

Computational Intelligence and AI in Games: A New IEEE Transactions

Simon M. Lucas

Games have long been seen as an ideal test-bed for the study of AI. Until recently, most of the academic work in the area focused on traditional board games and card games, the challenge being to beat expert human players. Following the release of Pong in the early 1970s, the last few decades have seen a phenomenal increase in the quality, diversity and pervasiveness of video games. The value of the worldwide computer and video games market is estimated to be \$USD25bn annually¹, and is predicted to grow rapidly over the next decade. This presents academic researchers [1] and game developers with the challenge of developing next generation game AI. The future will see games with intelligent empathetic characters who get to know you, and work hard to optimise your experience while playing the game. New titles such as *Left 4 Dead*² have already made important steps in this direction. Superior game AI will lead to more entertaining and immersive games and also add value to the burgeoning serious games market.

Traditional games are constrained by the ability of the players to manipulate the game objects such as pieces on a chess-board or cards from a deck. The rules of a game specify how the game objects interact and fundamental combinatorics quickly produce enormous game trees; this ensures that complex decision-making processes are required in order to play such games to a high standard. Implementing computer programs that are able to defeat their human opponents has been the subject of a great deal of AI research, epitomised by outstandingly successful programs such as *Deep Blue* for Chess [2] and Chinook for Checkers [3],[4]. These programs expend most of their effort on game-tree search, and for this reason are often called *brute force* approaches. This term although widely used [5] is somewhat inaccurate: there is a great deal of subtlety in the way these algorithms are implemented and fine-tuned. The levels of play achieved demonstrate extraordinary prowess in computer science and engineering for the authors of those systems. On a broader note, they force us to reconsider the nature of intelligence.

Game-tree search has limited applicability when the game has a large branching factor or the state-space is continuous. Hence the techniques applied to chess and checkers do not work well for Go, and are even less applicable to video games where the state-space is for practical purposes continuous. Traditional games usually have well-defined fixed rules, known in advance to the players.

Contemporary video games do not fit this paradigm. When playing a video game such as *Crysis* or *Grand Theft Auto* the player engages in a series of missions, learning about the game as he plays. While some aspects of the game-play remain constant for a given game and even a particular genre of game, there are other aspects of the game that you learn while playing. For example, the principle that it is good to kill enemies, but bad to be killed yourself may be fixed, but the details of how this is achieved may need to be learned during game play; this includes learning about the geography of the game area as well as the capabilities of the weapons and other power-ups, and the behaviour of the enemies.

For researchers making the transition from traditional game AI to video-game AI, it is a whole new world, but one full of exciting possibilities. Fortunately, in recent years a number of Application Programming Interfaces (APIs) have been published enabling researchers to conduct experiments in these complex virtual environments without having to implement the game. This is good news for research in all the more challenging aspects of machine intelligence. Classic machine learning textbooks such as Reinforcement Learning [6] include algorithms that will directly solve a range of toy problems such as *Mountain Car* and *Inverted Pendulum*. Such algorithms can be applied with some effort to learn high performing game strategies for board games such as Othello. However, to apply these algorithms to achieve similarly high performance on video games such as 3D car racing, first-person shooter battles in games such as Unreal Tournament, war game strategy such as DefCon, or even to software agent Ms Pac-Man requires an immense effort and is the subject of much ongoing research. Interested readers can participate in these and other competitions at future IEEE Computational Intelligence and Games conferences.

AI can be applied to most aspects of game development and design including automated content creation, procedural animation, adaptive lighting, intelligent camera control, and perhaps the most obvious example, controlling the Non-Player Characters (NPCs). A natural instinct for researchers is to focus on making the NPCs more intelligent. This is not the prime interest of commercial game AI, however, which predominantly seeks to make games more fun. This begs the question: *What is the aim of the NPCs?* If it were simply to defeat the human player by killing him as quickly as possible, then that would be easy to measure, and depending on the game setup and

¹Source: PricewaterhouseCoopers

²<http://valve.com>

the agent's view of the world, may be easy to achieve.³ If this led to a game where the player dies quickly every time then it is unlikely to be a best seller. A more realistic and commercially viable aim for the NPCs is to offer a challenge that is fun and engaging. Naturally, this is harder to define. Attempting to measure the level of fun provided by a game is an important emerging discipline. Until we have widely accepted metrics for fun (which may never happen) we are left with subjective judgments about the value of some of this research. Inevitably some of the papers in this journal will have more qualitative evaluations than is the norm for other IEEE Transactions. Authors can help their cause here by providing additional content to accompany their manuscripts, such as videos, or even better, playable versions of the games.

Contemporary video games offer increasingly sophisticated virtual worlds which often have a great deal of artistic merit, and sometimes breathtaking beauty. As virtual worlds become more complex, however, it becomes harder to program game AI that can match this level of sophistication by delivering game characters with sufficiently interesting and credible behaviour. In the industry this is known as the *uncanny valley*. This refers to the fact that people are willing to ascribe an emotional state to simply animated cartoon characters like Bugs Bunny, but as the graphical rendering becomes more realistic the player becomes less engaged with the characters unless they can backup their good looks with appropriate behaviour.

A scene from the hit video game F.E.A.R. [7] is eloquently described by Schaeffer *et al* [8], and those of you who have played the game might relate to that. However, the dramatic tension for me falls well short of that created by a similar flanking manoeuvre by the velociraptors on Robert Muldoon (played by actor Bob Peck) in Jurassic Park. Even the best games have some way to go to reach the levels of immersion achieved by film.

As game environments become more physically realistic, so the control of game characters may draw more heavily on robotics research. We can look forward to the day when shooting an enemy in the leg makes the character limp, not because it has been programmed to, but because it is the only way it can walk given the now limited muscle contractions. This opens up a plethora of possibilities for open-ended immersive gameplay. Viewing the videos of *Big Dog*⁴ it can be hard to believe that one is watching a robot, and not the legs of two actors in some strange incarnation of a headless pantomime horse. Already the middleware to enable physically believable animation in games is starting to appear, with companies such as Natural Motion at the forefront of this⁵.

The range of existing games is incredibly diverse and the scope for improved AI is enormous. Better AI will lead to new genres of game, and improved performance on a range of current challenges. In the past fields which share much

in common have been too fragmented, and this journal has an important role to play in helping to bring various sub-fields together. There is much to be gained from a greater cross-fertilisation of ideas. For example, in last few years Monte Carlo methods have made truly astonishing progress in the world of Computer Go [9] [10] [11], and more recently in General Game Playing [12], but they have a much wider application than this. Given more research they should be applicable to a wide variety of problems in video game intelligence.

One of the most important aspects of machine intelligence that is embarrassingly absent (embarrassing for the machine learning community that is) is *learning*. There seems to be so much scope to apply learning methods in games, and yet successful examples of this in video games are few and far between. Until recently game strategy learning was split into two almost disjoint fields: reinforcement learning (in particular temporal difference learning), and evolutionary learning. There is still much to be understood about the relative merits of these techniques on a range of challenging problems. Learning of game strategies and more generally control strategies has a long history of research in both camps with early work on TDL [13] [14] and on evolutionary learning [15] [16]. A famously successful application of TDL was Tesauro's TD Gammon [17], which was followed up by an evolutionary approach to the same problem by Pollack and Blair [18]. Following that, evolutionary methods have also been successfully applied to checkers [19] and chess [20]. Recent work has attempted to gain greater understanding of the relative merits of evolution versus TDL [21] and to investigate how these methods may be combined to produce more efficient learning algorithms [22]. Nonetheless, there is still a gulf to be crossed in making learning work well within commercial video games, where it needs to be rapid, robust, and highly flexible — essentially more human-like. Interesting prototype solutions have been proposed and we can look forward to more advances in the next few years.

While most games are developed to be fun to play, there exists a class of games developed as simplified models for the study of economic or social behaviour. Game theory [23] attempts to find analytical solutions for these models using concepts such as rational or bounded-rational agents and Nash equilibria [24]. However, there are limits to the complexity of the game models and of the agents that can be analysed using this approach, and over the last few decades there has been great interest in understanding the behaviour of these models by simulating their evolution through time. Until now much of the work in this field has focused on relatively simple games such as the iterated prisoner's dilemma [25], [26]. However, the computing power is now readily available to deal with rich and detailed scenarios inhabited by potentially sophisticated agents where even more complex and interesting behaviour may emerge.

The journal publishes archival quality original papers in all aspects of computational intelligence and AI related to all types of games. To name some examples, these

³It is much harder to achieve with an *embodied* agent: one which sees the same kind of first-person view that a human player sees.

⁴See <http://www.bostondynamics.com>

⁵<http://naturalmotion.com>

include computer and video games, board games, card games, mathematical games, games that model economies or societies, serious games with educational and training applications, and games involving physical objects such as robot football and robotic car racing. Emphasis will also be placed on the use of these methods to improve performance in, and understanding of the dynamics of games, as well as gaining insight into the properties of the methods as applied to games. It will also include using games as a platform for building intelligent embedded agents for real world applications.

The journal builds on a scientific community that has already been active in recent years with the development of new conference series such as the IEEE Symposium on Computational Intelligence in Games (CIG), and Artificial Intelligence and Interactive Digital Entertainment (AIIDE), as well as special issues on games in journals such as the IEEE Transactions on Evolutionary Computation [27]. When setting up the journal a decision was made to include both Artificial Intelligence (AI) and Computational Intelligence (CI) in the title. AI seeks to simulate intelligent behaviour in any way that can be programmed effectively. Some see the field of AI as being all-inclusive, while others argue that there is nothing *artificial* about real intelligence as exhibited by higher mammals [28]. CI more specifically deals with algorithms and architectures that enable intelligent behaviour to emerge via statistical processes. CI methods are usually driven by data or experience. In fields such as natural language processing and machine translation, the statistical methods have significantly outperformed their more structured grammar driven alternatives. Whether we will see such shifts in commercial game AI in the near future remains to be seen, but provides fertile grounds for research.

I would like to thank David Fogel, Jim Keller, Xin Yao, and Vincenzo Piuri for all their encouragement and support in getting this new journal started in such a timely manner. Jo-Ellen Snyder, Sonal Parikh, and Bill Colacchio at IEEE CIS headquarters have been a great help dealing with all the necessary administrative and production tasks. Working with me at the University of Essex, Sharmila Woollam has been an invaluable editorial assistant. At the time of writing, the journal has achieved an average time to first decision of approximately six weeks. I thank the associate editors and reviewers for their hard work and dedication in achieving this. Finally, I encourage all authors working in this exciting, innovative, and economically important field to submit their high quality work to this journal and play their part in making it an outstanding success.

REFERENCES

- [1] J. E. Laird, "Research in human-level AI using computer games," *Communications of the ACM*, vol. 3, no. 8, pp. 32–35, 2002.
- [2] F. Hsu, *Behind Deep Blue*. Princeton University Press, 2002.
- [3] J. Schaeffer, N. Burch, A. K. Yngvi Björnsson, M. Mueller, R. Lake, P. Lu, and S. Sutphen, "Checkers is solved," *Science*, vol. 317, pp. 1518 – 1522, September 2007.

- [4] J. Schaeffer, *One Jump Ahead: Computer Perfection at Checkers*. Springer-Verlag, 2009.
- [5] S. L. Epstein, "Game playing: The next moves," in *AAAI/IAAI*, 1999, pp. 987–993.
- [6] R. Sutton and A. Barto, *Introduction to Reinforcement Learning*. MIT Press, 1998.
- [7] J. Orkin, "Three states and a plan," in *Game Developer Conference*, 2006.
- [8] J. Schaeffer, V. Bulitko, and M. Buro, "Bots get smart," *IEEE Spectrum*, pp. 44 – 49, 12/01 2008. [Online]. Available: <http://www.spectrum.ieee.org/dec08/7011>
- [9] B. Brüggemann, "Monte Carlo Go," *Computer Go Magazine*, 1993.
- [10] S. Gelly, Y. Wang, R. Munos, and O. Teytaud, "Modification of uct with patterns in Monte-Carlo Go," INRIA, Tech. Rep. Technical Report 6062, 2006.
- [11] Y. Wang and S. Gelly, "Modifications of uct and sequence-like simulations for Monte-Carlo Go," in *IEEE Symposium on Computational Intelligence and Games*, 2007, pp. 175 – 182.
- [12] M. R. Genesereth, N. Love, and B. Pell, "General game playing: Overview of the aaii competition," *AI Magazine*, no. 2, p. 62–72, 2005.
- [13] A. Samuel, "Some studies in machine learning using the game of checkers," *IBM Journal of Research and Development*, vol. 3, pp. 211 – 229, 1959.
- [14] D. Michie, "Trial and error," in *In Science Survey, part 2*. Penguin, 1961, pp. 129–145.
- [15] L. Fogel, A. Owens, and M. Walsh, "Artificial intelligence through a simulation of evolution," in *Biophysics and Cybernetic Systems: Proceedings of the 2nd Cybernetic Sciences Symposium*, M. Maxfield, A. Callahan, and L. Fogel, Eds. Washington DC: Spartan Books, (1965), pp. 131 – 155.
- [16] G. Burgin, "On playing two-person zero-sum games against non-minimax players," *IEEE Transactions on Systems Science and Cybernetics*, vol. 5, pp. 369–370, 1969.
- [17] G. Tesauro, "Temporal difference learning and TD-gammon," *Communications of the ACM*, vol. 38, no. 3, pp. 58–68, 1995.
- [18] J. Pollack and A. Blair, "Co-evolution in the successful learning of backgammon strategy," *Machine Learning*, vol. 32, pp. 225–240, 1998.
- [19] K. Chellapilla and D. Fogel, "Evolving neural networks to play checkers without expert knowledge," *IEEE Transactions on Neural Networks*, vol. 10, no. 6, pp. 1382–1391, 1999.
- [20] D. Fogel, T. Hays, S. Hahn, and J. Quon, "An evolutionary self-learning chess program," *Proceedings of the IEEE*, pp. 1947 – 1954, December 2004.
- [21] T. P. Runarsson and S. M. Lucas, "Co-evolution versus self-play temporal difference learning for acquiring position evaluation in small-board go," *IEEE Transactions on Evolutionary Computation*, vol. 9, pp. 628 – 640, 2005.
- [22] S. Whiteson and P. Stone, "Evolutionary function approximation for reinforcement learning," *Journal of Machine Learning Research*, vol. 7, pp. 877–917, 2006.
- [23] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- [24] J. Nash, "Equilibrium points in n-person games," *Proceedings of the National Academy of Sciences*, vol. 36, pp. 48 – 49, 1950.
- [25] R. Axelrod, *The Evolution of Cooperation*. Basic Books Inc., 1984.
- [26] G. Kendall, X. Yao, and S. Y. Chong, *The Iterated Prisoners' Dilemma: 20 Years On*. World Scientific, 2007.
- [27] D. B. Fogel, A. D. Blair, and R. Miiikkulainen, Eds., *IEEE Transactions on Evolutionary Computation Special Issue on Evolutionary Computation and Games*, vol. 9, no. 6, 2005.
- [28] J. Hawkins, *On Intelligence*. Holt, 2005.

BIOGRAPHY

Simon M. Lucas (SMIEEE) is a reader in computer science at the University of Essex (UK). His main research interests are evolutionary computation, games, and pattern recognition, and he has published widely in these fields with over 130 peer-reviewed papers. He was chair of IAPR Technical Committee 5 on Benchmarking and

Software (2002 - 2006) and is the inventor of the scanning n-tuple classifier, a fast and accurate OCR method. He was appointed inaugural chair of the IEEE CIS Games Technical Committee in July 2006, has chaired or co-chaired many international conferences, including the first IEEE Symposium on Computational Intelligence and Games in 2005. He is an associated editor of IEEE Transactions on Evolutionary Computation, and the Springer Journal of Memetic Computing. He was an invited keynote speaker or tutorial speaker at IEEE CEC 2007, IEEE WCCI 2008, IEEE CIG 2008, PPSN 2008, and IEEE CEC 2009. He leads the newly established Game Intelligence Group at the University of Essex.