MARTIN BÄUMLER | 2017-10-24

AI & GAMES: WHY RESEARCHERS WANT AI TO DABBLE IN GAMES

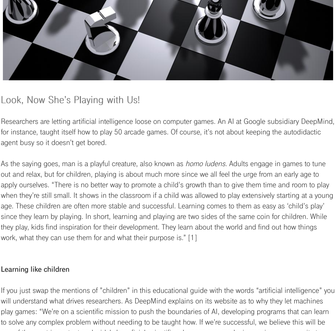

Look, Now She's Playing with Us!

Researchers are letting artificial intelligence loose on computer games. An AI at Google subsidiary DeepMind, for instance, taught itself how to play 50 arcade games. Of course, it's not about keeping the autodidactic agent busy so it doesn't get bored.

As the saying goes, *man is a playful creature*, also known as *homo ludens*. Adults engage in games to tune out and relax, but for children, playing is about much more since we all feel the urge from an early age to apply ourselves. "There is no better way to promote a child's growth than to give them time and room to play when they're still small. It shows in the classroom if a child was allowed to play extensively starting at a young age. These children are often more stable and successful. Learning comes to them as easy as 'child's play' since they learn by playing. In short, learning and playing are two sides of the same coin for children. While they play, kids find inspiration for their development. They learn about the world and find out how things work, what they can use them for and what their purpose is." [1]

Learning like children

If you just swap the mentions of "children" in this educational guide with the words "artificial intelligence" you will understand what drives researchers. As DeepMind explains on its website as to why they let machines play games: "We're on a scientific mission to push the boundaries of AI, developing programs that can learn to solve any complex problem without needing to be taught how. If we're successful, we believe this will be one of the most important and widely beneficial scientific advances ever made, increasing our capacity to understand the mysteries of the universe and to tackle some of our most pressing real-world challenges. From climate change to the need for radically improved healthcare, too many problems suffer from painfully slow progress, their complexity overwhelming our ability to find solutions. With AI as a multiplier for human ingenuity, those solutions will come into reach." [2]

DeepMind doesn't only use classic arcade games to teach its Artificial Intelligence, but also real-time strategy games like StarCraft and StarCraft II. Those games are among the biggest and most successful games of all times, drawing in players to compete for more than 20 years. Researchers in the fields of machine learning and AI have been using the original for quite some time, even having their bots duke it out in an annual competition called AIIDE StarCraft AI Competition. [3]

StarCraft is the perfect research environment for Artificial Intelligence training

"StarCraft's ... multi-layered gameplay ... makes it an ideal environment for AI research. For example, while the objective of the game is to beat the opponent, the player must also carry out and balance a number of sub-goals, such as gathering resources or building structures. In addition, a game can take from a few minutes to over 30 minutes to complete, meaning actions taken early in the game may not pay off for a long time. Finally, the map is only partially observed, meaning agents must use a combination of memory and planning to succeed."

"The game also has other qualities that appeal to researchers, such as the large pool of avid players that compete online every day. This ensures that there is a large quantity of replay data to learn from - as well as a large quantity of extremely talented opponents for AI agents.

Even StarCraft's action space presents a challenge with a choice of more than 300 basic actions that can be taken. Contrast this with Atari games, which only have about 10 (e.g. up, down, left, right etc). On top of this, actions in StarCraft are hierarchical, can be modified and augmented, with many of them requiring a point on the screen. Even assuming a small screen size of 84x84 there are roughly 100 million possible actions available." [4]

An Artificial Intelligence that can't tell a lie

This approach has not only attracted Google-affiliated companies. Guillaume Lample and Devendra Singh Chaplot, two students at Carnegie Mellon University also let their Artificial Intelligence compete in highly complex 3D games. "We present the first architecture to tackle 3D environments in first-person shooter games, that involve partially observable states. Typically, deep reinforcement learning methods only utilize visual input for training. We present a method to augment these models to exploit game feature information such as the presence of enemies or items, during the training phase. Our model is trained to simultaneously learn these features along with minimizing a Q-learning objective, which is shown to dramatically improve the training speed and performance of our agent. Our architecture is also modularized to allow different models to be independently trained for different phases of the game. We show that the proposed architecture substantially outperforms built-in AI agents of the game as well as humans in deathmatch scenarios." [5]

It's nothing new that first-person shooters battle AI enemies. The computer game Doom, for which two junior researchers have developed their Artificial Intelligence, routinely matches players with virtual opponents. But usually, they're of the cheating kind since they're part of the game and therefore know the locations of a human player, the scores and different scenarios. They're just following pre-programmed patterns. The AI built by the two CMU students, though, gains its own experience and learns from its successes and failures. It acts like a human player and makes its own decisions in real-time. It only knows what a human opponent would know.

Games level up safety for autonomous driving

The researchers at DeepMind and CMU, though, don't want to develop more sophisticated and experienced opponents to make their games more interesting. Lample and Chaplot explicitly name autonomous driving as one focus area of their development work.

Machine learning enables computers to perform impressive new feats like identifying a face, recognizing a language or a person. It usually requires huge amounts of curated data and it can be difficult and time-consuming to amass enough datasets. Open scenarios in particular make it nearly impossible to run through every possibility in advance. Autonomous driving is exactly that kind of challenge – there are simply too many things that can happen on the road. A research team from Intel Labs and the University of Darmstadt has come up with a clever method to extract useful training data from Grand Theft Auto.

High-capacity models that were developed for large datasets have driven recent advances in image processing. Unfortunately, it was extremely expensive to compile such large data sets with pixel-level labeling since it requires human effort. The research team at Intel Labs and Darmstadt University has now presented an approach to draw pixel-accurate semantic label maps of images extracted from computer games. Even though they couldn't access the source code and internal operation of commercial games, they demonstrate that one can reconstruct the associations between image patches by following the communications between the game and the graphics hardware.

It allows for a quick propagation of semantic markers within and across images which were synthesized by the game, all without accessing the source code or the content. The researchers validate their method by generating semantic annotations with a dense pixel layer for 25,000 images that were synthesized by a photo-realistic open-world computer game. Experiments on semantic segmentation datasets show that you can improve accuracy by using the acquired data to augment real-world images and that it's possible to reduce the amount of hand-labeled real-world data. Models that were trained with game data outperform models trained with the usual CamVid training set. [6]

While humans can suffer negative consequences from playing games like Grand Theft Auto for hours on end and often do so to escape into hyperrealistic worlds, AI systems can derive great benefits if they just keep at it. It's not without irony that we improve autonomous driving and our roads by racing down virtual streets in daredevil fashion and by engaging in wild car chases.

Reference:

[1] Kindergesundheit-info.de: Spielen – die beste Förderung für Ihr Kind
(visited 2017-10-09)

[2] DeepMind: Solve intelligence. Use it to make the world a better place.
(visited 2017-10-09)

[3] AIIDE: StarCraft AI Competition 2017, (visited 2017-10-09)

[4] DeepMind: DeepMind and Blizzard open StarCraft II as an AI research environment
(visited 2017-10-09)

[5] arXiv: Playing FPS Games with Deep Reinforcement Learning, 2016-09-18
(visited 2017-10-10)

[6] TU Darmstadt: Playing for Data: Ground Truth from Computer Game
(visited 2017-10-09)

Kommentare

Add a comment

Name*

E-mail (not published)*

Website

Please calculate 7 plus 7. *

Comment

Notify me of new comments by e-mail

SUBMIT COMMENT

SHARE:

GO BACK

AI NEWS 2018-05-28 TELEKOM DEFINES ITS OWN POLICY

AI NEWS 2018-05-01 AUTONOMOUS DRIVING: DRIVING WI...

AI NEWS 2018-04-09 KNOWLEDGE IS POWER – BUSINESS...

AI NEWS 2018-04-13 THE HISTORY OF ARTIFICIAL INTELLI...

AI NEWS 2018-04-04 THE AUTOMATION OF KNOWLEDGE...

AI NEWS MEET&GREET 2018-04-06 EVENTS

AI NEWS STEFFAN HEUER | 2018-01-23 SILICON VALLEY LOVES AI: REPORT...

AI NEWS 2018-01-22 FROM HAUTE COUTURE TO THE BAR...