Integrating learned preferences from historical solutions into vehicle routing

Dissertation submitted in fulfilment of the requirements for the award of the degree of

Doctor of Business Economics

by

Rocsildes Canoy

Supervisors:

Prof. dr. Tias Guns

Prof. dr. Yves Molenbruch

Vrije Universiteit Brussel—December 2022

Abstract

Among the processes involved in transportation logistics, the last-mile distribution of goods to the final recipients is often considered to be the costliest and the most complex. Various contributing factors make the process especially challenging, including strict and evolving customer demands, the unpredictability of weather and traffic conditions, and the rapid rise of e-commerce sales and urbanization, among others.

One way to approach the last-mile problem is by route optimization. There exists a rich body of literature that deals with real-life, practical versions of traveling salesman and vehicle routing problems. The proposed methodologies aim to find optimal routes for a given set of delivery vehicles and an objective function to optimize, subject to real-world limitations and constraints.

Nowadays, route scheduling and optimization software have become a common tool in the industry, mainly due to the immense progress in technology and research. Despite the technological advancement, however, users of optimization software are still known to exert a substantial amount of effort in finding satisfactory routing solutions. Just the task of modeling the problem itself can already become tedious, with numerous objectives and constraints to consider. Also, it is difficult and nearly impossible to incorporate all of the route planners' pref-

erences and the drivers' local knowledge in the optimization process. On top of that, with their own knowledge and personal preferences, drivers are known to oftentimes deviate from the routes that were originally planned for them.

The goal of this research is to use past observations of human experts to create routes that are as close as possible to the solutions that the experts would produce. The idea is to exploit and adapt techniques from machine learning to learn implicit preferences from previous solutions. Indeed, past routes contain expert preferences, as they are essentially software solutions that have been conscientiously modified by the route planners.

For learning, we propose a Markovian approach that models preferences as transition probabilities between stops/locations. To effectively capture preferences that evolve over time, we introduce several weighing techniques that assign varying weights to the historical instances. We likewise develop higher order Markov models to capture higher order preferences. We demonstrate how to best combine the learned preferences and some other objective criterion (e.g., distances) using structured output prediction. Furthermore, we endeavour to enhance our approach by introducing a neural network architecture that allows the model to consider other contextual information.

We conclude this dissertation with a review of our main contributions, recommendations on possible topics for future studies, and discussions on how our proposed approaches can be applied in the industry and in other fields of research.