eludamos
Journal for Computer Game Culture
Vol. 10, No. 1 (2019)

http://www.eludamos.org

Games and AI: Paths, Challenges, Critique

Mathias Fuchs, Andreas Sudmann

Eludamos. Journal for Computer Game Culture. 2019; 10 (1), pp. 1–7

Games and AI: Paths, Challenges, Critique

MATHIAS FUCHS AND ANDREAS SUDMANN

In 2006, the first-person shooter *F.E.A.R.* makes headlines in the gaming world. One feature in particular attracts much attention: the non-playable characters seem to behave intelligently to a degree yet unseen in computer games. From earlier productions like *No One Lives Forever 1* & *2* (2000, 2002), players were already familiar with NPCs that are able to seek cover under fire and to leave it at random in order to shoot back at the player. In *F.E.A.R.* that happens too, but in a much more realistic manner. Computer-controlled enemies attack players in a coordinated way. If one member of the enemy team comes closer, he gets supportive fire by his team members. If the player attacks them, enemy forces remain in cover until they are immediately threatened.

Ten years later, an AI system called *AlphaGo* beats the human world champion Kim Sung Yong in the ancient board game Go in five rounds—final score: 4-1. The global community of Go players is perplexed, almost shocked, even though the victory did not totally come out of the blue. Already in October 2015, an earlier version of *AlphaGo* was able to beat the European Go champion Fan Hui. However, Hui's playing level was significantly lower than that of Kim Sung Yong (2-dan out of possible 9-dan levels).

As these introductory examples illustrate, the relationship between artificial intelligence (AI) and games can basically be studied from two perspectives: The first is the implementation of AI technologies in games, in order to improve the game experience in one way or another, for example with the intention to make it more believable, more immersive, or simply more enjoyable. The second is the use of games as a benchmark, a learning or test environment to evaluate, but also demonstrate, the current state of AI technologies. Both perspectives have gained enormous importance in recent years—technically, but also culturally and economically.

Of course, games and AI technologies have been closely related for decades. Even several years before AI became the official name for a respective research program in 1956, when the Dartmouth conference took place, Alan Turing's famous test has been set up as an "imitation game" (Turing 1950). Instead of deciding, whether a machine can be intelligent or not, Turing proposed a test to evaluate if a machine can at least *simulate* or *imitate* intelligence successfully. In practice, despite reports to the contrary, no machine has so far passed Turing's test. Nevertheless, there is no question that the communication skills of programmable devices have improved enormously. For example, one can say of Google's Duplex system that it eerily resembles the personal assistant Samantha in Spike Jonze's movie *Her* (2014).

Indeed, the gap between AI as sci-fi and AI as empirical technology seems to be narrowing (Ernst et al. 2019). And whether it comes to personal language assistants, self-driving cars, or spectacular competitions between humans and machines in

Eludamos Vol. 10, No. 1 (2019)
pp. 1–7

games like Go, most of the progress has been enabled by a specific approach of AI: machine learning, i.e. statistical prediction methods which can solve tasks and problems without being explicitly programmed for this purpose (Engemann/Sudmann 2018). Such approaches have been established for decades, even in application to games. Already in the early 1950s it was Arthur Samuel who was able to use machine learning algorithms in order to teach a computer how to play checkers (Samuel 1959).

A particularly successful approach of machine learning is the use of artificial neuronal networks whose principal function is loosely based on a model of biological brains and whose most important characteristic is the massive parallelism of information processing (LeCun/Bengio/Hinton 2015). This property enables AI systems to succeed in areas where they have traditionally been inferior to humans, such as dealing with blurred and unsecure information.

The principal technology of artificial neural networks (ANN) has also already been developed back in the 1950s (and even before). However, apart from a short boom phase, the ANN approach was generally considered to be a dead end of AI research.[1] This has changed considerably since 2012, at the latest. At that time, a research team at the University of Toronto was able to win the so-called ImageNet Competition by training so-called convolutional neural networks (CNN) with GPU processors, i.e., fast, parallel-operating hardware, as has been used and developed for modern game consoles and PCs. In other words: The specific hardware requirements of the gaming industry were indirectly instrumental in enabling significant advances in the field of AI, in particular with regard to the field of computer vision. Furthermore, in 2013, DeepMind, the company responsible for AlphaGo, introduced a system to train an AI to master classic console games such as *Breakout* or *Pong* better than humans by means of a specific machine learning approach called reinforcement learning and purely by observing what happened on the screen and incorporating a single requirement: Maximize the score.

The fact that duels between man and machine are particularly effective in raising public awareness is hardly surprising. Games like Go or Chess are regarded as a great intellectual challenge, and if the machine is able to outperform the human being on this terrain, this obviously makes a lasting impression. Ironically, new conditions of machine learning methods and especially neuronal networks have ensured that computers are becoming better and better in terms of those tasks which to master is generally easy for humans, but not for machines: for example, to recognize images and differentiate sounds, in other words, to have a simple common-sense understanding of the environment. In this respect, too, games have an important epistemic-technical function: the hope is that AI systems can be trained with games as simulations and that this training experience can also be used for the empirical world. Exactly this form of transfer, however, does not work. For example, it is at least at this moment in time impossible to train self-driving cars with games and reinforcement learning to be used for all variations of real roads and environments.

This issue of *Eludamos* is, however, not so much a technical account of what AI is now able to accomplish or not, but rather it is a critical investigation into the socio-cultural contexts of AI research, including its phantasmatic and ideological dimensions. In addition, it also takes a look at the economic conditions of AI as a finally profitable technology.

The European, Asian, as well as US government provide huge amounts of funding for industry and research. Also do tech companies like Google/Alphabet, Baidu, or Facebook massively invest in AI technologies. Indeed, there is no question that we currently witness an unprecedented AI gold rush, even though the revenues generated by smart technologies are still relatively modest. The worldwide AI revenues in 2018 are estimated at anything between 8 billion US dollars (tractica 2019) and a more optimistic 9.51 billion US dollars (statistica 2019).

This is much less than the 135 billion US dollars that that the games industry is said to have made in 2018 (gamesindustry.biz 2018). Yet the conclusion that games sell ten times more than mainly AI-powered products is misleading, since the borders between AI and games technology have become quite blurry and the respective research fields are also not mutually exclusive. Of course, this specific development cannot be ascribed to the current boom in machine learning and artificial neural networks alone. For example, in the last twenty years or so the design and implementing of game worlds and gameplay has still been largely based on classical, symbolic forms of AI as opposed to machine learning approaches. However, there are numerous indications that the latter will increasingly influence and shape all operations of the game industry, as we can already witness with regard to the testing and marketing of games (Ariyurek/Betin-Can/Surer 2019).

The big players in the game industry including leading hardware manufacturers are also working on AI solutions in non-game environments. For example, Unity Technologies, the company behind the Unity game engine, has also recently begun using machine learning tools (called Unity ML-Agents Toolkit) and by doing so strategically addressing a new clientele: the automotive industry. The expansion of the business model also has an ethically significant dimension. In the past, the company did not have to worry too much about whether virtual humans were killed by virtual cars. Now they deal with real cars potentially killing real humans. The company is about to leap out of the magic circle of designing purely for play without existential threats and into the domain of real-world impacts. That is why they are now levelling up their ethical code. In “Unity’s Guiding Principles for Ethical AI” (Unity 2018), the company formerly specialized in computer game development, states: “It’s our mission every day to democratize development, solve hard problems, and enable success for creators in diverse industries such as games, automotive, architecture, engineering, and construction (AEC), media and entertainment.“

Why does the producer of one of the most popular game engines care about the automotive and construction industry? A closer look at recent *Unity* apps shows that this is not merely a company’s diversification, but that there is a strong logical link in between computer games and autonomous vehicles. The *Unity* demo “autonomous cars” demonstrates that making a vehicle run without human control is a challenge for cars both in games and on physical roads in the empirical world. Reinforcement learning is used for both of these situations. However, cars that have been trained with games (or simulations) cannot yet cope with the challenges of real roads and environments. And yet it is significant to find Audi and Toyota as mentors/sponsors on board of the games company, praising *Unity* for its benefits for the car manufacturers. Unity Technologies, on the other hand, is proud to show up at and actively participate in AI conferences like the AAAI-19 in Hawaii or the NeurIPS in Montréal (Unity blog 2019). Other companies act similarly. Ubisoft, Microsoft and NVIDIA are also very active in the AI market. NVIDIA, for example, has a few years ago been exclusively associated with graphics cards. Now they produce deep learning-tailored hardware like the NVIDIA V100 Tensor Core GPU and their "Inception Programme […] nurtures dedicated and exceptional startups who are revolutionizing industries with advances in AI and data science" (NVIDIA 2019).

The rhetorical claim that games and AI would revolutionize industries and eventually our lives, sounds like an echo from half a century ago. Shannon's speculations about AIs playing chess, Samuel's checkers automation, and many other game-based proofs for the simulation of a human problem-solving capability preceded the symbiosis of games/gaming and AI that we now observe becoming more and more a reality. While many promises of the 1950s led to disappointments for the time being, there were sometimes predictions that turned out to be relatively correct. In 1957, Herbert Simon predicted four major developments of AI that were to be accomplished during the following decade: Besides mathematical problem solving, psychoanalytical understanding and computerized composition, he announced that a chess playing program will defeat the world champion (Newell/Simon 1961). In fact, it took a bit longer than the ten years Simon expected, but in 1996/1997 IBM's *Deep Blue* was able to defeat world champion Kasparov.

The challenge of creating a "strong AI" (Russell/Norvig 1995: 27) or a "human-level AI", as others put it, has at times been dismissed as unrealistic, especially during the AI winter of the 1970s, but today we witness a revival of this idea when scholars discuss what is now called "open-ended AI" (Uber 2019). It seems important in the context of game studies to carefully revisit the terminology used in AI discourses. When do we and when should we talk about "AI", "classical AI", "Artificial General Intelligence"/AGI (Goertzel/Pennachin 2007), or about "Computational Intelligence" (cf. 9th Computer Science and Electronic Engineering Conference, Computational Intelligence and Games). What do such Nietzschean creations like "Uber AI" try to suggest? Do games like *Horizon: Zero Dawn* (2018) prepare us for a co-evolution with AIs when we are playing against super-intelligences like GAIA[2]? First in digital games and later in real life? How do serious purpose implementations of AI differ from ludic implementation in the entertainment sector? Are games AIs just mimicking research done in the Watson Labs, the Uber Labs or Elon Musk's laboratories? In what ways do games play a relevant role in the shaping of AI?

Games and game technology seem to pave the path to arrive at complex levels of simulation, creation, and contextualisation. Computer games development strives to increase the level of realism or of make-belief and so does AI. Furthermore, clusters of gaming consoles have been used to perform tasks that formerly could only run on supercomputers (Escribano 2012). Artificial neural network technologies are amongst those approaches favoured by the games industry to implement machine learning in game design, for example to make NPCs look and react more human-like—or at least to create the impression of human-like behaviour. Procedurally generated environments based on modern AI algorithms modify our perception of what artificial worlds could look like and shift the border of distinction between man-made 3D environments and machine-built levels.

This issue of *Eludamos* presents papers investigating the relation of games and AI from the viewpoints of sociological, systematic, philosophical, psychological, design, and programming scholarship. The articles cover topics like AI conversation, emotional game characters, procedural content generation, gamification, and human computer interaction, amongst others. With a field as wide and growing as the one of games and AI research, there will inevitably be a lot of questions to be asked in the near future that we do not think about today but that will become relevant and urgent in the future. For the time being, however, it may be enough to give correct answers to familiar questions.

Games and Movies Cited

Guerrilla Games (2017) *Horizon Zero Dawn* (PlayStation 4).

Jonze, Spike (2014) *Her* (Warner Bros. Pictures).

Monolith Productions (2005) *F.E.A.R.* (PC, PlayStation 3, Xbox 360).

Monolith Productions (2000) *No One Lives Forever* (PC).

Monolith Productions (2002) *No One Lives Forever 2* (PlayStation 2).

References

Ariyurek, S., Betin-Can, A., and Surer, E. (2019) Automated Video Game Testing Using Synthetic and Human-Like Agents. Available from: https://arxiv.org/abs/1906.00317v1 [Accessed 13 Jan. 2020].

Engemann, Ch., Sudmann, A., eds. (2018) *Machine Learning. Medien, Infrastrukturen und Technologien der Künstlichen Intelligenz*. Bielefeld: Transcript.

Ernst, Ch., Kaldrack, I., Schröter. J., and Sudmann, A. (2019) Künstliche Intelligenzen. Introduction. Special Issue „Künstliche Intelligenzen“, *Zeitschrift für Medienwissenschaft* 21, pp. 10-19.

Escribano (2012) Gamification as the Post-Modern Phalanstère - Is the Gamification Playing With Us or Are We Playing With Gamification? In Zachariasson, P. & Wilson, T. (eds.) *The Video Game Industry: Formation, Present State, and Future*. Routledge.

gamesindustry.biz (2018) Global games market value rising to $134.9bn in 2018. Available from: https://www.gamesindustry.biz/articles/2018-12-18-global-games-market-value-rose-to-usd134-9bn-in-2018 [Accessed 13 Jan. 2020].

Goertzel, B. and Pennachin, C. (2007) *Artificial General Intelligence*. Springer.

LeCun, Y., Bengio, Y and Hinton, G. (2015) Deep Learning. *Nature*, 521 (May), pp. 436–444.

Newell, A. and Simon, H. A. (1961) GPS, a program that simulates human thought. In Billing, H. (ed.) *Lernende Automaten*. Oldenbourg, pp. 109–124.

NVIDIA (2019). Deep learning and AI Startup Incubator. Available from: https://www.nvidia.com/en-us/deep-learning-ai/startups/ [Accessed 13 Jan. 2020].

Russell, S. J. and Norvig, P. (1995) *Artificial Intelligence. A Modern Approach*. Prentice Hall.

Samuel, A. L. (1959) Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development* 3, pp. 211–29.

statistica (2019) Worldwide artificial intelligence market revenue. Available from: https://www.statista.com/statistics/621035/worldwide-artificial-intelligence-market-revenue/ [Accessed 13 Jan. 2020].

tractica (2019) Artificial Intelligence Software Market to Reach 105.8 Billion in Annual Worldwide Revenue by 2025. Available from: https://www.tractica.com/newsroom/press-releases/artificial-intelligence-software-market-to-reach-105-8-billion-in-annual-worldwide-revenue-by-2025/ [Accessed 13 Jan. 2020].

Turing, A. (1950) Computing machinery and intelligence. *Mind*, Vol. LIX (236), 1 October 1950, pp. 433–460.

Uber (2019) Uber AI. Available from: https://www.uber.com/de/de/uberai/ [Accessed 13 Jan. 2020].

Unity (2018) Introducing Unity’s Guiding Principles for Ethical AI. Available from: https://blogs.unity3d.com/2018/11/28/introducing-unitys-guiding-principles-for-ethical-ai/?_ga=2.147125054.881379825.1564238426-901192280.1564238426 [Accessed 13 Jan. 2020].

Unity blog (2019) Available from: https://blogs.unity3d.com/2019/01/18/fostering-ai-research-meet-us-at-aaai-19/ [Accessed 13 Jan. 2020].

Notes

[1] The establishment of backpropagation (as the central learning algorithms in the field of artificial neural networks) at the end of the 1980s did indeed bring about only a brief upswing in neural networks. Computers were still too slow, and data infrastructures were lacking in order to train networks for more advanced learning tasks.

[2] GAIA was Project Zero Dawn’s governing A.I., “the most powerful, most advanced AI ever created.”The game tells us that GAIA, the AI, had at some point to get rid of the humans in order to save the planet.