


DeepCube solver approach might go beyond cube into other research

June 18 2018, by Nancy Owano



An illustration of DeepCube. The training and solving process is split up into ADI and MCTS. First, we iteratively train a DNN by estimating the true value of the input states using breadth-first search. Then, using the DNN to guide exploration, we solve cubes using Monte Carlo Tree Search. Credit: arXiv:1805.07470 [cs.AI]

Unleashing ways for a machine to solve the Rubik's Cube? Numerous teams can stand up and say been there, done that. We have seen lots of headlines, too, on how they clocked in to set time records. So what's the big deal about the latest machine-solving-cube story?

David Grossman in *Popular Mechanics* remarked that the California scientists took things to the third dimension with an algorithm that can

[figure out](#) how to solve a Rubik's Cube.

A team from University of California Irvine are behind an approach that drew special attention. "Solving the Rubik's Cube Without Human Knowledge" is the title of their paper, which describes their exploration, and the paper is on arXiv.

Stephen McAleer, Forest Agostinelli, Alexander Shmakov and Pierre Baldi are the authors.

"We introduce Autodidactic Iteration: a novel reinforcement learning algorithm that is able to teach itself how to solve the Rubik's Cube with no human assistance."

Paul Lilly in *HotHardware*: Machines typically use a self-teaching method based on a rewards system. Researchers feed the [machine](#) the rules of the game, and then it uses a rewards process to determine if a move was a good one or a bad one,

However, as the authors wrote, "for many combinatorial optimization environments, rewards are sparse and episodes are not guaranteed to terminate."

They took the Autodidactic Iteration path. They said, "In order to solve the Rubik's Cube using reinforcement learning, the algorithm will learn a policy. The policy determines which move to take in any given state."

MIT Technology Review pinned down how it works. "Given an unsolved cube, the machine must decide whether a specific move is an improvement on the existing configuration. To do this, it must be able to evaluate the move. Autodidactic iteration does this by starting with the finished cube and working backwards to find a configuration that is similar to the proposed [move](#)."

The authors wrote that "DeepCube discovered a notable amount of Rubik's Cube knowledge during its training process, including the knowledge of how to use complex permutation groups and strategies similar to the best human 'speed-cubers'."

Their training machine was a 32-core Intel Xeon E5-2620 server with three NVIDIA Titan XP GPUs. They called their solver DeepCube.

Lilly's assessment: It's not a perfect solution to the problem, but is flawless in terms of accuracy.

The team stated in the paper's abstract that "Our algorithm is able to solve 100% of randomly scrambled cubes while achieving a median solve length of 30 moves —less than or equal to solvers that employ human domain knowledge."

Why this matters: it's a cube solving story and more. The team mentioned additional goals.

"Besides further work with the Rubik's Cube, we are working on extending this method to find approximate solutions to other combinatorial optimization [problems](#) such as prediction of protein tertiary structure. Many combinatorial optimization problems can be thought of as sequential decision making problems, in which case we can use reinforcement learning."

MIT Technology Review said the new approach tackled "an important problem in computer science—how to solve complex problems when help is minimal."

Ideally, said Lilly, "it could lead to finding cures for diseases, if the method is able to work as well on such things as it does with solving a Rubik's Cube."

MIT Technology Review: "The real test, of course, will be how this approach copes with more complex problems such as protein folding. We'll be watching to see how it does."

More information: Solving the Rubik's Cube Without Human Knowledge, arXiv:1805.07470 [cs.AI] arxiv.org/pdf/1805.07470.pdf

Abstract

A generally intelligent agent must be able to teach itself how to solve problems in complex domains with minimal human supervision. Recently, deep reinforcement learning algorithms combined with self-play have achieved superhuman proficiency in Go, Chess, and Shogi without human data or domain knowledge. In these environments, a reward is always received at the end of the game, however, for many combinatorial optimization environments, rewards are sparse and episodes are not guaranteed to terminate. We introduce Autodidactic Iteration: a novel reinforcement learning algorithm that is able to teach itself how to solve the Rubik's Cube with no human assistance. Our algorithm is able to solve 100% of randomly scrambled cubes while achieving a median solve length of 30 moves—less than or equal to solvers that employ human domain knowledge.

© 2018 Tech Xplore

Citation: DeepCube solver approach might go beyond cube into other research (2018, June 18) retrieved 20 March 2024 from <https://techxplore.com/news/2018-06-deepcube-solver-approach-cube.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.
