

# The Computer Games Manifesto

Martin MAGNUSSON

*Department of Computer and Information Science  
Linköping University, 581 83 Linköping, Sweden  
E-mail: marma@ida.liu.se*

**Abstract.** We propose computer game AI as a suitable topic for research with the long term goal of artificial general intelligence in mind. Modern games include large virtual worlds and a growing but unsatisfied demand for human-like computer characters inhabiting them. Significant research challenges are present in building an agent capable of intelligent timely behaviour in such complex and dynamic environments. However, games can help turn these challenges into concrete problems forming stepping stones that enable incremental progress. Game AI might thus serve to revive the research field's original ambitious goal of human-level AI and help drive progress towards this difficult task.

**Keywords.** Artificial general intelligence, computer games, simulation

## 1. Introduction

Much artificial intelligence research is directed towards special-purpose technologies that provide great benefits in their application niches, and there are many good reasons to focus on narrow domains. Even very specialized problems attacked by artificial intelligence researchers, e.g. in the area of automated planning, are still exceedingly difficult. One need also consider the usefulness of the resulting technologies, a concern which is aided by a concentration of effort on specific real world tasks. Besides, if intelligence is an elusive term to start with, *general intelligence* seems fuzzier still.

But some researchers have expressed unease as to whether matters of generality have received insufficient attention ([2], [7]). And even though the concept of artificial general intelligence (AGI) may be hard to define, some applications, such as natural language understanding, are believed to be fundamentally conditional upon it.

We suggest that the use of artificial intelligence to enhance computer games is another application that requires generality and simultaneously enjoys unusually many beneficial properties as a research topic. The games we have in mind are not the classical board games that are characterized by determinism and complete knowledge. Nor are we thinking of present day commercial computer games with their exclusive reliance on ad hoc techniques to hide a complete lack of character intelligence behind a facade of appearances. Rather they are games that result from the application of modern and future AI technologies to inhabit the complex and dynamic simulated worlds of computer games with human-like autonomous intelligent agents.

Some of these games will only ever be experimental platforms, sharing many of the favourable characteristics of the simulation environments that are commonplace in

other AI research, while adding an element of motive, involvement, and fun. Others might result in technologies that help migrate some genuine intelligence into commercial games, opening up a new space of game design possibilities and creating an entirely new genre of games that introduces AI to the general public.

Below, we discuss some particularly useful properties of games used as test beds for research in AGI and take a brief look at previous work on the topic before concluding with a look towards the future.

## 2. Accessibility

Many games consist of simulated agents in simulated environments. Simulations are safe, cheap, do not break down or require spare parts, need never be charged, and their assembly requires no competence outside a computer scientist's. Experiments with them can be run with controlled input and repeated thousands of times at speeds unlimited by external reality. Their only *real* problem is that they are *not* real. When replacing work on physical robots by simulation one risks sweeping important problems under blanket assumptions, which will later reemerge to frustrate any attempt to utilize the developed solutions in practice. Fortunately, for those working with computer games this is a non-issue. The simulated game environment *is* the real environment. The application *is* a virtual world, and solutions benefit the very same sort of simulation in which they were developed.

These factors contribute to the accessibility of game AI research for the researcher himself. But the research results are also made accessible to fellow researchers and potential users. A simulated environment is pure software that can be readily made available on-line. Identical copies can be evaluated by interested peers and any research claims independently verified. Vague claims of capabilities can be put to the test and unrealistic assumptions uncovered in a spirit of complementing theory by allowing the results to speak for themselves.

## 3. Challenges

Although most simulations simplify the real world, game environments are still surprisingly complex. Much more so than classical benchmark problems such as the blocks world and the three wise men puzzle, often used in AI research. Such benchmark problems help clarify specific difficulties that pose tough challenges for any generally intelligent system. However, criticism is often raised over their limited scope (a classical example being Hayes' [4]), and doubts have been voiced over the scalability of proposed solutions. In contrast, many computer games involve intricate dynamics where restrictive simplifying assumptions do not hold and call for the integration of many different capabilities necessary for general intelligence, some of which we list below.

**Planning** helps the game and its characters adapt to new situations created in the interaction with human players, even if unanticipated by the game designers. Goal directed agents also have a "will" of their own and are able to perform sequences of actions to achieve their goals. These are clear indications of intelligence that increase the characters believability to humans who observe them.

**Natural language understanding** could enable an entirely new level of dialog interaction hitherto unseen in games, which usually stick to “canned” phrases. Players would elicit important information through the language understanding and question answering capabilities of computer characters in order to solve game puzzle elements.

**Epistemic reasoning** will be required by game inhabitants, who should not hope to have complete knowledge of a dynamic environment that is constantly influenced by other characters and human players. Instead, they need to reason about how to obtain the knowledge they require to carry out their business, e.g. by scouting an area or using a speech act. Any kind of “cheating” through direct lookup in the underlying simulation data structures would of course defeat the purpose of using games as a research challenge, but would also be evident to human players observing the resulting “clairvoyance”.

**Agent architectures** are required as the agents inhabiting game environments are considered complete autonomous individuals who need some way of incorporating all aspects from perception to action. Robustness is vitally important in a context with unpredictable human players aiding or interfering with the agent’s plans.

**Self-awareness** in limited forms would seem to have a uniquely important status in games. Often, the entire purpose of the game characters is to act as if they were “real” live inhabitants of an alternative reality. They are, in effect, stand-ins for human actors. Any form of self-awareness would help sustain the illusion.<sup>1</sup>

#### 4. Applications

There is, in fact, a *need* for AI in commercial computer games. Graphical realism has already seen tremendous development and has long been a vital selling point in the marketing of games. As game graphics technology advances, the failure of game AI technology to exhibit similar progress makes such development a high priority that could be used to differentiate oneself from the competition. Perhaps it will provide the much needed innovation called for by John Riccitiello, chief executive of the the largest producer of commercial computer and video games Electronic Arts, who said that current games are “boring people to death” [13].

Though, academic research on games has often replaced human involvement by trying to construct an intelligent agent to solve the game itself. A well-known example is the simple Wumpus world in the leading AI textbook “Artificial Intelligence: A Modern Approach” [11] by Stuart Russel and Peter Norvig, who credit Michael Genesereth for suggesting the use of the Wumpus game.

A game environment whose solution would be much more challenging is involved in John McCarthy’s proposal that the computer game Lemmings provides research opportunities in common sense reasoning and that it “[...] allows its phenomena to be studied with a minimum of tedious technicalities and which provide a maximum of scientific information for the work expended” [8].

The suggestion by Eyal Amir at the University of Illinois at Urbana-Champaign and Patrick Doyle at Stanford University is even more challenging. They propose the use

---

<sup>1</sup>Of course, should progress proceed beyond illusion we ought no longer keep the AI imprisoned in games.

of text adventure games as a vehicle of research in cognitive robotics where the agent would start out with severely limited knowledge, not knowing what actions are available to it, what fluents it should use to represent the environment, nor even the purpose of the game apart from the commonsensical knowledge that games have a purpose [1]. These are some significant challenges, though they say a computer game “[...] allows us to examine them in a controlled environment in which we can easily change the problems to be solvable, and then gradually increase the difficulty step by step”.

An alternative to AI taking the role of the human players is its use to animate virtual characters in interactions with human players. Such work has seen a relatively recent surge of interest, and one advocate of this view is Alexander Nareyek, currently at the interactive intelligences Labs at the National University of Singapore. He notes that “[game] characters can be seen as *agents*, their properties perfectly fitting the AI agent concept” [9]. Nareyek’s work with the Excalibur system uses the settings of game environments to investigate the application of automated planning in highly dynamic worlds with incomplete knowledge and the need to reason about resources.

Another example is the Soar/Games project at the University of Michigan. They report uncovering new research challenges after coupling the Soar artificial general intelligence architecture to Quake 2 and Descent 3 [5]. Their emphasis is on generality in their attempts to build reusable rule bases for agent behaviour. John Laird’s and Michael van Lent’s enthusiasm for the use of computer games in AGI research is evident in their paper “Human-level AI’s Killer Application: Interactive Computer Games” [6].

Few true AI technologies have so far been applied to commercial games, with the prototypical exception being path finding using A\* search. Instead, most game developers try to emulate intelligent behaviour through relatively simple mechanisms such as ad hoc implementations of finite state machines. Such deterministic technologies often require the developer to foresee and explicitly encode a response for every conceivable possibility, enforcing a relatively rigid and limited game environment in order for this task not to grow impossibly complex.

However, there are exceptions. The innovative game Creatures [3] incorporated genetic algorithms and artificial neural networks in an artificial life simulation and Black & White [12] included reinforcement learning in autonomous agents that could be trained by the player. Another example is described by Jeff Orkin, currently at the MIT Media Lab but previously at Monolith Productions. He applied a planning algorithm reminiscent of STRIPS to enable enemy soldiers in Monolith’s game F.E.A.R. to plan sequences of actions for taking cover and attacking the player [10]. Severe real-time requirements limited the planning model, but it is nevertheless a demonstration of feasibility. A look at the positive reception of all of these games provides encouragement for further such efforts.

## 5. Future Directions

When working with games interesting problems abound, and most of them call for new research in artificial general intelligence. Some open questions are the reasoning about in-game resource use, planning to satisfy the long term goals of the game, representation of and reasoning with knowledge of the virtual world, what to do about conflicting and inconsistent information, how to learn from experience or from being taught by a player

and the integration of perception, reasoning, planning, execution, and failure recovery in an agent architecture. Moreover, all of these topics need to be treated in the context of real time operation of persistent agents that accumulate experiences over long time periods. Such difficult problems at least tend to have concrete instantiations in the structure of game environments that make them easier to think about, discuss, and hopefully to solve.

Finally, computer games span the range between the utterly simple and the dauntingly complex, from discrete puzzles like Tetris to highly dynamic social simulations like The Sims. Oddly, players find the entire continuum satisfying. Game AI research can take advantage of this fact by identifying or constructing games of a suitable complexity for the AI techniques under investigation, creating a significant challenge that is still clearly solvable. No threshold capability must be surpassed before applications become possible. One might start right away, attacking problems that are relatively simple while moving steadily towards bigger tasks. As research progresses, over years and decades, human players will find it increasingly hard to separate artificial characters from other human players and one might even come to approximate a sort of Turing test.

Thus, computer game AI could provide fertile ground for a seed of renewed interest in the research fields original vision. Perhaps it might even pique gamers' and students' interest in AI research, and we will need good help to succeed in our pursuit of the extraordinarily ambitious goal of building artificial general intelligence.

## References

- [1] Eyal Amir and Patrick Doyle. Adventure games: a challenge for cognitive robotics. In *AAAI'02 workshop on Cognitive Robotics*, 2002.
- [2] Mark Baard. AI founder blasts modern research. <http://www.wired.com/science/discoveries/news/2003/05/58714>, May 2003.
- [3] Stephen Grand, Dave Cliff, and Anil Malhotra. Creatures: Artificial life autonomous software agents for home entertainment. In *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, pages 22–29, 1997.
- [4] Patrick J. Hayes. The naive physics manifesto. In Donald Michie, editor, *Expert Systems in the Micro-Electronic Age*, pages 242–270. Edinburgh University Press, 1979.
- [5] John E. Laird and Michael van Lent. Developing an artificial intelligence engine. In *Proceedings of the Game Developers' Conference*, pages 577–588, 1999.
- [6] John E. Laird and Michael van Lent. Human-level AI's killer application: Interactive computer games. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, pages 1171–1178, 2000.
- [7] John McCarthy. AI needs more emphasis on basic research, 1983. <http://www-formal.stanford.edu/jmc/basicresearch/basicresearch.html>.
- [8] John McCarthy. Partial formalizations and the Lemmings game. Technical report, Stanford University, Formal Reasoning Group, 1995. <http://www-formal.stanford.edu/jmc/lemmings.html>.
- [9] Alexander Nareyek. Intelligent agents for computer games. In T. A. Marsland and I. Frank, editors, *Computers and Games, Second International Conference CG 2000*, pages 414–422, 2000.
- [10] Jeff Orkin. Agent architecture considerations for real-time planning in games. In *Proceedings of the First Artificial Intelligence and Interactive Digital Entertainment Conference AIIDE 2005*, pages 105–110, 2005.
- [11] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, Englewood Cliffs, NJ, 2nd edition, 2003.
- [12] Lionhead Studios. Black & white. <http://www.lionhead.com/bw/>, 2007.
- [13] Nick Wingfield. EA chief cites need for more innovative games. *The Wall Street Journal*, page B4, July 9 2007.