

Improved approach to the 'traveling salesperson problem' could improve logistics and transport sectors

April 26 2022



Credit: Unsplash/CC0 Public Domain

A new approach to solving the traveling salesperson problem—one of the most difficult questions in computer science—significantly

outperforms current approaches.

A notorious theoretical question that has puzzled researchers for 90 years, the traveling salesperson problem also has real relevance to industry today. Essentially a question about how best to combine a set of tasks so that they can be performed in the fastest and most efficient way, finding good solutions to the problem can greatly help improve sectors such as transport and logistics.

Researchers from the University of Cambridge have developed a hybrid, data-driven approach to the problem that not only produces high-quality solutions, but at a faster rate than other state-of-the-art approaches. Their results are presented this week at the International Conference on Learning Representations.

"The importance of global logistics system was brought home to us during the pandemic," said Dr. Amanda Prorok from Cambridge's Department of Computer Science and Technology, who led the research. "We're highly reliant on this kind of infrastructure to be more efficient—and our [solution](#) could help with that as it targets both in-warehouse logistics, such as the routing of robots around a warehouse to collect goods for delivery, and those outside it, such as the routing of goods to people."

The traveling salesperson problem involves a notional delivery driver who must call at a set number of cities—say, 20, 50 or 100—that are connected by highways all in one journey. The challenge is to find the shortest possible route that calls at each destination once and to find it quickly.

"There are two key components to the problem. We want to order the stops, and we also want to know the cost, in time or distance, of going from one stop to another in that order," said Prorok.

Twenty years ago the route from the warehouse to the destinations might have been fixed in advance. But with today's availability of real-time traffic information, and the ability to send messages to the driver to add or remove delivery locations on the fly, the route may now change during the journey. But minimizing its length or duration still remains key.

There's often a cost attributed to waiting for an optimal solution or hard deadlines at which decisions must be taken. For example, the driver cannot wait for a new solution to be computed—they may miss their deliveries, or the traffic conditions may change again.

And that is why there is a need for general, anytime combinatorial optimization algorithms that produce high-quality solutions under restricted computation time.

The Cambridge-developed hybrid approach does this by combining a [machine learning model](#) that provides information about what the previous best routes have been, and a "metaheuristic" tool that uses this information to assemble the new route.

"We want to find the good solutions faster," said Ben Hudson, the paper's first author. "If I'm a driver for a courier firm I have to decide what my next destination is going to be as I'm driving. I can't afford to wait for a better solution. So that's why in our research we focused on the trade-off between the computational time needed and the quality of the solution we got."

To do this, Hudson came up with a Guided Local Search algorithm that could differentiate routes from one city to another that would be costly—in time or distance—from routes that would be less costly to include in the journey. This enabled the researchers to identify high-quality, rather than optimal, solutions quickly.

They did this by using a measure of what they call the "global regret"—the cost of enforcing one decision relative to the cost of an optimal solution—of each city-to-city route in the Guided Local Search algorithm. They used machine learning to come up with an approximation of this "regret."

"We already know the correct solution to a set of these problems," said Hudson. "So we used some machine learning techniques to try and learn from those solutions. Based on that, we try to learn for a new problem—for a new set of cities in different locations—which paths between the cities are promising.

"When we have this information, it then feeds into the next part of the algorithm—the part that actually draws the routes. It uses that extra information about what the good paths may be to build a good solution much more quickly than it could have done otherwise."

The results they came up with were impressive. Their experiments demonstrated that the hybrid, data-driven approach converges to optimal solutions at a faster rate than three recent learning-based approaches for the traveling salesperson problem.

In particular, when trying to solve the problem when it had a 100-city route, the Cambridge method reduced the mean optimality gap from 1.534% to 0.705%, a two-fold improvement. When generalizing from the 20-city problem route to the 100-city problem [route](#), the method reduced the optimality gap from 18.845% to 2.622%, a seven-fold improvement.

"A lot of logistics companies are using routing methods in real life," said Hudson. "Our goal with this research is to improve such methods so that they produce better solutions—solutions that result in lower distances being traveled and therefore lower carbon emissions and reduced impact

on the environment."

More information: Graph Neural Network Guided Local Search for the Traveling Salesperson Problem.

openreview.net/forum?id=ar92oEosBIg

Provided by University of Cambridge

Citation: Improved approach to the 'traveling salesperson problem' could improve logistics and transport sectors (2022, April 26) retrieved 20 April 2024 from

<https://techxplore.com/news/2022-04-approach-salesperson-problem-logistics-sectors.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.