

Human-Level AI's Killer Application

Interactive Computer Games

John E. Laird and Michael van Lent

■ Although one of the fundamental goals of AI is to understand and develop intelligent systems that have all the capabilities of humans, there is little active research directly pursuing this goal. We propose that AI for interactive computer games is an emerging application area in which this goal of human-level AI can successfully be pursued. Interactive computer games have increasingly complex and realistic worlds and increasingly complex and intelligent computer-controlled characters. In this article, we further motivate our proposal of using interactive computer games for AI research, review previous research on AI and games, and present the different game genres and the roles that human-level AI could play within these genres. We then describe the research issues and AI techniques that are relevant to each of these roles. Our conclusion is that interactive computer games provide a rich environment for incremental research on human-level AI.

Over the last 30 years, research in AI has fragmented into more and more specialized fields, working on more and more specialized problems, using more and more specialized algorithms. This approach has led to a long string of successes with important theoretical and practical advancements. However, these successes have made it easy for us to ignore our failure to make significant progress in building human-level AI systems. Human-level AI systems are the ones that you dreamed about when you first heard of AI: HAL from *2001, A Space Odyssey*; DATA from *Star Trek*; or CP30 and R2D2 from *Star Wars*. They are smart enough to be both triumphant heroes and devious villains. They seamlessly integrate all the human-level capabilities: real-time response, robustness, autonomous intelligent interaction with their environment, planning,

communication with natural language, commonsense reasoning, creativity, and learning.

If this is our dream, why isn't any progress being made? Ironically, one of the major reasons that almost nobody (see Brooks et al. [2000] for one high-profile exception) is working on this grand goal of AI is that current applications of AI do not need full-blown human-level AI. For almost all applications, the generality and adaptability of human thought is not needed—specialized, although more rigid and fragile, solutions are cheaper and easier to develop. Unfortunately, it is unclear whether the approaches that have been developed to solve specific problems are the right building blocks for creating human-level intelligence. The thesis of this article is that interactive computer games are the killer application for human-level AI. They are the application that will need human-level AI. Moreover, they can provide the environments for research on the right kinds of problem that lead to the type of incremental and integrative research needed to achieve human-level AI.

Computer-Generated Forces

Given that our personal goal is to build human-level AI systems, we have struggled to find the right application for our research that requires the breadth, depth, and flexibility of human-level intelligence. In 1991, we found computer-generated forces for large-scale distributed simulations as a potential application. Effective military training requires a complete battle space with tens if not hundreds or thousands of participants. The real world is too expensive and dangerous to use for continual training,

Human-level AI systems are the ones that you dreamed about when you first heard of AI: HAL from 2001, A Space Odyssey; DATA from Star Trek; or CP30 and R2D2 from Star Wars.

and even simulation is prohibitively expensive and cumbersome when fully manned with humans. The training of 4 pilots to fly an attack mission can require over 20 planes plus air controllers. The military does not even have a facility with 20 manned simulators, and if it did, the cost in personnel time for the other pilots and support personnel to train these four pilots would be astronomical. To bypass these costs, computer-generated forces are being developed to populate these simulations. These forces must integrate many of the capabilities we associate with human behavior—after all, they are simulating human pilots. For example, they must use realistic models of multiple sensing modalities, encode and use large bodies of knowledge (military doctrine and tactics), perform their missions autonomously, coordinate their behavior, react quickly to changes in the environment, and dynamically replan missions. Together with researchers at the University Southern California Information Sciences Institute and Carnegie Mellon University, we set off to build human-level AIs for military air missions (Tambe et al. 1995). In 1997, we successfully demonstrated fully autonomous simulated aircraft (Jones et al. 1999), and research and development continues on these systems by Soar Technology, Inc. Although computer-generated forces are a good starting application for developing human-level AI, there are extremely high costs for AI researchers to participate in this work. It requires a substantial investment in time and money to work with the simulation environments and to learn the extensive background knowledge, doctrine, tactics, and missions. Furthermore, much of the current funding is for building and fielding systems and not for conducting research.

Computer Games

In late 1997, we started to look for another application area, one where we could use what we learned from computer-generated forces and pursue further research on human-level intelligence. We think we have found it in interactive computer games. The games we are talking about are not Chess, Checkers, Bridge, Othello, or Go, which emphasize only a few human capabilities such as search and decision making. The types of game we are talking about use the computer to create virtual worlds and characters for people to dynamically interact with—games such as *Doom*, *Quake*, *Tomb Raider*, *Starcraft*, *Myth*, *Madden Football*, *Diablo*, *Everquest*, and *Asheron's Call*.

Human-level AI can have an impact on these games by creating enemies, partners, and sup-

port characters that act just like humans. The AI characters can be part of the continual evolution in the game industry toward more realistic gaming environments. Increasing realism in the graphic presentation of the virtual worlds has fueled this evolution. Human-level AI can expand the types of experiences people have playing computer games by introducing synthetic intelligent characters with their own goals, knowledge, and capabilities. Human-level AI can also recreate the experience of playing with and against humans without a network connection. Current players of computer games are driven to networked games because of the failings of the computer characters. In massively multiplayer online games, human-level AIs can populate the worlds with persistent characters that can play the game alongside humans, providing opportunities for interesting interactions that guide players in the game and enhance the social dynamics between players. Our hypothesis is that populating these games with realistic, human-level characters will lead to fun, challenging games with great game play.

From the AI researcher perspective, the increasing realism in computer games makes them an attractive alternative to both robotics in the real world and homegrown simulations. By working in simulation, researchers interested in human-level AI can concentrate on cognitive capabilities and finesse many of the pesky issues of using real sensor and real motor systems; they must still include some sensor modeling to get realistic behavior, but they don't have to have a team of vision researchers on their staff. They can pursue AI research in worlds that are becoming increasingly realistic simulations of physical and social interactions, without having to create these worlds themselves. Computer games are cheap (\$49.95), reliable, and sometimes surprisingly accessible, with built-in AI interfaces. Moreover, computer games avoid many of the criticisms often leveled against simulations. They are real products and real environments on their own that millions of humans vigorously interact with and become immersed in. Finally, unlike military simulations, we do not need to hunt out experts on these games; they surround us.

Another reason for AI researchers to work in computer games is that if we don't start working in this area, the computer game industry will push ahead without us (Woodcock 2000). Already there are at least five AI Ph.D.s working in the industry (Takahashi 2000). AI researchers have the opportunity to team with an aggressive, talented, and caffeine-charged industry in the pursuit of human-level AI. Here is a list of

reasons for AI researchers to take the computer game industry seriously (Laird 2000a).

First, computer game developers are starting to recognize the need for human-level AI. Synthetic human-level characters are playing an increasingly important role in many genres of computer games and have the potential to lead to completely new genres.

Second, the computer game industry is highly competitive, and a strong component of this competition is technology. AI is often mentioned as the next technology that will improve games and determine which games are hits. Thousands of new computer games are written every year with overall development time averaging nine months to two years, so technological advances sweep through the industry quickly. Already, many computer games are marketed based on the quality of their AI. This field is one in which AI will have a significant impact.

Third, game developers are technologically savvy, and they work hard to stay current with technology. *AI programmer* is already a common job title on game development teams.

Fourth, the game industry is big. In terms of gross revenue, the computer game industry is bigger than the movie industry (Croal and Totilo 1999).

Fifth, computer game hardware is going to provide cheap, high-end computation power for AI in computer games in the next five years. The newest PC 3D video boards and the next-generation consoles, such as Sony's Playstation 2 and Microsoft's X-box, move the entire graphics pipeline off the increasingly powerful central processing unit, freeing it for AI. It is not at all unthinkable that in five years, there will be dedicated AI processors in game consoles—we just have to tell them what we need.

Sixth, computer games need help from academic AI. The current emphasis in computer game AI is on the illusion of humanlike behavior for limited situations. Thus, most, if not all, of the current techniques that are used for controlling game AIs (such as big c functions or finite-state machines) will not scale up. However, just as computer game graphics and physics have moved to more and more realistic modeling of the physical world, we expect that game developers will be forced into more and more realistic modeling of human characters. Moreover, as researchers, we can get a step ahead of the game designers by using their environments for research on human-level AI.

One thing that is missing in the computer game field is significant research funding. Some of the military funding to support computer-generated forces is spilling over to com-

puter games research,¹ and some of the biggest computer game companies (for example, Electronic Arts in England and Sony) have started research centers that include research in AI. More funding could become available as more game developers discover they need help with the AI in their products to push for a competitive advantage. Much of the research could get done in nontraditional ways, with the involvement of undergraduates, game developers, and game players. Thus, we can move AI research out of the labs and into the hands of millions.

Related Research on Computer Games

Other researchers have argued that great game play comes from “believable” agents. These agents don't necessarily have to be human level in their intelligence, as long as they have a façade of intelligence supported by great personality. Joe Bates's (1992) OZ research group at Carnegie Mellon University and Barbara Hayes-Roth's group at Stanford University (Hayes-Roth and Doyle 1998) have worked on developing believable agents for interactive fiction and related computer games. Their research emphasized personality, AI agent-to-human interaction, and shallow but broad agents. We think these aspects are important but want to emphasize that computer games provide an arena for attempting to also build knowledge-rich, complete, integrated AI that incorporates many “deep” capabilities.

John McCarthy has also argued that interactive computer games should be considered as a topic of study for AI, where we can study how an AI system could play a game (his example is LEMMINGS, JR., a real-time scheduling and resource-allocation game) and solve problems that a human attempts.² Other researchers have used other computer games such as *Pengi* (Agre and Chapman 1987) and *SIMCITY* (Fasciano 1996). Our extension is to propose research on the AI characters that are part of the game. Clearly, these efforts are related because human-level AI characters often require the skills of human players. One advantage of creating game characters is that we can influence how games are made and played. The current emphasis on violence in computer games is partially owed to the inability of the AI in these games to support more interesting social interactions. Although computer games will probably always include violence, human-level AI in games will give the game designers freedom to explore other forms of player-character interaction.

RoboCup (Asada et al. 2000) is another relat-

...
***computer
game
hardware is
going to
provide cheap,
high-end
computation
power for AI
in computer
games in the
next five
years.***




Figure 1. AI Roles in Game Genres with Illustrative Links to Their Associated Research Problems.

ed project where competitors develop AI systems to defeat other AI systems in both real robotic and simulated soccer games. In RoboCup, the goal is to build the best soccer-playing robots, not to create the best game play or humanlike behavior. RoboCup is stimulating the development of integrated systems but none with the variety of capabilities we expect to see in interactive computer games.

Computer Game Genres

In this section, we review the major genres of computer games to which human-level AI is relevant. There are other game genres, such as hunting games, fishing games, and lifelike creatures games (Stern 1999), where deer-level, fish-level, or dog-level AI is necessary. For each of the genres in this section, we discuss the different roles that human-level AI can play: enemies, partners, support characters, strategic opponents, low-level units, and commentators. Other roles are possible, but these are the most common. In the following sections, we go through these roles and discuss how AI could improve the games and how these games provide research problems for human-level AI.

Finally, we review the areas of AI that are applicable to these problems. This information is collected together in figure 1.

Although we list specific genres, the genres are fuzzy concepts, with many games being hybrids, incorporating components of multiple genres. For example, there are strategy games (*Dungeon Keeper*) that allow the human to “jump in the body” of one of their units and play as if it is an action game for a while. Also, there are action games where you must also manage resources and multiple units (such as *Battlezone*). Although there will be a continual blurring of the genres, the basic roles for AI stay the same: enemies, partners, support characters, strategic opponents, units and commentators.

Action Games

Shortly after landing on an alien surface, you learn that hundreds of your men have been reduced to just a few. Now you must fight your way through heavily fortified military installations, lower the city's defenses, and shut down the enemy's war machine.

—*Quake II*

Action games involve the human player con-



Figure 2. A Screenshot from the Popular Action Game *Half-Life* from Valve.

(Reproduced with permission from Valve Software.)

trolling a character in a virtual environment, usually running around and using deadly force to save the world from the forces of evil. These games vary in the perspective that the human has of his/her character, be it *first person*, where the human sees what the character would see, or *third person*, where the player looks over the shoulder of the character. Popular examples include *Doom*, *Quake*, *Descent*, *Half-Life* (figure 2), *Unreal*, and *Tomb Raider*. In pure action games, AI is used to control the enemies, which are invariably alien monsters or mythical creatures. Realism in graphics has been the point of competition for these games; however, the graphics race seems to have run its course, with better AI becoming the point of comparison. Recent games, such as *Rainbow Six*, have extended the genre so that the human player can be part of a team, which includes either human or AI partners.

Role-Playing Games

Immerse yourself in a...world, where nations hang in the balance of your actions, dark prophecies test your resolve, and heroic dreams can be fulfilled at last.

—*Baldur's Gate*

In role-playing games, a human can play different types of character, such as a warrior, a magician, or a thief. The player goes on quests, collects and sells items, fights monsters, and expands the capabilities of the character (such as strength, magic, quickness), all in an extended virtual world. Example games include *Baldur's Gate*, *Diablo*, and *Ultima*. Recently, massively multiplayer role-playing games have been created where thousands of people play and interact in the same game world: *Ultima Online*, *Everquest*, and *Asheron's Call*. In both types of role-playing game, AI is used to control enemies, as with action games; partners who travel and adventure with the players; and supporting characters, such as shopkeepers. The



Figure 3. A Screenshot from the Upcoming Adventure Game *Rubu Tribe* by Outrage Entertainment.
(Reproduced with permission from Outrage Entertainment.)

massively multiplayer games provide an additional opportunity to use AI to expand and enhance the player-to-player social interactions, perhaps with AI-controlled kings who war by sending player-controlled knights to battle each other.

Adventure Games

Aye, 'tis a rollicking piratey adventure that's sure to challenge the mind and shiver a few timbers!

—*The Curse of Monkey Island*

Adventure games, and the related genre of interactive fiction, move further from action games as they deemphasize armed combat and emphasize story, plot, and puzzle solving. In these games, players must solve puzzles and interact with other characters as they progress through an unfolding adventure that is determined in part by their actions. Early adventure games, such as *Adventure* and *Zork*, were totally text based, but more recent games sport 3D graphics (sometimes using the graphics engines developed for action games). Example

games include the Infocom series, *King's Quest*, and many games from Lucas Arts, such as *Full Throttle*, *Monkey Island*, and *Grim Fandango* as well as *Rubu Tribe* (figure 3) from Outrage. AI can be used to create realistic supporting goal-driven characters that the player must interact with appropriately to further their progress in the game. One of the Holy Grails of interactive fiction is to have a computer director who can dynamically adjust the story and plot based on the actions of the human. The majority of games have fixed scripts and use many tricks to force the human player through essentially linear stories. However, a few games, such as *Blade Runner*, have incorporated some autonomy and dynamic scripting into their characters and story line (Castle 1998).

Strategy Games

Players must successfully construct and rule their medieval empire while they engage in real-time tactical warfare over land, sea, and air.

—*Warcraft*

In strategy games, the human controls many units (usually military units, such as tanks, or the ever-present alien war machines) to do battle from a god's eye view against one or more opponents. Strategy games include reenactments of different types of battle: historical (*Close Combat*, *Age of Empires*), alternative reality (*Command and Conquer*), fictional future (*Starcraft*), and mythical (*Warcraft*, *Myth*). The human is often faced with problems of allocating resources, scheduling production, and organizing defenses and attacks (Davis 1999). AI is used in two roles: (1) as a control for the detailed behavior of individual units that the human commands and (2) as a strategic opponent that must play against the human. The AI needs of the individual units differ from the enemies and partners of action and role-playing games because they are not meant to be autonomous but are meant to be good soldiers who "follow orders."

God Games

You're in charge of creating an entire city from the ground up—and the sky's the limit.

—*SimCity 3000*

God games give the player godlike control over a simulated world. The human can modify the environment and, to some extent, its inhabitants. The entertainment comes by observing the effects of his/her actions on individuals, society, and the world. *SimCity* is the classic example of a god game where the human acts as mayor, and the AI controls individual units or citizens of the simulated city. *The Sims* is probably the most intriguing example (figure 4). The player creates individual characters (units) that have significant autonomy, with their own drives, goals, and strategies for satisfying these goals, but God (the human player) can come in and stir things up by managing both the individual characters and their environment.

Team Sports

Welcome to *Madden NFL 97*, the game that captures the excitement of a 30-yard touchdown pass, the strategy of a well-executed scoring drive, and the atmosphere of a crisp autumn afternoon in the stadium.

—*Madden NFL 97*

Team sports games have the human play a combination of coach and player in popular sports, such as football (Whatley 1999), basketball, soccer, baseball, and hockey. AI is used in two roles that are similar to the roles in strategy games: The first is unit-level control of all the

individual players. Usually, the human controls one key player, such as the quarterback, while the computer controls all the other members of the team (figure 5). A second role is the strategic opponent, which, in this case, is the opposing coach. One unique aspect of team sport games is that they also have a role for a commentator, who gives the play-by-play and color commentary of the game (Frank 1999).

Individual Sports

Rip up the course on inline skates, speed on the street luge, pull serious air on the skateboard, and shred courses on the mountain bike.

—*ESPN Extreme Games*

For individual competitive sports, such as driving, flying, skiing, and snowboarding, the computer provides a simulation of the sport from a first- or third-person perspective. The human player controls a participant in the game who competes against other human or computer players. The computer player is more like an enemy in an action game than a strategic opponent or unit from a strategy game because the game is usually a tactical, real-time competition. Individual sports can also require commentators.

Roles

In each of the game genres described earlier, AI is used in a variety of different roles to populate the game environment.

Tactical Enemies

In early games, the tactics of the computer-controlled enemies were generally limited to running directly at the player. Later enemies were scripted or controlled by simple finite-state machines. In these early games, the enemies were made more challenging, not with improved intelligence but with bigger guns, tougher hides, and superior numbers. They also usually "cheated" by being able to see through walls or out the back of their heads. More recently, games such as *Half Life* (Birdwell 1999), *Descent 3*, *Quake III*,³ and *Unreal Tournament* have incorporated path planning and many tactics that make these enemies more humanlike. Our own research (Laird 2000b; Laird and van Lent 1999) has concentrated on building enemies for *Quake II* that have the same strengths and weaknesses as human players. To beat them, you have to outthink them as much as you have to outshoot them. Our SOAR QUAKEBOT is essentially a real-time expert system that has multiple goals and extensive tactics and knowledge of the game. It is built within the SOAR architecture and has over 800

Building human-level enemies for these games requires solving many general AI problems and integrating the solutions into coherent systems.



Figure 4. AI Characters, Called Sims, in the God Game The Sims.

(Reproduced with permission, courtesy Electronic Arts.)

rules. While it explores a level, it creates an internal model of its world and uses this model in its tactics to collect nearby weapons and health, track down an enemy, and set ambushes. It also tries to anticipate the actions of human players by putting itself in their shoes (creating an internal model of their situation garnered from its perception of the player) and projecting what it would do if it were the human player.

Building human-level enemies for these games requires solving many general AI problems and integrating the solutions into coherent systems. The enemies must be autonomous. They must interact with complex dynamic environments, which requires reactive behavior, integrated planning, and commonsense

reasoning. As they advance, they will also need models of high-level vision that have the same strengths and weaknesses as humans. One common complaint among game players is that the enemy AI is cheating, which destroys the game-playing experience. For example, if the human is in a dark room, the AI would be cheating if it could easily sense, identify, and locate the human. However, if the human is backlit by a bright hall, the AI enemy should be able to easily sense and locate the human but possibly not identify him/her. This element is important for game play so that the same tactics and behaviors that work well with humans work well with AI enemies.

There are many other applications of AI to building intelligent enemies. Because of the



Figure 5. A Team Composed of AI-Controlled Players in Madden NFL™ 2000 Football.
(Reproduced with permission from Electronic Arts Inc.)

extended geography of the environment, they must navigate and use path planning, spatial reasoning, and temporal reasoning. As the games become more complex, the enemies will need to plan, counterplan, and adapt to the strategies and tactics of their enemies, using plan recognition, opponent modeling techniques, and learning. Their responses need to be within the range of humans in terms of reaction times and realistic movement. One can even imagine adding basic models of emotions, where the enemies get “mad,” get “frustrated,” and change their behavior as a result.

Partners

Creating AI-controlled partners involves many of the same research issues as tactical enemies. However, enemy AI systems emphasize autonomy, and partners emphasize effortless cooperation and coordination between the human player and the AI partner. Current games restrict the human to using specific commands to interact with partners, such as defend, attack, and follow me—commands much more limited than used in human-to-human interactions. In the extreme, the need for full human-to-human interaction brings in speech recogni-

tion and natural language processing and even gesture recognition. The partner AI must coordinate its behavior, understand teamwork, model the goals of the human, and adapt to his/her style. Building such partners can build on previous research in AI in these areas but within the context of all the other cognitive activities involved in playing the game.

Support Characters

Support characters are usually some of the least sophisticated AI characters in computer games, but they have the most promise to improve games and are the most interesting in terms of developing human-level AI. They currently have sets of canned responses that they spit back to the user based either on menu-selected questions or keywords. The most complex ones, such as in *Blade Runner* (Castle 1998) have some autonomy and some simple goals, but they are extremely narrow goals with limited sets of behaviors for achieving these goals.

Adding other AI-controlled support characters could help populate the games with interesting opportunities for interaction that guide the player along various plot lines. Because these characters need to exist in a virtual world

and generally play a human role in this world, they provide a useful first step toward human-level AI. In this role, support characters must interact with, and adapt to, the environment; interact with, and adapt to, human players and other support characters; and provide humanlike responses, possibly including natural language understanding and generation. Because of these requirements and because these support characters are most directly playing the role of embodied virtual humans, they require a wide range of integrated AI capabilities, including everything from natural language to path planning to teamwork to realistic movement.

Strategic Opponents

When creating strategic opponents for strategy games and team sports games, most game developers have had to resort to "cheating" to make the opponent challenging. Often, strategic opponents are given extra units or resources or additional information about the map or the human player's position, or they play the game by a different set of rules. Even with these advantages, most strategic opponents are predictable and easily beaten once their weaknesses are found. Strategic opponents for team sports games face an additional difficulty in that their style of play must match a real-world team about which the human players are likely to be knowledgeable.

The tasks a strategic opponent must perform can be divided into two categories: (1) allocating resources and (2) issuing unit-control commands. Involved in both of these tasks is the development of a high-level strategy. Creating this strategy, which is where current strategic opponents are weakest, involves integrated planning, commonsense reasoning, spatial reasoning, and usually plan recognition and counterplanning to react to the human's attack. One of the most important aspects of strategy creation is the coordination of multiple types of a unit into a cohesive strategy. Once the plan is decided, the strategic opponent must determine how to best use limited resources (mined minerals or substitute players on a team) to compose an attack force appropriate to implement the battle plan. This re-

source allocation involves scheduling production and temporal reasoning about when the resulting units will be available. The strategic opponent must also issue commands to the newly created individual units, causing them to carry out the battle plan. Controlling a large force of units with only a single mouse is a significant part of the challenge for human players. Because of the human's limited input ability, the strategic opponent must enforce humanlike limitations, such as reaction times and realistic movements, when issuing commands to make the battle fair.

Units

In strategy games, god games, and team sports games, AI is used to control individual units. Generally, these units are given high-level commands from either the human player or the strategic opponent and need to carry out these commands. Units are usually controlled by finite-state machines (or large *c* functions) that are augmented with special routines for path planning (Cavazza 1999) and path following. In addition to following orders, units often need some ability to act autonomously. For example, a platoon of marines moving from one position to another should not ignore an enemy tank. Instead, they should autonomously choose to attack if appropriate or else find a new path. This semiautonomous behavior involves commonsense reasoning and perhaps coordination with other units. Because there can be hundreds of units active in a game at one time, the issues of computational and memory overhead are particularly important for unit AI (Atkin, Westbrook, and Cohen 1999).

Commentators

The role of the commentator is to observe the actions of the AI and the human and generate natural language comments suitable to describe the action (Frank 1999). In the RoboCup competition, there is a separate competition for commentator agents (Binsted 1998). Although sports games, both team and individual, are the most obvious genres for commentators, they can also be found in some action games, such as *Unreal Tournament*. The obvious challenge for a

commentator is to create a natural language description of the ongoing action in the game. The description can include both the moment-to-moment action as well as key tactical and strategy events that can require complex plan recognition and a deep understanding of the game.

Resource and Development Issues

A constant issue for game developers is the need to meet the limited computational power, in both memory and processing power, available in the average home computer or video game console. These resource issues can be finessed within the academic research community when the goal is just to do research on human-level AI independent of the commercial applications. However, we encourage researchers to take resource issues seriously because the more accessible our research is, the more likely it is that game developers and other industries will understand the need for research on human-level AI and AI techniques in general. Our experience with the SOAR QUAKEBOT has driven us to research on comparisons of SOAR with other architectures (Bhattacharyya and Laird 1999; Wallace and Laird 1999) and the overall efficiency of SOAR. The SOAR QUAKEBOT requires 3 megabytes of random-access memory and 10 percent of the processing power of a 400-megahertz WINDOWS NT PENTIUM II.

An additional constraint is that these AI systems must be developed at moderate cost. A game company will not be able to spend more than one person-year on development of the AI for a game. We need to develop techniques for quickly building and customizing human-level AI systems. Research on software engineering, knowledge acquisition, and machine learning will definitely play a role.

Conclusion

From a researcher's perspective, even if you are not interested in human-level AI, computer games offer interesting and challenging environments for many, more isolated, research problems in AI. We are most interested in human-level AI and wish to leverage computer games to rally support for research in human-level AI. One

attractive aspect of working in computer games is that there is no need to attempt a "Manhattan Project" approach with a monolithic project that attempts to create human-level intelligence all at once. Computer games provide an environment for continual, steady advancement and a series of increasingly difficult challenges. Just as computers have inexorably gotten faster, computer game environments are becoming more and more realistic worlds, requiring more and more complex behavior from their characters. Now is the time for AI researchers to jump in and ride the wave of computer games.

Acknowledgments

The authors are indebted to the many students and staff who have worked on the SOAR/Games project, most notably Steve Houchard, Karen Coulter, Mazin Assanie, Josh Buchman, Joe Hartford, Ben Houchard, Damion Neff, Kurt Steinkraus, Russ Tedrake, and Amy Unger.

Notes

1. See, for examples, www.stomped.com/published/jcal979197050_1_1.html.
2. John McCarthy, 1998, Partial Formalizations and the Lemmings Game. www.formal.stanford.edu/jmc/lemmings.html.
3. G. Keighley, 1999, The Final Hours of *Quake III Arena*: Behind Closed Doors at id Software, GameSpot. Available at www.gamespot.com/features/btg-q3/index.html.

References

- Agre, P. E., and Chapman, D. 1987. *Pengi*: An Implementation of a Theory of Activity. In Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87), 268-272. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Asada, M.; Veloso, M.; Tambe, M.; Noda, I.; Kitano, H.; and Kraetzschmar, G. K. 2000. Overview of RoboCup-98. *AI Magazine* 21(1): 9-19.
- Atkin, M. S.; Westbrook, D. L.; and Cohen, P. R. 1999. Capture the Flag: Military Simulation Meets Computer Games. In *Papers from the AAAI 1999 Spring Symposium on Artificial Intelligence and Computer Games*, 1-5. Technical Report SS-99-02. Menlo Park, Calif.: AAAI Press.
- Bates, J. 1992. Virtual Reality, Art, and Entertainment. *Presence: The Journal of Teleoperators and Virtual Environments* 1(1): 133-138.

- Bhattacharyya, S., and Laird, J. E. 1999. Lessons for Empirical AI in Plan Execution. Paper presented at the Sixteenth International Conference on Artificial Intelligence (IJCAI-99) Workshop on Empirical AI, 31 July-6 August, Stockholm, Sweden.
- Binsted, K. 1998. Character Design for Soccer Commentary. Paper presented at the RoboCup Workshop, 2-3, 9 July, Paris, France.
- Birdwell, K. 1999. The CABAL: Valve's Design Processing for Creating Half-Life. *Game Developer* 6(12): 40-50.
- Brooks, R. A.; Breazeal, C.; Marjanovic, M.; Scassellati, B.; and Williamson, M. 2000. The Cog Project: Building a Humanoid Robot. In *Computation for Metaphors, Analogy, and Agents*, ed. G. Nehaniv, 52-87. Lecture Notes on Artificial Intelligence 1562. Berlin: Springer-Verlag.
- Castle, L. 1998. The Making of *Blade Runner*, Soup to Nuts! Paper presented at the Computer Game Developers' Conference, 4-8 May, Long Beach, California.
- Cavazza, M.; Bandi, S.; and Palmer, I. 1999. "Situating AI" in Video Games: Integrating NLP, Path Planning, and 3D Animation. In *Papers from the AAAI 1999 Spring Symposium on Artificial Intelligence and Computer Games*, 6-12. Technical Report SS-99-02. Menlo Park, Calif.: AAAI Press.
- Croal, N., and Totilo, S. 1999. Who's Got Game? *Newsweek* 134(11): 46.
- Davis, I. 1999. Strategies for Strategy Game AI. In *Papers from the AAAI 1999 Spring Symposium on Artificial Intelligence and Computer Games*, 24-27. Technical Report SS-99-02. Menlo Park, Calif.: AAAI Press.
- Fasciano, M. J. 1996. Real-Time Case-Based Reasoning in a Complex World. Technical Report, TR-96-05, Computer Science Department, University of Chicago.
- Frank, I. 1999. Explanations Count. In *Papers from the AAAI 1999 Spring Symposium on Artificial Intelligence and Computer Games*, 77-80. Technical Report SS-99-02. Menlo Park, Calif.: AAAI Press.
- Hayes-Roth, B., and Doyle, P. 1998. Animate Characters. *Autonomous Agents and Multi-Agent Systems* 1(1): 195-230.
- Jones, R. M.; Laird, J. E.; Nielsen, P. E.; Coulter, K. J.; Kenny, P. G.; and Koss, F. V. 1999. Automated Intelligent Pilots for Combat Flight Simulation. *AI Magazine* 20(1): 27-42.
- Laird, J. E. 2000a. Bridging the Gap between Developers and Researchers. *Game Developer* 7(8): 34.
- Laird, J. E. 2000b. It Knows What You're Going To Do: Adding Anticipation to a QUAKEBOT. In *Papers from the AAAI 2000 Spring Symposium on Artificial Intelligence and Interactive Entertainment*, 41-50. Techni-

cal Report SS-00-02. Menlo Park, Calif.: AAAI Press.

Laird, J. E., and van Lent, M. 1999. Developing an Artificial Intelligence Engine. Paper presented at the Game Developers' Conference, 15-19 March, San Jose, California.

Stern, A. 1999. AI Beyond Computer Games. In *Papers from the AAAI 1999 Spring Symposium on Artificial Intelligence and Computer Games*, 77-80. Technical Report SS-99-02. Menlo Park, Calif.: AAAI Press.

Takahashi, D. 2000. Artificial Intelligence Gurus Win Tech-Game Jobs. *The Wall Street Journal*, March 30, 2000, B14.

Tambe, M.; Johnson, W. L.; Jones, R. M.; Koss, F.; Laird, J. E.; Rosenbloom, P. S.; and Schwamb, K. 1995. Intelligent Agents for Interactive Simulation Environments. *AI Magazine* 16(1): 15-39.

Wallace, S., and Laird, J. E. 1999. Toward a Methodology for AI Architecture Evaluation. In *Intelligent Agents VI*, eds. N. R. Jennings and Y. Lespérance, 117-132. Berlin: Springer-Verlag.

Whatley, D. 1999. Designing around Pitfalls of Game AI. Paper presented at the Game Developers' Conference, 15-19 March, San Jose, California.


Woodcock, S. 2000. Game AI: The State of the Industry. *Game Developer* 7(8): 24-32.



John Laird is a professor of electrical engineering and computer science at the University of Michigan. He received his Ph.D. in computer science from Carnegie Mellon University in 1983. He is a fellow of the American Association for Artificial Intelligence and a founder of Soar Technology, Inc. His e-mail address is laird@umich.edu.



Michael van Lent is a postdoctoral fellow at the University of Michigan. He received his Ph.D. from the University of Michigan in 2000. His research interests include machine learning and AI in computer games. His e-mail address is vanlent@umich.edu.



Artificial Intelligence and Mobile Robots

Case Studies of Successful Robot Systems

Edited by David Kortenkamp, R. Peter Bonasso, and Robin Murphy

The mobile robot systems described in this book were selected from among the best available implementations by leading universities and research laboratories. These are robots that have left the lab and been tested in natural and unknown environments. They perform many different tasks, from giving tours to collecting trash. Many have distinguished themselves (usually with first or second-place finishes at various indoor and outdoor mobile robot competitions).

Each case study is self-contained and includes detailed descriptions of important algorithms, including pseudo-code. Thus this volume serves as a recipe book for the design of successful mobile robot applications. Common themes include navigation and mapping, computer vision, and architecture.

6 x 9, 400 pp., ISBN 0-262-61137-6

**To order, call 800-356-0343 (US and Canada) or (617) 625-8569.
Distributed by The MIT Press, 55 Hayward, Cambridge, MA 02142**