EasyChair Preprint

№ 4921

Artificial Intelligence in Video Games

Ahmed Behe Eddine Cherif and Marwen Ben Salem

January 22, 2021

Artificial Intelligence in Video Games

Ahmed Behe Eddine Cherif, Marwen Ben Salem

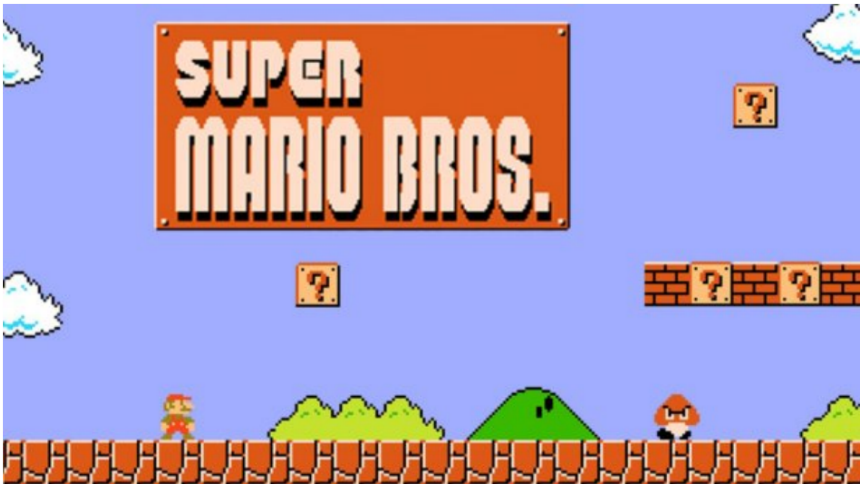

Figure 1: *The first Super Mario Bros. level.[9]*

Abstract

This work will explore three methods of machine learning that make it possible to train an algorithm to the extent that it can play the video game Super Mario and outperform human players. The aim is to find out which method of machine learning is best suited and what the differences are.

Keywords: Supervised Learning, Unsupervised Learning. Supervised Learning Reinforcement Learning. Artificial Intelligence, Video Games.

1 Introduction

Video games developed from technical experiments at universities in the 1950s to one of the most influential recreational forms of the 21st century. [2].
Video games differ in a number of factors. Some allow playing against human opponents, such as Pong from 1972 or several online multiplayer games today. However, there have been computer-controlled opponents since the beginning (1971). [14]
Artificial intelligence (AI) and machine learning, on the other hand, have only gained increasing importance in many different areas of application in recent years. In the video game industry, AI is predominantly used for the behavior of non-player characters (NPCs) in games [13].
A given pattern is often used for this task, which reacts to certain player behavior but does not learn beyond that and adapts to new circumstances. Techniques involving decision trees and path finding are often used. [14] A learning AI would enable the player, in addition to the standardized difficulty settings (easy, normal, difficult), new experiences and create an individual gaming experience. In this thesis, an AI is to be trained with the help of supervised, unsupervised and reinforcement learning (methods of machine learning) to beat human records in the game Super Mario.

2 State of the art

The following section aims at the three most popular methods:

- Supervised Learning.
- Unsupervised Learning.
- Reinforcement Learning.

for machine learning. A critical comparison will then identify the most suitable method.
DeepMind Technologies from Google has already taught an AI how to play Atari with deep reinforcement learning [18]. In early 2019, DeepMind released the AI called AlphaStar, which defeated a professional player for the first time in Star-Craft II. [12] This was also trained with deep reinforcement learning, which will be an important approach to solving the problem and is a frequently used method.

What is machine learning?

People act intuitively in tasks such as: facial recognition of relatives, driving a car or playing video games. This depends on the experience, knowledge and curiosity already gained. The concept of intuition is a challenge for algorithms. Machine learning consists of an algorithm that solves a problem based on training data and possibly past experience and optimizes itself in the process. [1].

A model is generated that is equipped with many parameters, which are adapted and optimized through learning in training or through experience. The resulting algorithm can either make predictions or conclusions and / or act descriptively in order to collect further knowledge. Another important part of machine learning is to make this algorithm “intelligent” by learning to adapt to new conditions in a changeable environment [1]. A suitable example would be self-driving cars, which have to constantly adapt to new changes like others Adapt cars, cyclists or pedestrians and make predictions accordingly, react and learn from them.

2.1 Supervised Learning

In the publication Deep Learning for video games [16], deep learning methods and examples were presented for various genres of games. For the supervised learning method, a model is trained with sample data. In the process, decisions are to be made by the al-

gorithm. The correct decisions are known well in advance. The differences between the correct answer and the answer from the algorithm are compared. If the algorithm was wrong, this negative result is used to update the model. All training data are classified in advance or belong to a certain group. In this process, the algorithm has to recognize the classification and categorize it correctly for further data. More model receives data, the algorithm becomes more reliable. In the end, the goal is to ensure that this model can deal well with unknown training data with right decisions.
In order to get training data for these methods, for example : Recorded people playing, another option is already available [15], it is the use of already existing data that another algorithm has already achieved. This data can be interpreted by an algorithm, that's called agent. The agent can recognize what the game looks like and what actions were taken in what situation to achieve a predefined output (for example: reach the end of the level or find a number of coins) (for example of game the first Super Mario level)." [16]
The publication [10] describes how the algorithm exactly recognizes and treats a particular game and its elements, and what effects it has on the player when only a few elements are obscured or made unrecognizable. The playing time of a level can increase from 2 minutes to 20 minutes. The biggest aspect here is with experiences and knowledge applied to recognize certain objects like a ladder and associate with them like climbing a ladder to reach higher objects [10]. These experiences can be derived from the training data and made available to the algorithm.
Supervised learning is also often used to learn state transitions in a game. The next state is predicted, it's planning the next move [6]. This can be used to play chess.

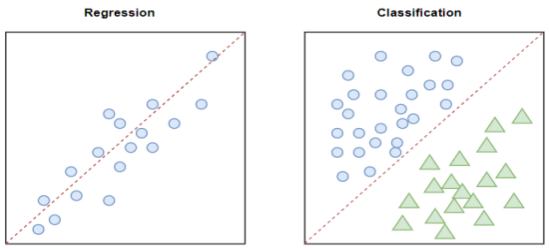

Figure 2: Regression and classification in supervised learning [11]

2.2 Unsupervised Learning

In unsupervised learning, the goal for the algorithm is to recognize and learn a pattern in the data based on the characteristics. By distributing characteristics, similar data can be grouped together (cluster analysis) or reduced (compression). Another task is to create artificial data that closely resembles the original data. Unlike supervised learning, unsupervised learning does not involve known or correct answers given beforehand. This method is more commonly used for games with very few rewards, such as text-based adventure games where that's have to recognize textual connections. [16]

2.3 Supervised Learning Reinforcement Learning

The reinforcement learning method includes some important ones Elements:

- Agent.
- Environment
- Reward
- Action
- State

Figure 3: Cluster analyse in Unsupervised Learning [11]

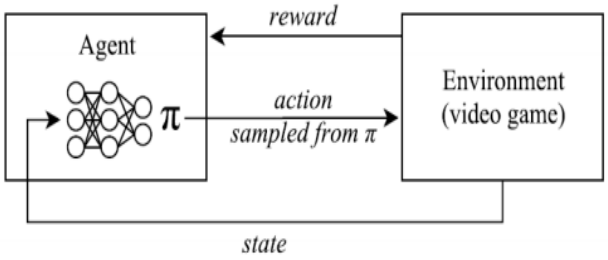

Figure 4: Reinforcement Learning [16]

All related as shown in the figure above. The agent receives a certain state from the environment. The agent then interacts with the environment. The agent receives a limited number of actions for each training session. With these, certain behaviors are learned that are needed to achieve the goal. The environment gives the agent a signal in the form of a reward. The aim is to get as many rewards as possible. This learns the neural network behind the agent. The rewards can appear regularly or irregularly, depending on the game principle. In games where points are collected in order to crack a high score, rewards are repeatedly returned to the agent in the course of the game (regular rewards). However, the only reward can be to complete a certain level or defeat an opponent (irregular reward) [16]

3 Comparison of the most relevant approaches

The two most relevant approaches to the question would be on the one hand supervised learning and on the other hand reinforcement learning. This section describes the advantages and disadvantages of the two methods.
Supervised learning is pretty good at applying best practices to existing data. This could work well for the game Super Mario in that a human has already played the game and the paths and behaviors derived from it are used to make the algorithm learn. The disadvantage is that this process is very time-consuming, since the training data must first be recorded and processed. Furthermore, it can happen that some game elements such as paths, collectibles or secrets were not discovered in the training data and the algorithm cannot find them either. This can mean that only parts of a game are played or the agent does not know how to deal with new scenarios. More training data from other players could solve this problem, which in turn takes a lot of time. Furthermore, the agent will for the most part only orient itself on the paths already given and will not find out further elements that the game offers by itself. [16] However, an approach with Self Supervised Prediction, where the agent lets the agent explore the world with a form of supervised learning without

rewards but with the agent's "curiosity" [8]. The first level of Super Mario is used to train the algorithm in order to become faster and better in further levels with the experience already gained.
The basic elements that reinforcement learning needs can be modeled well in video games. The game character can be used as an agent, the environment is for example, a certain level that should be completed successfully. Most games also rely on several reward strategies, coins can be collected in Super Mario and the goal of the level should be achieved.
In open world games, games in which the player can do or discover things in a world outside of a certain mission, reinforcement learning is more difficult because the reward system is structured differently than in linear games [16] Here the algorithm becomes much more complex. However, this is not the case in Super Mario. Games that give little or no rewards present a challenge. According to the motto "the journey is the goal", after a reward has been received, a decision must be made as to how the previous actions or actions of the agent that led to its success are to be rewarded [16] An example of a game with no rewards at all would be Human Fall Flat. This is a platform puzzle game with the aim of reaching the exit which can be reached by solving physics puzzles. No coins, points, diamonds or anything else are collected here and there are no opponents to fight against. There are various approaches to solving this problem, For example, Markov Decision Process (MDP), temporal difference (TD) learning or Q-learning which would go beyond the scope of this work. [16] With Super Mario this would still have to be considered in part, since the journey and past experiences have to be rewarded.

4 Description of the solution details

Unsupervised learning is out of the question for the problem, since clustering or reducing data would not lead to the goal. Both supervised learning and reinforcement learning can be considered for solving the problem. Depending on the requirements, e.g. Complexity, time required, costs etc. are placed on the algorithm, one method is more suitable than the other. A suggested solution would be that also by DeepMind in the game Go [7] was used to combine both methods. With this idea, the algorithm can first be trained with the help of already existing data, which concerns the basic functions of the game. In Super Mario this would be running, crouching, jumping, collecting coins and dodging enemies. Reinforcement learning is now used to minimize the effort of all possible ways or collectibles. This method allows the agent to independently explore the level to find more rewards. Abbreviations found, collectibles and the actions taken are rewarded.

5 Conclusion

In this work it was shown how the video game Super Mario could be played by an algorithm using the methods supervised learning and reinforcement learning, for example Beat personal bests. With supervised learning, an algorithm can follow a certain path using specified training data without independently discovering new game elements. Unsupervised learning is certainly more suitable for clustering or reducing data. Reinforcement learning relies on an algorithm that independently explores the world of the game and adapts its behavior through rewards.
It turned out that supervised learning and reinforcement learning would be best suited for the game. Like Deepmind [7] already showed, these can be used in combination to solve different problems and to achieve a defined complexity.
The approach described corresponds to my personal requirements. Alternatively, an implementation with reinforcement learning could save time, since the inclusion of the player would be omitted. Furthermore, the approach of [8] is of great interest, since self-supervised learning is used here and the agent takes the learned and found interaction elements in the game and uses them in other levels in order to continue playing faster and more efficiently.

References

[1] E ALPAYDIN. "Introduction to Machine Learning". In: *(Adaptive 261 Computation and Machine Learning), vol. 4. MIT Press.* (2020).

[2] A Astinus. "Der Retro- und Kult PC-Spiel Report: Von Klassikern wie Donkey Kong, Centipide, Scramble und Moon Patrol". In: (2013).

[3] S . H Ather. "Reinforcement Learning: Super Mario, AlphaGo and beyond". In: *https : / / medium . com / @HussainAther / reinforcement - learning - super - mario - alphago-and-beyond-fdd676d1bf43* (Oct. 2018).

[4] Amato Christopher and Shani Guy. "High-level reinforcement learning in strategy games". In: *AAMAS '10: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1. https://dl.acm.org/doi/abs/10.5555/1838206.1838217* (May 2010), pp. 75–82.

[5] Wiles Daniel Johnson andJanet. "Computer games with intelligence". In: *10th IEEE International Conference on Fuzzy Systems. https : / / ieeexplore . ieee . org / abstract / document/1008909* (2002).

[6] Ha David and Schmidhuber Jurgen. "World Models". In: *https://arxiv.org/pdf/1803.10122.pdf* (2018).

[7] Silver David et al. "Mastering the game of Go with deep neural networks and tree search". In: *https://pubmed.ncbi.nlm.nih.gov/26819042/* (2016).

[8] Pathak Deepak et al. "Curiosity-driven Exploration by Self-supervised Prediction". In: *https : // pathak22 . github . io / noreward-rl/resources/icml17.pdf* (2017).

[9] Michel Dennis. "Super Mario Maker 2: So brutal habt ihr Level 1-1 noch nie gesehen". In: *https://www.gamepro.de/artikel/super-mario-maker-2-hoellisches-level,3346714.html* (2019).

[10] Rachit Dubey et al. "Investigating Human Priors for Playing Video Games". In: *https://arxiv.org/abs/1802.10217* (2018).

[11] ECLOUDVALLEY. "Introduction to Machine Learning: Is 292 AutoML Replacing Data Scientists?" In: *https://www.ecloudvalley.com/mlintroduction/* (2019).

[12] Arulkumaran Kai, Cully Antoine, and Togelius Julian. "AlphaStar: an evolutionary computation perspective". In: *GECCO '19: Proceedings of the Genetic and Evolutionary Computation Conference Companion. https://dl.acm.org/doi/10.1145/3319619.3321894* (July 2019), pp. 314–315.

[13] Merrick Kathryn and Lou Maher Mary. "Motivated reinforcement learning for non-player characters in persistent computer game worlds". In: *ACE '06: Proceedings of the 2006 ACM SIGCHI international conference on Advances in computer entertainment technology. https://dl.acm.org/doi/abs/10.1145/1178823.1178828* (Mar. 2006).

[14] Böhm Matthias. "Wie wird künstliche Intelligenz in Videospielen eingesetzt?" In: *https : // googleseo . de / blog / wissensdatenbank / wie - wird - kuenstliche - intelligenz - in - videospielen-eingesetzt/* (2020).

[15] Bogdanovic Miroslav et al. "Deep Apprenticeship Learning for Playing Video Games". In: *Twenty-Ninth AAAI Conference on Artificial Intelligence. https://aaai.org/ocs/index.php / WS / AAAIW15 / paper / view / 10159 / 10213* (2015), pp. 7–9.

[16] Justesen Niels et al. "Deep Learning for Video Game Playing". In: *IEEE Transactions on Games 12. https : // ieeexplore . ieee . org / abstract / document / 8632747* (Feb. 2019), pp. 1–20.

[17] R. SATHYA and A ABRAHAMS. “Comparison of supervised and unsupervised learning algorithms for pattern classification. 80”. In: *http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.278.5274&rep=rep1&type=pdf* (2013).

[18] Mnih Volodymyr et al. “Playing Atari with Deep Reinforcement Learning”. In: *https://arxiv.org/pdf/1312.5602v1.pdf* (Dec. 2013), pp. 1–9.

war-submission-example