AI4media

ARTIFICIAL INTELLIGENCE FOR THE MEDIA AND SOCIETY

ROADMAP ON AI TECHNOLOGIES & APPLICATIONS FOR THE MEDIA INDUSTRY

SECTION: "AI FOR GAMES"

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951911

info@ai4media.eu www.ai4media.eu

Authors	David Melhart (modl.ai) Ahmed Khalifa (modl.ai) Daniele Gravina (modl.ai)

This report is part of the deliverable D2.3 - "*AI technologies and applications in media: State of Play, Foresight, and Research Directions*" of the AI4Media project.

You can site this report as follows:

F. Tsalakanidou et al., Deliverable 2.3 - AI technologies and applications in media: State of play, foresight, and research directions, AI4Media Project (Grant Agreement No 951911), 4 March 2022

This report was supported by European Union´s Horizon 2020 research and innovation programme under grant number 951911 - AI4Media (A European Excellence Centre for Media, Society and Democracy).

AI for Games

Current status

During the last decade ***game AI*** as a research area has been becoming more and more popular in both academic and industrial circles. This is evidenced by increasing interest in AI by industry stakeholders and the large volume of academic studies, which has been supported by an active and healthy research community — roots of which point well beyond the past decade to at least since the start of the IEEE CIG and the AIIDE conference series in 2005. While initially most of the work published at academic venues was concerned with learning to play a particular game or using search/planning algorithms to play a game without learning, gradually, a number of new applications for AI in games and for games have come to complement the original focus on AI for playing games[1,2].

Since the early days of the field, papers on procedural content generation, player modelling, game data mining, human-like playing behaviour, automatic game testing and so on have become commonplace within the community. While academic interest in game AI applications has been on the rise in the last decade, industry and academia do not necessarily attempt to solve the same problems with the same approaches. Nevertheless, it may be that more traditional algorithmic solutions emerging from industry can inspire new approaches in academia and vice versa. Consequently, what we see today is a healthy indication of a parallel progress with a certain degree of collaboration.

Figure 1: The increase in the game development cost over the past four decades[3].

While it can be argued that ***Non-Player Characters*** (NPC) have been solved to a satisfactory degree with sophisticated behaviour trees, game AI research can offer much more to the industry. The multidisciplinary nature of game AI and a more pragmatic and holistic view of the

[1] G. N. Yannakakis. Game AI revisited. In Proceedings of the 9th conference on Computing Frontiers, pages 285–292. ACM, 2012

[2] G. N. Yannakakis, and J. Togelius. *Artificial intelligence and games*. Springer, 2018.

[3] Image source: R. Koster, VentureBeat, The cost of games (2018) https://venturebeat.com/2018/01/23/the-cost-of-games/

game AI problem have shifted academic and industrial interests in recent years. It seems that we have long reached an era where the primary focus of the application of AI in the domain of games is not on agents and NPC behaviours. The focus has, instead, started to shift towards interweaving game design and game technology by viewing the role of AI holistically and integrating aspects of ***procedural content generation*** and ***player modelling*** within the very notion of game AI[78]. While AI can help to make better games, this does not necessarily happen through better, more human-like or believable NPCs. Notable examples of non-NPC AI in games include *No Man's Sky* (Hello Games, 2016) and its procedural generation of a quintillion different planets, *Nevermind* (Flying Mollusk, 2016) with its affective-based game adaptation via a multitude of physiological sensors, and – more contemporarily – *Watch Dogs Legion* (Ubisoft, 2020) with procedurally generated characters and missions. But there might be other AI roles with game design and game development that are still to be found by AI.

Beyond playing games and content generation, AI might be able to play the role of a ***design assistant***, a ***data analyst***, a ***playtester***, a ***game critic***, or even a ***game director***. Large industry players like Ubisoft and King are working on integrating AI into their data analytics[4], content recommendation[5] and moderation[6], player modelling[7], and level design[8] pipelines in games like *Tom Clancy's The Division 1-2* (Ubisoft, 2016; 2019), *For Honor* (Ubisoft, 2017; 2018), *Candy Crush Saga* (King, 2012) and *Candy Crush Soda Saga* (King, 2014). Indeed the AI assistance in data processing, content generation, testing, and moderation is becoming more and more sought-after by industry stakeholders as new "game as a service" business models turn game development and maintenance into an ongoing and costly process. Companies like modl.ai are aiming to meet the demand with sophisticated AI tools that aid crucial parts of these processes.

Finally, beyond the games industry, video games are still used as benchmarks for more complex AI architectures. In recent years, we saw tech giants like OpenAI and the Google-affiliated DeepMind test their algorithms on esports games. As great milestones in both game AI and AI in general, OpenAI's *OpenAI Five* became the first AI system to defeat the world champions in *Dota 2* (Valve Corporation, 2013)[9] and DeepMind's *AlphaStar* was able to perform at a grand-master level in *StarCraft II* (Blizzard Entertainment, 2010)[10]. While advancements can be considered great stepping stones towards more complex general AI, they show the strength of games as benchmarks for AI applications beyond the games industry.

---

[4] A, Canossa, et al. "Like a DNA string: Sequence-based player profiling in Tom Clancy's the Division." Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment. Vol. 14. No. 1. 2018.
[5] L. Cao, et al. "Debiasing Few-Shot Recommendation in Mobile Games." ORSUM@ RecSys. 2020.
[6] A. Canossa, et al. "For Honor, for Toxicity: Detecting Toxic Behavior through Gameplay." Proceedings of the ACM on Human-Computer Interaction (CHI PLAY), 2021.
[7] D. Melhart, et al. "Your gameplay says it all: modelling motivation in Tom Clancy's The Division." 2019 IEEE Conference on Games (CoG). IEEE, 2019.
[8] V. Volz, et al. "Capturing local and global patterns in procedural content generation via machine learning." 2020 IEEE Conference on Games (CoG). IEEE, 2020.
[9] C. Berner, et al. "Dota 2 with large scale deep reinforcement learning." arXiv preprint arXiv:1912.06680, 2019.
[10] O. Vinyals, et al. "Grandmaster level in StarCraft II using multi-agent reinforcement learning." Nature 575.7782, 2019.

Research challenges

There are a number of challenges still open in different fields of game research. The intersection of key areas such as game playing, content generation, and player modelling hide numerous pitfalls when we are considering combining these methods to achieve more complex goals (e.g. combining game playing and player modelling in AI-assisted automated testing). While NPC AI might be a solved issue for gameplay purposes as far as the industry is concerned, game-playing agents that test games or replace human players still hold value to stakeholders. While research often focuses on limited testbeds and single-use applications, a great challenge lies in scaling these pipelines to production environments.

Game playing is one of the core research areas of game AI. Even though today's landscape is more diverse in terms of research topics, game playing still takes up the majority of published papers in the field[11]. While early studies and even recent high-profile examples (see DeepMind and OpenAI above) are focusing on playing to win (Figure 2), introducing other goals into the system is not straightforward. The main areas which can intersect game playing are playing for experience, player modelling, and Procedural Content Generation (PCG).

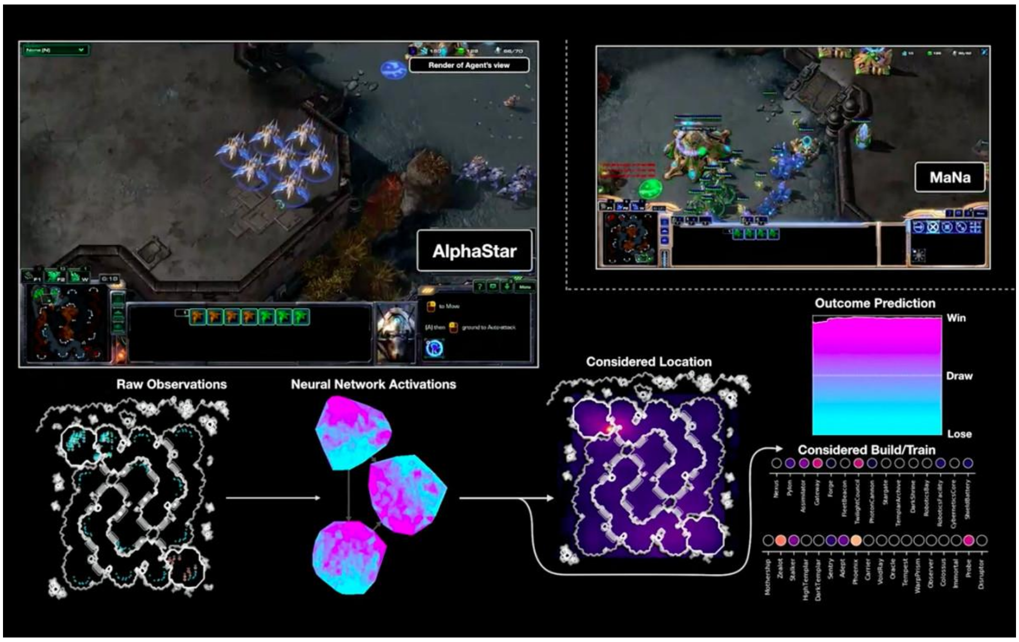

Figure 2: Visual representation of the AlphaStar agent while it plays against human champion MaNa. The agent sees the game map and predicts what actions it should make to lead it to victory [12].

In the intersection of playing to win, playing for experience, and player modelling the focus is on ***believable play***. There is a direct link between player modelling and believable agents, as research carried out for the modelling of human, human-like, and supposedly believable playing behaviour can inform the construction of more appropriate models for players. However, an agent cannot be believable or existent to augment the game's experience, if it is not proficient.

[11] G. N. Yannakakis, and J. Togelius. *Artificial intelligence and games*. Springer, 2018.
[12] Image source: DeepMind - https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii

Being able to play a game well is in several ways a precondition for playing games in a believable manner. Nevertheless, player models can inform and update believable agent architectures. ***Models of behavioural, affective and cognitive aspects of gameplay*** can improve the human-likeness and believability of any agent controller—whether it is ad-hoc designed or built on data derived from gameplay. While the link between player modelling and believable agent design is obvious and direct, research efforts towards this integration within games are still sparse. However, the few efforts made on the imitation of human game playing for the construction of believable architectures have resulted in successful outcomes.

Procedural Content Generation (PCG) is one of the areas of recent academic research on AI in games, which bears the most promise for incorporation into commercial games. A number of recent games have been based heavily on PCG, including independent ("indie") game production successes such as *Spelunky* (Mossmouth, 2009) and Minecraft (Mojang, 2011), and mainstream AAA games such as *Diablo III* (Blizzard Entertainment, 2012), *Civilization V* (2K Games, 2010), *No Man's Sky* (Hello Games, 2016), Borderlands (Gearbox Software, 2009), and *Watch Dogs Legion* (Ubisoft, 2020) (see Figure 3). Some games heavily based on PCG and developed by researchers have been released as commercial games on platforms such as Steam and Facebook. Where PCG intersects game playing we can see agents that are capable of playing a game proficiently, which can be useful for ***simulation-based testing***, i.e., the testing of newly generated game content by playing through that content with an agent (see section on "*Reinforcement learning*"). Moreover, if an agent is trained to perform well in only a single game environment, it is easy to overspecialise the training and arrive at a policy/behaviour that will not generalise to other levels. Therefore, it is important to have a large number of environments available for training. PCG can help with this, potentially providing an infinite supply of test environments for different agents.

Figure 3: Procedural Content Generation is used for the generation of maps in Civilization games[13].

[13] Image source: The Scientific Gamer - https://scientificgamer.com/the-procedural-generation/

Societal and media industry drivers

Vignette: AI for content moderation in multiplayer games

Ella is a content moderator on a new multiplayer team shooter game. Ella's job is to investigate and evaluate cases of in-game harassment and cheating. In the past, these instances were investigated based on tickets submitted by players (see Figure 4), which made the process slow and unreliable. On one hand, an open ticket system was prone to abuse, while on the other, many victims didn't report harassment or cheating but stopped playing the game altogether instead. However, today Ella is working with automated AI tools, which analyze in-game behaviour, game telemetry, and chat logs to flag problematic interactions. The new system is able to identify not just severe cases of harassment but more subtle toxic behaviour, which used to go mostly unreported in the past. While the AI tool is not banning people from the game automatically, it alleviates much of the strain of having to look through gameplay logs manually, allowing Ella to focus on more crucial parts of the investigation. Additionally to identifying toxic behaviour, the tool Ella uses also flags potential instances of cheating and abuse of game mechanics. To retain players – beyond catching toxic behaviour – it is especially important in multiplayer games to minimise cheating and exploits. While cheaters can be reprimanded and banned if necessary, it is also important to identify exploits that can be patched out in later content updates. This way the AI-assisted monitoring of the game aids not just Ella's work as a content moderator but her designer and developer colleagues' as well.

Figure 4: Traditional in-game reporting tool for players in For Honor (Ubisoft, 2017)[14].

Future trends for the media sector

As highlighted in the previous sections, papers on procedural content generation, player modelling, game data mining, human-like playing behaviour, automatic game testing, and so on have become commonplace within the community, and industry adaptation of these methods are already on their way. However, almost all research projects in the game AI field are very specific. Most published papers describe a particular method—or a comparison of two or more methods—for performing a single task (playing, modelling, generating, etc.) in a single game. This is problematic in several ways, both for the scientific value and for the practical applicability

[14] A. Canossa, et al. "For Honor, for Toxicity: Detecting Toxic Behavior through Gameplay." Proceedings of the ACM on Human-Computer Interaction (CHI PLAY), 2021.

of the methods developed and studies made in the field. If an AI approach is only tested on a single task for a single game, how can we argue that is an advance in the scientific study of artificial intelligence? And how can we argue that it is a useful method for a game designer or developer, who is likely working on a completely different game than the one the method was tested on?

The focus of generality solely on play is very narrow as the possible roles of AI and general intelligence in games are many, including game design, content design, and player experience design. The richness of the cognitive skills and affective processes required to successfully complete these tasks has so far been largely ignored by game AI research. We thus argue, that while the focus on general AI needs to be retained, research on ***general game AI*** needs to expand beyond mere game playing. The new scope for general game AI beyond game-playing broadens the applicability and capacity of AI algorithms and our understanding of intelligence as tested in a creative domain that interweaves problem-solving, art, and engineering.

For general game AI to eventually be truly general, we argue that we need to ***extend the generality of general game playing*** to all other ways in which AI is (or can be) applied to games. We should develop methods that can model, respond to, and/or reproduce the very large variability among humans in design style, playing style, preferences, and abilities. This generality can be embodied in the concept of general game design, which can be thought of as a final frontier of AI research within games[15]. It is important to note that we are not arguing that more focused investigations into methods for single tasks in single games are useless; these are often important as proofs-of-concept or industrial applications and they will continue to be important in the future, but there will be an increasing need to validate such case studies in a more general context. We are also not envisioning that everyone will suddenly start working on general methods. Rather, we are positing generalisations as a long-term goal for our entire research community.

Finally, the general systems of game AI that we envision ought to have a ***real-world use***. There is a risk that by making systems too general we might end up not finding applications of these general systems to any specific real-world problem. Thus, the system's applicability (or usefulness) sets our core constraint towards this vision of general game AI. More specifically, we envision general game AI systems that are nevertheless integrated successfully within specific game platforms or game engines.

Goals for next 10 or 20 years

The shift of focus towards generalised AI solutions is a trend that is already taking place in other fields of AI research. So-called *foundation models* are on the rise in areas such as natural language processing and computer vision. These models are large pre-trained networks leveraging deep learning and transfer learning to provide easily adaptable AI for broad applications. It is not surprising that many of these models are the driving force behind innovation in many industries. While ***foundation models for games and gameplay*** are not existing yet, seeing the success of these models in other fields, the path seems clear. In the

[15] J. Togelius and G. N. Yannakakis. General General Game AI. In 2016 IEEE Conference on Computational Intelligence and Games (CIG). IEEE, 2016.

future, we are expecting to see both the more widespread integration of existing foundation models into games and the rise of ***game-based models*** in areas such as level design, believable game playing, and player modelling. The adaptation of existing models has already begun with games such as *AI Dungeon* (Latitude, 2019), which incorporates GPT-2 and GPT-3 foundational language models.

In the next decades, we expect to see innovations in different areas of the games industry driven by these foundation models. As the shift towards a "games as a service" model continues, the need for testing new content will exceed human capacity. ***Testing games*** for bugs, ***balancing player experience and behaviour***, and other issues is important in game development, and one of the areas where game developers are already looking for AI assistance. In the future, we expect to see the rise of foundation models which are built to navigate ad-hoc game spaces, test puzzles, or provide generalised feedback of player states. For the particular case of finding bugs and exploits in games, one of the research challenges is to find a good and representative coverage of problems, so as to deliver an accurate picture to the development team of how many problems there are and how easy they are to run into and allow prioritisation of which problems to fix.

More than a decade ago, the outstanding feature of *Left 4 Dead* (Valve Corporation, 2008) was its ***AI director***, which adjusted the onslaught of zombies to provide a dramatic challenge curve for players. While simple and literally a single dimension of player experience was tracked, the AI director proved highly effective. In the future, new advancements in deep learning will allow for much room for creating more sophisticated AI directors; the experience-driven PCG framework[16] is one potential way within which to work towards this.

We also expect to see the rise of ***AI-Based Game Design***. This could be seen as an opportunity to showcase AI methods in the context of games, but it could also be seen as a way of advancing game design. Most classic game designs originate in an era where there were few effective AI algorithms, there was little knowledge among game designers about those AI algorithms that existed, and CPU and memory capacity of home computers was too limited to allow anything beyond simple heuristic AI and some best-first search to be used. One could even say that many classic video game designs are an attempt to design around the lack of AI—for example, the lack of good dialog AI for NPCs led to the use of dialog trees, the lack of AIs that could play First-person shooter (FPS) games believably and competently led to FPS game designs where most enemies are only on-screen for a few seconds so that you do not notice their lack of smarts, and the lack of level generation methods that guaranteed balance and playability led to game designs where levels did not need to be completable. The persistence of such design patterns may be responsible for the relatively low utilisation of interesting AI methods within commercial game development. Advancements in computation and AI optimisation in the next decades will certainly make AI methods far more accessible in the future. By starting with AI and designing a game around it, new design patterns that actually exploit some of the recent AI advances can be found.

[16] G. N. Yannakakis and J. Togelius. Experience-driven procedural content generation. Affective Computing, IEEE Transactions on, 2(3):147–161, 2011.

Finally, as AI becomes ***more ubiquitous and easily adaptable*** in the future, we expect to see the rise of ***AI-Assisted Game Development***. By combining the AI techniques mentioned above, new co-creation algorithms will be able to aid developers in providing faster iteration times and greater customisation in content creation. Beyond leveraging computational creativity to create new content and various AI methods to test on the fly, the future will most probably see a push for more personalised content as foundational player models are becoming more widely used. Future AI applications might be able to create new content or fine-tune the experience of any game for any emotional state, completing a meaningful affective loop, which defines a framework that is able to successfully elicit, detect and respond to the cognitive, behavioural and emotive patterns of the players[17].

[17] P. Sundstrom. Exploring the affective loop. PhD thesis, Stockholm University, 2005.

AI4media

ARTIFICIAL INTELLIGENCE FOR THE MEDIA AND SOCIETY

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951911

info@ai4media.eu www.ai4media.eu