SEMAPHORE

Newsletter of the Maritime Law

Association of Australia and New Zealand

Innovative Solution to “Travelling Salesperson Problem”

A University of Cambridge-led research team has developed a modern solution to the “travelling salesperson problem”, that is being heralded as offering significant efficiencies to operations in the transport and logistics sectors.

Authored by Benjamin Hudson et al, the 20-page “Graph Neural Network Guided Local Search for the Traveling Salesperson Problem” paper was recently presented at the International Conference on Learning Representations.

The four-person research team’s approach to the “problem” – determining, as fast as possible, the shortest-possible route for a notional delivery driver to call at a set number of destinations – entailed developing a hybrid, data-driven solution.

They combined a machine-learning model – which provides information about what the previous best routes have been – and a “metaheuristic” tool – that uses this information to assemble the new route – explains the paper’s first author, Ben Hudson.

“We want to find the good solutions faster,” he says.

“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, Mr Hudson says a “guided local search algorithm” was developed to differentiate routes from one city to another based on their “costliness” in either time or distance. This enabled the researchers to quickly identify “high-quality”, rather than “optimal”, solutions.

The team also incorporated a measure of what they termed the “global regret” – the cost of enforcing one decision relative to the cost of an optimal solution for each city-to-city route. Machine learning then provided an approximation of this “regret”.

“We already know the correct solution to a set of these problems. 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 team’s experiments have reportedly demonstrated that their hybrid, data-driven approach delivers optimal solutions at a faster rate than three recent learning-based approaches put forward for the problem.

In particular, the researchers report their method:

- reduces the mean optimality gap from 1.534% to 0.705% when trying to solve the problem when it had a 100-city route – a two-fold improvement

- reduces the optimality gap from 18.845% to 2.622% when generalising from the 20-city to the 100-city problem route – a seven-fold improvement

“A lot of logistics companies are using routing methods in real life. Our goal with this research is to improve such methods so that they produce better solutions – solutions that result in lower distances being travelled and therefore lower carbon emissions and reduced impact on the environment.”

Team lead researcher, Dr Amanda Prorok, says the importance of a functioning global logistics system has been fully “brought home” during the pandemic.

“We’re highly reliant on this kind of infrastructure to be more efficient,” says Dr Prorok.

“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.”

June 2022