the objects that occur in its current region, the agent perceives the boundaries to each neighboring region, along with characteristics that let it assess their usefulness for the task at hand.

Exploration

The ICARUS agent performs two distinct types of exploration, one in physical space and the other in action space. Given a new map, human players move around and explore the mission map. The synthetic agent uses different exploration skills depending on whether the goal is to reach a region that is outside, inside a building, or on a rooftop. The system also learns topological connections between regions as it encounters them, which it stores as new specific skills that encode route knowledge. The agent prefers more recently acquired skills, which we exploit to ensure prioritized application of the exploration strategies.

Sometimes the agent finds itself in a situation where objects block its path through a gateway or to a goal. Then the system performs exploration in its action space, attempting to find the right activity to overcome the obstacle. We provide the agent with a set of candidate actions that include turning, strafing, jumping, ducking, and climbing. The agent tries each of these, in turn, until it finds the correct method for the current obstacle, in which case it learns a specific skill that it executes directly in future situations.

Learning

Although some games contain characters that appear realistic at first, their static knowledge makes them seem less realistic, challenging, and interesting as the game progresses. A common response is to manually encode multiple behaviors, vary agent parameters, or use learning algorithms with little prior knowledge. These often make the characters more complicated and difficult to play, but not necessarily more believable. In contrast, we believe that the knowledge-driven learning associated with a cognitive architecture can result in synthetic characters that improve their behavior throughout games much as do their human counterparts.

ICARUS uses a learning method, explained in detail elsewhere (Nejati et al. 2006), that uses background knowledge to explain and construct new skills from traces of successful problem solutions. The architecture records the agent's behavior during exploration, along with the states it observed along the way. The system attempts to explain how the agent's steps achieved the goal by chaining off known skills and concepts, accounting for unexplainable changes by instantiating generic templates for basic action models. ICARUS transforms the explanation into skills it can use later to achieve similar goals. These encode general strategies, route knowledge, and methods for overcoming obstacles.

Experience with the Urban Combat Player

Rather than attempting to quantify our agent's believability, we took a more qualitative approach to evaluation. We ran the agent on a variety of scenarios that involved different levels of problem difficulty. Each scenario had the same goal of defusing an IED, but involved different obstacle types, obstacle configurations, and device locations. The agent learned cumulatively over different scenarios at a given level. The initial scenarios relied mostly on knowledge-driven exploration, but still demonstrated goal-directed behavior. For example, if the agent saw the IED, it tried to overcome obstacles that were blocking its path to the device. In later runs, the agent used skills learned from previous scenarios to find paths and overcome obstacles with little or no search.

The agent was able to smoothly combine its exploration strategies with learned knowledge and adapt to unexpected changes. For example, when it knew the location of the IED but started at an unknown location, it explored its surroundings until it found a familiar location and planned a route using learned knowledge. One weakness of the agent is that it retains perfect memories of all visited regions, while human players can require multiple visits before remembering them reliably. However, this was a minor problem for the agent's believability, since *Urban Combat* did not have a large map.

Related Research

Recent years have seen increasing interest in the construction of believable agents for computer games and other simulated environments. Laird and van Lent (2000) present a lucid analysis of the reasons for this trend, so we will not repeat them here. However, we should review some specific systems that bear similarities to the agent we have reported.

Perhaps the closest research comes from Laird (2001), who reports a synthetic agent for the game *Quake*. His system was similar to our own, although his effort emphasized the benefits of anticipating an enemy's movements, which we have not addressed. In other work, Thomas et al. (2003) report an agent that supports reactive and planned navigation in a simulated city. They organized maps hierarchically and model unreliable memory, but their system is not cast within a broader theory of the cognitive architecture.

Other applications of believable agents have included instruction and interactive drama. For example, Rickel et al. (2002) describe an agent that functions as an instructor in an immersive virtual world, incorporating realistic human motion, task-oriented dialogue, and expression of emotions. Mateas and Stern (2005) report *Facade*, a system for interactive drama which controls characters in a simulated apartment that carry out dialogues, get into arguments, and express emotions. Both support abilities for interacting with other agents that we have not addressed in our work.

Conclusions

The popularity of first-person shooter games has motivated developers to build more believable agents, and in this paper we described an architectural response to this challenge. We argued that using a cognitive architecture provides agents with psychologically plausible structures and

processes, making them more believable if they also have similar sensory-motor abilities and knowledge as humans. We also claimed that ICARUS' distinctive features, such as its hierarchical organization of knowledge, goal-directed but reactive behaviors, further support this objective.

Our initial experience with constructing an ICARUS agent for the game *Urban Combat* led to promising results, and we believe that extensions to the architecture can produce even more desirable effects. In future work, we intend to limit perception to reflect human abilities, add the capacity to pursue multiple goals, and support control of multiple game entities. We continue to believe that cognitive architectures offer a promising approach to efficiently creating believable characters for interactive games.

Acknowledgements

This paper reports research sponsored by DARPA under agreement FA8750-05-2-0283. The authors' views and conclusions should not be interpreted as representing policies or endorsements, expressed or implied, of DARPA or the Government. We thank Tom Fawcett, Nicholas Gorski, Larry Holder, John Laird, Cynthia Matuszek, Dan Shapiro, Blake Shepard, and Michael Youngblood for their contributions.

References

- Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Lawrence Erlbaum.
- Jones, R. M., Laird, J. E., Nielsen, P. E., Coulter, K. J., Kenny, P., & Koss, F. (1999). Automated intelligent pilots for combat flight simulation. *AI Magazine*, 20, 27–41.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An Architecture for General Intelligence. *Artificial Intelligence*, 33, 1–64.
- Laird, J. E., & van Lent, M. (2000). Human-level AI's killer application: Interactive computer games. *Proceedings of the Seventeenth National Conference on Artificial Intelligence* (pp. 1171–1178). Austin, TX: AAAI Press.
- Laird, J. E. (2001). It knows what you're going to do: Adding anticipation to a Quakebot. *Proceedings of the Fifth International Conference on Intelligent Agents* (pp. 385–392). Montreal.
- Mateas, M., & Stern, A. (2005). Structuring content in the *Facade* interactive drama architecture. *Proceedings of the Conference on Artificial Intelligence and Interactive Digital Entertainment*. Los Angeles: AAAI Press.
- Nejati, N., Langley, P., & Konik, T. (2006). Learning hierarchical task networks by observation. *Proceedings of the Twenty-third International Conference on Machine Learning*. (pp. 665–672). New York: ACM Press.
- Rickel, J., Marsella, S., Gratch, J., Hill, R., Traum, D., & Swartout, B. (2002). Towards a new generation of virtual humans for interactive experiences. *IEEE Intelligent Systems*, July/August, 32–38.
- Thomas, R., & Donikian, S. (2003). A model of hierarchical cognitive map and human memory designed for reactive and planned navigation. *Proceedings of the Fourth International Space Syntax Symposium*. London.