*** 北米連携センター定期報告***

2014 年 5 月 28 日

●機械学習でさらなる省電力化図るグーグルのデータセンター

【ArsTechnica, 2014/05/28】

省電力性の高さで知られるグーグルのデータセンターは、ニューラルネットワークによる機械学習でその消費電力をさらに下げる試みを進めている。

同社データセンターの電力消費は「Power Usage Effectiveness (PUE)」で計測されているが、現在の平均 PUE は 1.12 で、ほぼ全ての電力がコンピューティングに消費されていることを示している。

同社のジョー・カヴァ副社長がブログで説明したところによると、データセンター・チームのエンジニアであるジム・ガオ氏が PUE を 30 秒毎に算出するために継続的に収集している IT 総負荷、外気温、冷却装置の消費電力などを PUE の予測・改善に用いることができると気付き、機械学習により PUE を 99.6%の精度で予測できるモデルを作り上げたのこと。

このモデルを活用することで、現状から消費電力を一層削減できる箇所を見つけ出すことが可能になったという。

（参考）本件報道記事

Google's machine-learning data centers make themselves more efficient

Neural networks drive Google's energy usage to even lower lows.

by Jon Brodkin - May 28 2014, 10:38am -0400

The mechanical plant at Google's data center in The Dalles, OR. Google continuously tracks performance of heat exchangers and other equipment in this image.

Google

Google's data centers are famous for their efficient use of power, and now they're (literally) getting even smarter about how they consume electricity. Google today explained how it uses neural networks, a form of machine learning, to drive energy usage in its data centers to new lows.

Google measures data center electricity usage by PUE (power usage effectiveness). A PUE of 2.0 would mean that "for every watt of IT power, an additional watt is consumed to cool and distribute power to the IT equipment," Google explains. Google's PUE across all of its data centers is an average of 1.12,

meaning nearly all of its energy is used for computing rather than overhead. One Google employee figured out how to get the number even lower, Google's blog explained.

"It all started as a 20 percent project, a Google tradition of carving out time for work that falls outside of one's official job description," Google data center VP Joe Kava wrote. "Jim Gao, an engineer on our data center team, is well-acquainted with the operational data we gather daily in the course of running our data centers. We calculate PUE, a measure of energy efficiency, every 30 seconds, and we're constantly tracking things like total IT load (the amount of energy our servers and networking equipment are using at any time), outside air temperature (which affects how our cooling towers work) and the levels at which we set our mechanical and cooling equipment. Being a smart guy—our affectionate nickname for him is 'Boy Genius'—Jim realized that we could be doing more with this data. He studied up on machine learning and started building models to predict—and improve—data center performance."

Gao's models are 99.6 percent accurate in predicting PUE. "This means he can use the models to come up with new ways to squeeze more efficiency out of our operations," Kava wrote. "For example, a couple months ago we had to take some servers offline for a few days—which would normally make that data center less energy efficient. But we were able to use Jim's models to change our cooling setup temporarily—reducing the impact of the change on our PUE for that time period."

In that case, an upgrade to electrical infrastructure required Google to re-route 40 percent of server traffic at a facility, but "through a combination of PUE simulations and local expertise, the team selected a new set of operational parameters that reduced the PUE by ~0.02 compared to the previous configuration," Gao wrote in a white paper published today.

Gao explained that "neural networks are a class of machine learning algorithms that mimic cognitive behavior via interactions between artificial neurons. They are advantageous for modeling intricate systems because neural networks do not require the user to predefine the feature interactions in the model, which assumes relationships within the data. Instead, the neural network searches for patterns and interactions between features to automatically generate a best¬fit model."

The neural network studies a variety of factors including total server load in kilowatts; the numbers of water pumps, cooling towers, chillers, dry coolers, and chilled water injection pumps in operation; cooling tower water temperature;

bulb temperatures; and outdoor humidity, wind speed, and wind direction. In one case, Google saved electricity by increasing the water supply temperature by 3 degrees fahrenheit. In another, the system detected erroneous readings from natural gas meters.
"Actual testing on Google DCs [data centers] indicate that machine learning is an effective method of using existing sensor data to model DC energy efficiency, and can yield significant cost savings," Gao wrote. "Model applications include DC simulation to evaluate new plant configurations, assessing energy efficiency performance, and identifying optimization opportunities."

Source:
http://arstechnica.com/information-technology/2014/05/googles-machine-learning-data-centers-make-themselves-more-efficient/

以　上