

A gnu way to control room temperature

5 December 2019, by Anna Boyle



A wild gnu, also known as a wildebeest.

Heating, ventilation, and air conditioning systems—called HVAC systems—can be a delicate balance. There are many factors to consider, from air flow between rooms to the effect of human body heat. In the past decade, researchers have turned to machine learning to optimize these systems. With smarter controllers, buildings can save on energy without sacrificing comfort.

There are currently two main approaches to the problem. In the first approach, the controller uses a detailed model of the [building](#) to manage its systems. However, the model takes a lot of effort to create. "A very good model of a building is hard to make, hard to maintain, and doesn't scale," says Mario Bergés, a professor of civil and environmental engineering. "Buildings are not all the same, so you have to create a model for each building."

The other approach involves generating vast amounts of data, which allows the controller to adapt to different building systems. In this case, the main obstacle is how long it takes. "You would need about 40 years of simulation data for a relatively complex building," Bergés says. "In the real world, you can't just spend 40 years trying to

figure out how to control a building."

To tackle these challenges, Bergés worked with Ph.D. student Bingqing Chen and a Dell collaborator. They developed a new solution, Gnu-RL, that incorporates the best of both approaches.

First, Gnu-RL completes offline pretraining using historical data. HVAC systems already have controls, so Gnu-RL learns to copy them. In this way, it avoids the complications of precise models and large amounts of data. "It only needs historical data, which we already have a lot of," Chen says.

Once the pretraining is completed, Gnu-RL can imitate the previous controller reliably. Next, it's taught to adapt and become better. Bergés and Chen applied a recently developed differentiable Model Predictive Control (MPC) policy. This policy rewards the agent for maximizing reward and minimizing cost, and the agent adjusts accordingly until it achieves the optimal controls for the HVAC system. This method is called [reinforcement learning](#)—which is why the solution has RL at the end of its name.

The first part of the name, on the other hand, comes from a more unconventional source. A gnu is a large, dark antelope from Africa. These animals are incredibly precocial, which means they are born in a relatively advanced state. "They can run away from predators within the same day that they're born," Chen says. "And Gnu-RL controls reasonably well at onset." This similarity made the name a natural choice.

Bergés and Chen back up this comparison with two tests. The first test was performed with a simulation of the intelligent workplace at the top of Margaret Morrison. "We had a 40-years to four-week improvement in terms of the training time," Bergés says. "And we also showed about a 6% improvement in [energy savings](#) without sacrificing comfort."

Bergés and Chen were so encouraged by the

simulation results that they decided to apply Gnu-RL to a real-world setting. For three weeks, they let Gnu-RL control the air flow of a conference room in Gates Center. The results of this test were equally promising. "It learned how to imitate the existing controller," Bergés says. "Then, in addition to that, it learned to pre-cool the space and provide comfort before people would arrive, which is something that the existing controller wasn't doing."

However, while their work is exciting, Bergés and Chen want to acknowledge the work of the researchers who came before them. "Our contribution is an application, so we're building upon others' work," Chen says. Most notably, Gnu-RL adopted the differentiable MPC policy developed by Brandon Amos and Zico Kolter. This policy allowed Gnu-RL to be both efficient and flexible.

Bergés and Chen presented their paper on Gnu-RL at the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys 2019). The conference took place in New York City on November 13 and 14.

Looking to the future, Bergés and Chen believe there's still room for Gnu-RL to grow. "We've been looking at relatively simple scenarios," Bergés says. "There may be complications as we try to control much more complex buildings, so that's still an open question. But at least we're pointing it in a direction that is new and that may spur a lot of research for how to address this problem."

More information: Bingqing Chen et al. Gnu-RL, *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation - BuildSys '19* (2019). [DOI: 10.1145/3360322.3360849](https://doi.org/10.1145/3360322.3360849)

Provided by Carnegie Mellon University,
Department of Civil and Environmental Engineering
APA citation: A gnu way to control room temperature (2019, December 5) retrieved 26 January 2022
from <https://techxplore.com/news/2019-12-gnu-room-temperature.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.