

Why artificial intelligence doesn't really exist yet

October 24 2019



Example of a learning agent in a computer game: the character is controlled by a human player. The eyes are the agents. The player is supposed to guide the agents in such a way that they execute a task, for example without running into an obstacle first. Training is based on a machine learning process; all the player does is outline rough requirements. Credit: RUB, Institut für Neuroinformatik

The processes underlying artificial intelligence today are in fact quite dumb. Researchers from Bochum are attempting to make them smarter.

Radical change, revolution, megatrend, maybe even a risk: artificial intelligence has penetrated all industrial segments and keeps the media busy. Researchers at the RUB Institute for Neural Computation have been studying it for 25 years. Their guiding principle is: in order for machines to be truly intelligent, new approaches must first render machine learning more efficient and flexible.

"There are two types of machine learning that are successful today: deep neural networks, also known as Deep Learning, as well as reinforcement learning," explains Professor Laurenz Wiskott, Chair for Theory of Neuronal Systems.

Neural networks are capable of making complex decisions. They are frequently utilized in image recognition applications. "They can, for example, tell from photos if the subject is a man or a woman," says Wiskott.

The architecture of such networks is inspired by networks of nerve cells, or neurons, in our brain. Neurons receive signals via several input channels and then decide whether they pass the signal in the form of an electrical pulse to the next neurons or not.

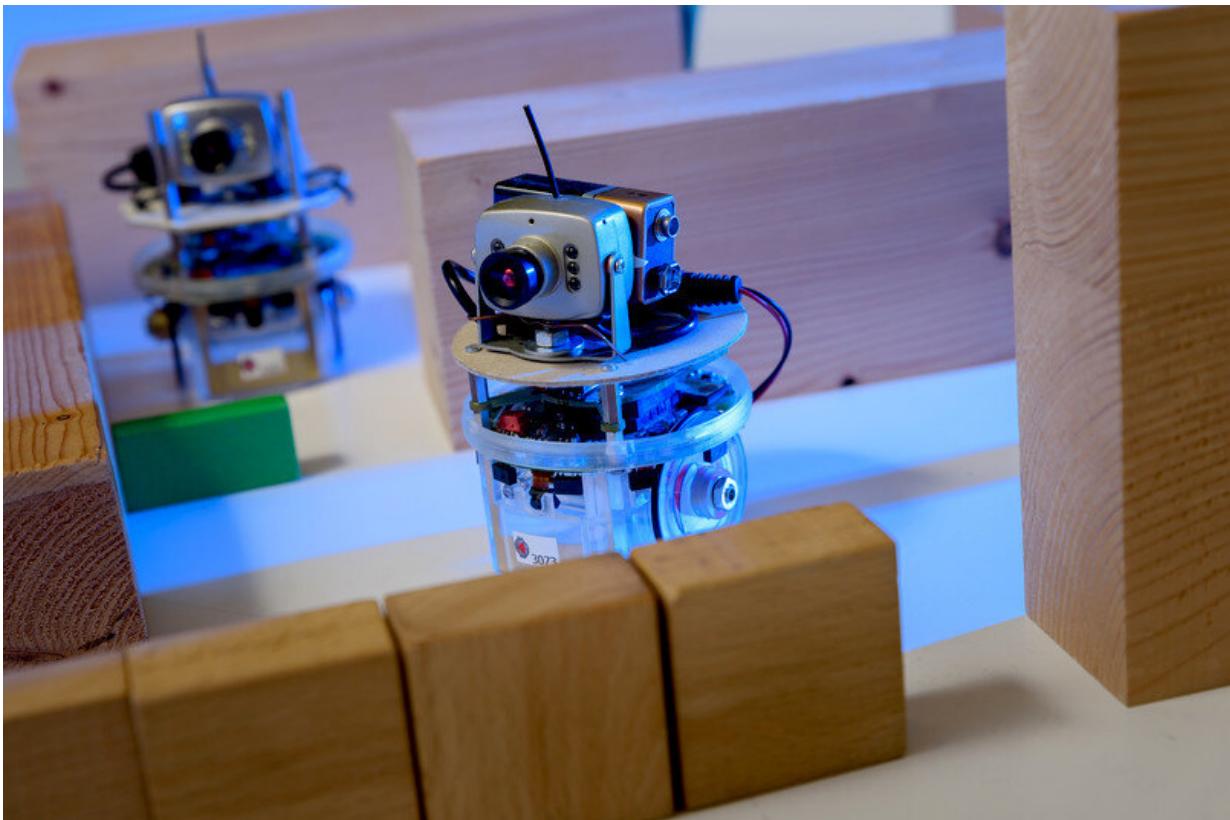
Neural networks likewise receive several input signals, for example pixels. In a first step, many artificial neurons calculate an output signal from several input signals by simply multiplying the inputs by different but constant weights and then adding them up. Each of these arithmetic operations results in a value that – to stick with the example of man/woman – contributes a little to the decision for female or male. "The outcome is slightly altered, however, by setting negative results to zero. This, too, is copied from nerve cells and is essential for the performance of neural networks," explains Laurenz Wiskott.

The same thing happens again in the next layer, until the network comes

to a decision in the final stage. The more stages there are in the process, the more powerful it is – neural networks with more than 100 stages are not uncommon. Neural networks often solve discrimination tasks better than humans.

The learning effect of such networks is based on the choice of the right weighting factors, which are initially chosen at random. "In order to train such a network, the input signals as well as what the final decision should be are specified from the outset," elaborates Laurenz Wiskott. Thus, the network is able to gradually adjust the weighting factors in order to finally make the correct decision with the greatest probability.

Reinforcement learning, on the other hand, is inspired by psychology. Here, every decision made by the algorithm – experts refer to it as the agent – is either rewarded or punished. "Imagine a grid with the agent in the middle," illustrates Laurenz Wiskott. "Its goal is to reach the top left box by the shortest possible route – but it doesn't know that." The only thing the agent wants is to get as many rewards as possible, otherwise it's clueless. At first, it will move across the board at random, and every step that does not reach the goal will be punished. Only the step towards the goal results in a reward.



Which route should the robot take? This decision is based on countless arithmetic operations. Credit: Roberto Schirdewahn

In order to learn, the agent assigns a value to each field indicating how many steps are left from that position to its goal. Initially, these values are random. The more experience the agent gains on his board, the better it can adapt these values to real life conditions. Following numerous runs, it is able to find the fastest way to its goal and, consequently, to the reward.

"The problem with these machine learning processes is that they are pretty dumb," says Laurenz Wiskott. "The underlying techniques date back to the 1980s. The only reason for their current success is that today we have more computing capacity and more data at our disposal."

Because of this, it is possible to quickly run the virtually inefficient learning processes innumerable times and feed neural networks with a plethora of images and image descriptions in order to train them.

"What we want to know is: how can we avoid all that long, nonsensical training? And above all: how can we make machine learning more flexible?" as Wiskott succinctly puts it. Artificial intelligence may be superior to humans in exactly the one task for which it was trained, but it cannot generalize or transfer its knowledge to related tasks.

This is why the researchers at the Institute for Neural Computation are focusing on new strategies that help machines to discover structures autonomously. "To this end, we deploy the principle of unsupervised learning," says Laurenz Wiskott. While [deep neural networks](#) and reinforcement learning are based on presenting the desired result or rewarding or punishing each step, the researchers leave learning algorithms largely alone with their input.

"One task could be, for example, to form clusters," explains Wiskott. For this purpose, the computer is instructed to group similar data. With regard to points in a three-dimensional space, this would mean grouping points whose coordinates are close to each other. If the distance between the coordinates is greater, they would be allocated to different groups.

"Going back to the example of pictures of people, one could look at the result after the grouping and would probably find that the computer has put together a group with pictures of men and a group with pictures of women," elaborates Laurenz Wiskott. "A major advantage is that all that's required in the beginning are photos, rather than an image caption that contains the solution to the riddle for training purposes, as it were."

The slowness principle

Moreover, this method offers more flexibility, because such cluster formation is applicable not only for pictures of people, but also for those of cars, plants, houses or other objects.

Another approach pursued by Wiskott is the slowness principle. Here, it is not photos that constitute the input signal, but moving images: if all features are extracted from a video that change very slowly, structures emerge that help to build up an abstract representation of the environment. "Here, too, the point is to prestructure input data," points out Laurenz Wiskott. Eventually, the researchers combine such approaches in a modular way with the methods of supervised learning, in order to create more flexible applications that are nevertheless very accurate.

"Increased flexibility naturally results in performance loss," admits the researcher. But in the long run, flexibility is indispensable if we want to develop robots that can handle new situations."

Provided by Ruhr-Universitaet-Bochum

Citation: Why artificial intelligence doesn't really exist yet (2019, October 24) retrieved 24 April 2024 from <https://techxplore.com/news/2019-10-artificial-intelligence-doesnt.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.