Continuous approximation models with temporal constraints and objectives

John Gunnar Carlsson
University of Southern California
jcarlsso@usc.edu

Project Objective

The purpose of this project was to discover new continuous approximation models for modern logistical problems in which time plays a significant role, with a specific focus on last-mile delivery. Famous examples of such problems include the *vehicle routing problem with time windows* (VRPTW) and the *cumulative travelling salesperson problem* (CTSP). The continuous approximation paradigm is a quantitative method for solving logistics problems in which one uses a small set of parameters to model a complex system, which results in simple algebraic equations that are easier to manage than (for example) large-scale optimization models. As a further benefit, one often obtains insights from these simpler formulations that help to determine what affects the outcome most significantly. Although continuous approximation models have been used for over 60 years in logistics systems analysis, there has been very little research conducted on their use to problems with temporal features such as those described above. To the best of our knowledge, this project was the first of its kind to incorporate these temporal features into the continuous approximation paradigm.

Problem Statement

In completing this project, the team designed simple and concise mathematical models for predicting trade-offs that arise in logistical problems with time constraints and objectives. Examples of these trade-offs include the relationships between time to completion of service, average or worst-case customer satisfaction, vehicle miles travelled (VMT), or greenhouse gas (GHG) emissions. Traditionally, these problems have been solved in a discrete setting, involving fixed sets of (for example) demand points, time periods, and service facility locations; one then solves them with an integer mathematical programming solver such as CPLEX or Gurobi. A drawback of this approach is that the problems are almost always *NP-hard*, and hence solving large-scale instances would require enormous computational efforts which likely increase exponentially with the problem instance size. A further drawback is that such models are often extremely complex, which hinders understanding of salient problem features and managerial insights. For these reasons, this project used tools from geospatial optimization, computational geometry, and geometric probability theory to discover simple *continuous approximation models* that identify the key problem attributes that affect them most significantly. A continuous approximation model is characterized by its use of continuous representations of input data and decision variables as density functions over time and space, and the goal is to approximate the objective function into an expression that can be optimized by relatively simple analytical operations. Such an approximation enables transforming otherwise high-dimensional decision variables into a low-dimensional space, allowing the optimal solution to be obtained with mere calculus, even when significant operational complexities are present. The results from such models often bear closed-form analytical structures that help reveal managerial insights.

Research Methodology

This project focused on studying the *cumulative traveling salesperson problem* (CTSP), which focuses on minimizing the total waiting time of all customers to be visited, rather than finding the shortest total route, which is the objective of the traditional traveling salesperson problem. In this variant, the goal is to minimize the sum of the arrival times at each customer, prioritizing quicker service to all customers over the traditional objective of minimizing the total distance or travel time of the tour. This shift in objective from the traditional Traveling Salesman Problem (TSP) makes the CTSP particularly relevant for scenarios where the speed of service delivery is crucial, such as in logistics for perishable goods. Our approach consisted of three phases: an upper bounding argument, a lower bounding argument, and extensions to further problem types:

- Initially, the team developed a method to determine upper bounds for the CTSP, employing a routing strategy that prioritizes visiting areas with higher customer densities before those with fewer customers. This approach, referred to as the "most dense to least dense" strategy, aims to minimize the total waiting time across all customers. By formalizing this strategy, we established upper bounds that serve as a benchmark to evaluate the efficiency of CTSP solutions.
- Next, we focused on establishing lower bounds for the CTSP, using probabilistic modeling and asymptotic analysis to understand how the problem's behavior changes as the number of customers increases. This phase was essential to ensure the feasibility of the proposed solutions.
- After addressing the CTSP, we extended our analysis to the Cumulative Capacitated Vehicle Routing Problem (CCVRP). This problem variant incorporates not only the CTSP's temporal objectives but also the capacity constraints of the vehicles. Our study of the CCVRP built upon our work on the CTSP, addressing the additional complexity introduced by vehicle capacities.

Results

The conclusions of our project, based on theoretical analysis as well as empirical validation, give a fast and simple formula for predicting the total cost of a CTSP tour or its variant, the CCVRP. Our main result is the following theorem:

Let $X_1, \dots, X_n$ be independent samples drawn from a probability density f with compact support. Let $L(X_1, \dots, X_n)$ denote the cost of the minimal CTSP tour through all points. We have

$$0.2935 < \liminf_{n\to\infty} \frac{L(X_1,\dots,X_n)}{n^{3/2}\int_{\mathcal{R}} \sqrt{f(x)}P(x)\,\mathrm{d}x} \leq \limsup_{n\to\infty} \frac{L(X_1,\dots,X_n)}{n^{3/2}\int_{\mathcal{R}} \sqrt{f(x)}P(x)\,\mathrm{d}x} < 0.92117$$

Where the function *P* is defined as

$$P(x) = \Pr(f(X) \leq f(x)) = \int_{x':f(x')\leq f(x)} f(x')\,\mathrm{d}x' .$$

In a nutshell, this theorem allows us to assert a few key points: First, we observe that the CTSP scales proportionally to $n^{3/2}$, revealing a diseconomy of scale as the number of customers increases. This indicates that as the problem size grows, the cumulative waiting time for customers increases at a rate faster than linear, underscoring the importance of efficient routing strategies to mitigate this effect. Second, we establish that the "most dense to least dense" routing rule described in our report is essentially optimal for minimizing the total waiting time in CTSP scenarios. This means that prioritizing areas with higher customer densities before moving to less dense areas is not just a heuristic but aligns closely with the optimal solution strategy. Finally, our proof techniques unveil an important distinction between solutions that are optimal for the traditional Traveling Salesman Problem (TSP) and those suitable for the CTSP. We demonstrate that solutions deemed system optimal for the regular TSP, based solely on minimizing the total travel distance or time, may not necessarily provide the best outcomes for the CTSP. This distinction arises because the CTSP introduces an additional layer of complexity by prioritizing the reduction of cumulative waiting time, which is not a consideration in the traditional TSP.