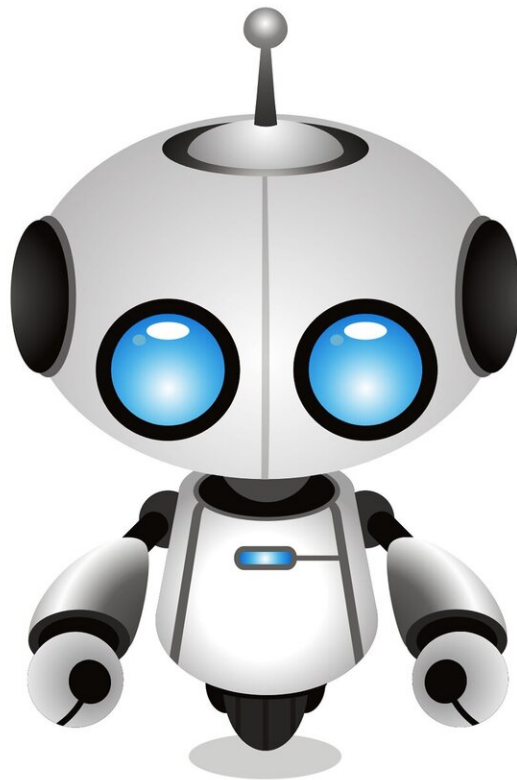# Learning from mistakes and transferable skills—the attributes for a worker robot

November 4 2019

Credit: CC0 Public Domain

Practise makes perfect—it is an adage that has helped humans become highly dexterous and now it is an approach that is being applied to robots.

Computer scientists at the University of Leeds are using the artificial intelligence (AI) techniques of automated planning and reinforcement learning to "train" a robot to find an object in a cluttered space, such as a warehouse shelf or in a fridge—and move it.

The aim is to develop robotic autonomy, so the machine can assess the unique circumstances presented in a task and find a solution—akin to a robot transferring skills and knowledge to a new problem.

The Leeds researchers are presenting their findings today (Monday, November 4) at the International Conference on Intelligent Robotics and Systems in Macau, China.

The big challenge is that in a confined area, a robotic arm may not be able to grasp an object from above. Instead it has to plan a sequence of moves to reach the target object, perhaps by manipulating other items out of the way. The computer power needed to plan such a task is so great, the robot will often pause for several minutes. And when it does execute the move, it will often fail.

Developing the idea of practise makes perfect, the computer scientists at Leeds are bringing together two ideas from AI.

One is automated planning. The robot is able to "see" the problem through a vision system, in effect an image. Software in the robot's operating system simulates the possible sequence of moves it could make to reach the target object.

But the simulations that have been "rehearsed" by the robot fail to capture the complexity of the real world and when they are implemented, the robot fails to execute the task. For example, it can knock objects off the shelf.

So the Leeds team have combined planning with another AI technique called reinforcement learning.

Reinforcement learning involves the computer in a sequence of trial and error attempts—around 10,000 in total—to reach and move objects. Through these trial and error attempts, the robot "learns" which actions it has planned are more likely to end in success.

The computer undertakes the learning itself, starting off by randomly selecting a planned move that might work. But as the robot learns from trial and error, it becomes more adept at selecting those planned moves that have a greater chance of being successful.

Dr. Matteo Leonetti, from the School of Computing, said: "Artificial intelligence is good at enabling robots to reason—for example, we have seen robots involved in games of chess with grandmasters.

"But robots aren't very good at what humans do very well: being highly mobile and dexterous. Those physical skills have been hardwired into the human brain, the result of evolution and the way we practise and practise and practise.

"And that is an idea that we are applying to the next generation of robots."

According to Wissam Bejjani, a Ph.D. student who wrote the research paper, the robot develops an ability to generalise, to apply what it has planned to a unique set of circumstances.

He said: "Our work is significant because it combines planning with reinforcement learning. A lot of research to try and develop this technology focuses on just one of those approaches.

"Our approach has been validated by results we have seen in the University's robotics lab.

"With one problem, where the robot had to move a large apple, it first went to the left side of the apple to move away the clutter, before manipulating the apple.

"It did this without the clutter falling outside the boundary of the shelf."

Dr. Mehmet Dogar, Associate Professor in the School of Computing, was also involved in the study. He said the approach had speeded up the robot's "thinking" time by a factor of ten—decisions that took 50 seconds now take 5 seconds.

Provided by University of Leeds