



## AN AI TAUGHT ITSELF TO PLAY A VIDEO GAME – FOR THE FIRST TIME, IT’S BEATING HUMANS

A SELF-TAUGHT AI BEAT HUMANS AT THEIR OWN GAME – HERE’S HOW THEY DID IT

By IMD Professor Amit Joshi and Maude Lavanchy

**IMD**  
Chemin de Bellerive 23  
PO Box 915,  
CH-1001 Lausanne  
Switzerland

Tel: +41 21 618 01 11  
Fax: +41 21 618 07 07  
[info@imd.org](mailto:info@imd.org)  
[www.imd.org](http://www.imd.org)

Since the earliest days of virtual chess and solitaire, video games have been a playing field for developing artificial intelligence (AI). Each victory of machine against human has helped make algorithms smarter and more efficient. But in order to tackle real world problems – such as automating complex tasks including driving and negotiation – these algorithms must navigate more complex environments than board games and learn teamwork. Teaching AI how to work and interact with other players to succeed had been an insurmountable task – until now.

[In a new study](#), researchers detailed a way to train AI algorithms to reach human levels of performance in a popular 3D multiplayer game – a modified version of Quake III Arena in Capture the Flag mode.

Even though the task of this game is straightforward – two opposing teams compete to capture each other’s flags by navigating a map – winning demands complex decision-making and an ability to predict and respond to the actions of other players.

This is the first time an AI has attained human-like skills in a first-person video game. So how did the researchers do it?

### [The robot learning curve](#)

In 2019, several milestones in AI research have been reached in other multiplayer strategy games. Five “bots” – players controlled by an AI – [defeated a professional e-sports team in a game of DOTA 2](#). Professional human players were [also beaten by an AI](#) in a game of StarCraft II. In all cases, a form of [reinforcement learning](#) was applied, whereby the algorithm learns by trial and error and by interacting with its environment.

The five bots that beat humans at DOTA 2 didn’t learn from humans playing – they were trained exclusively by [playing matches against clones of themselves](#). The improvement that allowed them to defeat professional players came from [scaling existing algorithms](#). Due to the computer’s speed, the AI could play in a few seconds a game that takes minutes or even hours for humans to play. This allowed the researchers to train their AI with 45,000 years of gameplay within ten months of real-time.

The Capture the Flag bot from the recent study also began learning from scratch. But instead of playing against its identical clone, a cohort of 30 bots was [created and trained in parallel](#) with their own internal reward signal. Each bot within this population would then play together and learn from each other. As David Silver – one of the research scientists involved – notes, AI is beginning to “remove the constraints of human knowledge... and create knowledge itself”.

The learning speed for humans is still much [faster than the most advanced deep reinforcement learning algorithms](#). Both OpenAI’s bots and DeepMind’s AlphaStar (the bot playing StarCraft II) devoured thousands of years’ worth of gameplay before being able to reach a human level of performance. Such training is [estimated to cost several millions of dollars](#). Nevertheless, a self-taught AI capable of beating humans at their own game is an exciting breakthrough that could change how we see machines.

## The future of humans and machines

AI is often portrayed replacing or [complementing human capabilities](#), but rarely as a fully-fledged team member, performing the same task as human beings. As these video game experiments involve machine-human collaboration, they offer a glimpse of the future.

Human players of Capture the Flag rated the bots as more collaborative than other humans, but players of DOTA 2 had a mixed reaction to their AI teammates. Some were quite enthusiastic, saying they felt supported and that they learned from playing alongside them. [Sheever](#), a professional DOTA 2 player, spoke about her experience teaming up with bots:

It actually felt nice; [the AI teammate] gave his life for me at some point. He tried to help me, thinking 'I'm sure she knows what she's doing' and then obviously I didn't. But, you know, he believed in me. I don't get that a lot with [human] teammates.

[Others were less enthusiastic](#), but as communication is a pillar of any relationship, improving human-machine communication will be crucial in the future. Researchers have already adapted some features to make the bots more "human friendly", such as making bots [artificially wait before choosing their character during the team draft](#) before the game, to avoid pressuring the humans.

But should AI learn from us or continue to teach themselves? Self-learning without imitating humans could teach AI more efficiency and creativity, but this could create algorithms more appropriate to tasks that don't involve human collaboration, such as warehousing robots.

On the other hand, one might argue that having a machine trained from humans would be more intuitive – humans using such AI could understand why a machine did what it did. As AI gets smarter, we're all in for more surprises.

[Amit Joshi](#) is Professor of Digital Marketing and Strategy at IMD.

[Maude Lavanchy](#) is Research Fellow at IMD.

This article was first published by [The Conversation](#).