


DeepMind sees promising AI results for data center cooling system

20 August 2018, by Nancy Owano



Credit: CC0 Public Domain

Back in March, 3M ran a presentation about data, reminding us there was not going to be anything like a slowing down of data, and then posed the question, ok, so how do we swallow that and at the same time envision a sustainable future? The link: [data centers and the challenges they pose to stay up and running and more environmentally friendly.](#)

The presentation delivered a riveting thought. "More data has been created in the last two years than ever before in human [history](#)."

We're not just talking about cute cats sitting in boxes or rappers with models, so park it. We are talking about data in surgery, daily number changes that affect farmer's crops, traffic support, crisis warnings of all sorts. The [data centers](#) are crucial to our ability to communicate.

US data centers in the US consume levels of energy that are staggering to the uninitiated.

"Storing, moving, processing, and analyzing data all require energy. Lots of it. The processors in the biggest data centers hum with as much energy as can be delivered by a large power station, 1,000 megawatts or more. And it can take as much energy again to keep the servers and surrounding buildings from [overheating](#)," as *YaleEnvironment360* put it.

The 3M presentation remarked that 38 percent of the electricity need of data centers is just to cool the electronics.

Over at DeepMind, they have been working toward a cooling solution. The DeepMind blog stated, "At DeepMind and Google, we believe that if we can use AI as a tool to discover new knowledge, [solutions](#) will be easier to reach." With that mission, they have been looking at how AI might step in to manage data center cooling.

They had the AI learn how to adjust a cooling system in order to reduce power consumption. Data centers' energy consumption has gone down as a result.

Behind the success of their system was feedback from data center operators, who had something on top of their wish list: too much operator effort and supervision needed for data center cooling had them asking if anything could be done without so much manual implementation and still achieve [energy savings](#)?

The news is that Google is putting a self-taught algorithm in charge of a part of its infrastructure.

Amanda Gasparik, Google data center engineer, Chris Gamble and Jim Gao, the latter two of DeepMind, wrote about the effort in the DeepMind blog—autonomous data center cooling. Actually, the effort had its start a while back.

MIT Technology Review provided some history: "Over the past couple of years, Google has been testing an algorithm that learns how best to adjust [cooling](#) systems—fans, ventilation, and other equipment—in order to lower power consumption. This system previously made recommendations to data center managers, who would decide whether or not to implement them, leading to energy savings of around 40 percent in those cooling systems."

Those were "human implemented" recommendations.

The news is that a system that work on in 2016 is another level. Back then, the goals were fundamental, gunning for (1) energy savings and (2) a cut in CO2 emissions.

Now comes the 2018 twist: They announced "our AI system is directly controlling data centre cooling."

Now in "multiple" Google data centers. the system has been in place for only a matter of months, but the blog claimed signs that the system was already delivering energy savings of about 30 percent—and further improvements are expected.

Why do they expect further improvements? Remember, this is AI. "That's because these systems get better over time with more data," AI can walk the walk. "Rules don't get better over time, but AI does," declared Dan Fuenffinger, one of Google's data center operators, in the blog.

What did he mean by that? "AI control system is finding yet more novel ways to manage cooling that have surprised even the data center operators." Fuefinnger noted he saw the AI "learn to take advantage of winter conditions and produce colder than normal water, which reduces the energy required for cooling within the data center."

Will Knight reported on an optimistic team. "DeepMind fed its new algorithm information gathered from Google data centers and let it determine what cooling configurations would reduce energy consumption. The project could generate millions of dollars in [energy](#) savings and may help the company lower its carbon emissions, says Joe Kava, vice president of data centers for Google." Nonetheless, some might feel this is a big gamble in placing a data center, with all its mission-critical expectations, in the hands of an algorithm.

No human element at all? Yes, there is a human player. The blog authors said their data center operators "are always in control and can choose to exit AI control mode at any time. In these scenarios, the control system will transfer

seamlessly from AI control to the on-site rules and heuristics that define the automation industry today." Simply put, the human override is always available—and designed to supersede any AI actions.

How does their system work?

"Every five minutes, our cloud-based AI pulls a snapshot of the data centre [cooling](#) system from thousands of sensors and feeds it into our deep neural networks, which predict how different combinations of potential actions will affect future [energy](#) consumption. The AI system then identifies which actions will minimise the [energy consumption](#) while satisfying a robust set of safety constraints. Those actions are sent back to the data centre, where the actions are verified by the local control system and then implemented."

Their design involves AI agents and a control infrastructure for safety and reliability,

They said they use eight mechanisms to ensure the system will behave properly.

One of them is that AI can be used to estimate uncertainty. For every potential action their AI agent calculates its confidence that this is a good action. Actions with low confidence are eliminated from consideration. Then there is two-level verification, where a local control system verifies the instructions against its own set of constraints.

More information:

deepmind.com/blog/safety-first...-industrial-control/

© 2018 Tech Xplore

APA citation: DeepMind sees promising AI results for data center cooling system (2018, August 20) retrieved 23 October 2021 from <https://techxplore.com/news/2018-08-deepmind-ai-results-center-cooling.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.