

Machine learning speeds up vehicle routing

December 11 2021, by Becky Ham



Credit: Unsplash/CC0 Public Domain

Waiting for a holiday package to be delivered? There's a tricky math problem that needs to be solved before the delivery truck pulls up to your door, and MIT researchers have a strategy that could speed up the solution.

The approach applies to vehicle routing problems such as last-mile delivery, where the goal is to deliver goods from a central depot to multiple cities while keeping travel costs down. While there are algorithms designed to solve this problem for a few hundred cities, these solutions become too slow when applied to a larger set of cities.

To remedy this, Cathy Wu, the Gilbert W. Winslow Career Development Assistant Professor in Civil and Environmental Engineering and the Institute for Data, Systems, and Society, and her students have come up with a machine-learning strategy that accelerates some of the strongest algorithmic solvers by 10 to 100 times.

The solver algorithms work by breaking up the problem of delivery into smaller subproblems to solve—say, 200 subproblems for routing vehicles between 2,000 cities. Wu and her colleagues augment this process with a new machine-learning algorithm that identifies the most useful subproblems to solve, instead of solving all the subproblems, to increase the quality of the solution while using orders of magnitude less compute.

Their approach, which they call "learning-to-delegate," can be used across a variety of solvers and a variety of similar problems, including scheduling and pathfinding for warehouse robots, the researchers say.

The work pushes the boundaries on rapidly solving large-scale vehicle routing problems, says Marc Kuo, founder and CEO of Routific, a smart logistics platform for optimizing delivery routes. Some of Routific's recent algorithmic advances were inspired by Wu's work, he notes.

"Most of the academic body of research tends to focus on specialized algorithms for small problems, trying to find better solutions at the cost of processing times. But in the real-world, businesses don't care about finding better solutions, especially if they take too long for compute,"

Kuo explains. "In the world of last-mile logistics, time is money, and you cannot have your entire warehouse operations wait for a slow algorithm to return the routes. An [algorithm](#) needs to be hyper-fast for it to be practical."

Wu, social and engineering systems doctoral student Sirui Li, and electrical engineering and computer science doctoral student Zhongxia Yan presented their research this week at the 2021 NeurIPS conference.

Selecting good problems

Vehicle routing problems are a class of combinatorial problems, which involve using heuristic algorithms to find "good-enough solutions" to the problem. It's typically not possible to come up with the one "best" answer to these problems, because the number of possible solutions is far too huge.

"The name of the game for these types of problems is to design efficient algorithms ... that are optimal within some factor," Wu explains. "But the goal is not to find optimal solutions. That's too hard. Rather, we want to find as good of solutions as possible. Even a 0.5% improvement in solutions can translate to a huge revenue increase for a company."

Over the past several decades, researchers have developed a variety of heuristics to yield quick solutions to combinatorial problems. They usually do this by starting with a poor but valid initial solution and then gradually improving the solution—by trying small tweaks to improve the routing between nearby cities, for example. For a large problem like a 2,000-plus city routing challenge, however, this approach just takes too much time.

More recently, machine-learning methods have been developed to solve the problem, but while faster, they tend to be more inaccurate, even at

the scale of a few dozen cities. Wu and her colleagues decided to see if there was a beneficial way to combine the two methods to find speedy but high-quality solutions.

"For us, this is where machine learning comes in," Wu says. "Can we predict which of these subproblems, that if we were to solve them, would lead to more improvement in the solution, saving computing time and expense?"

Traditionally, a large-scale vehicle routing problem heuristic might choose the subproblems to solve in which order either randomly or by applying yet another carefully devised heuristic. In this case, the MIT researchers ran sets of subproblems through a [neural network](#) they created to automatically find the subproblems that, when solved, would lead to the greatest gain in quality of the solutions. This process sped up subproblem selection process by 1.5 to 2 times, Wu and colleagues found.

"We don't know why these subproblems are better than other subproblems," Wu notes. "It's actually an interesting line of future work. If we did have some insights here, these could lead to designing even better algorithms."

Surprising speed-up

Wu and colleagues were surprised by how well the approach worked. In machine learning, the idea of garbage-in, garbage-out applies—that is, the quality of a machine-learning approach relies heavily on the quality of the data. A combinatorial problem is so difficult that even its subproblems can't be optimally solved. A neural network trained on the "medium-quality" subproblem solutions available as the input data "would typically give medium-quality results," says Wu. In this case, however, the researchers were able to leverage the medium-quality

solutions to achieve high-quality results, significantly faster than state-of-the-art methods.

For vehicle routing and similar problems, users often must design very specialized algorithms to solve their specific problem. Some of these heuristics have been in development for decades.

The learning-to-delegate method offers an automatic way to accelerate these heuristics for large problems, no matter what the heuristic or—potentially—what the problem.

Since the method can work with a variety of solvers, it may be useful for a variety of resource allocation problems, says Wu. "We may unlock new applications that now will be possible because the cost of solving the problem is 10 to 100 times less."

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Machine learning speeds up vehicle routing (2021, December 11) retrieved 19 April 2024 from <https://techxplore.com/news/2021-12-machine-vehicle-routing.html>

<p>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.</p>
--