Socially Optimal Personalized Routing with Preference Learning

PI: Phebe Vayanos
Co-PI: Maged M. Dessouky
Daniel J. Epstein Department of Industrial and Systems Engineering
University of Southern California, Los Angeles, CA 90089-0193
Email: phebe.vayanos@usc.edu, maged@usc.edu

Project Objective

Spurred by rapid population growth and city development, traffic congestion has become inescapable in metropolitan areas across the U.S. and its direct and indirect effects can be dire. With support lacking for taxes to fund expansion of the existing network, it is imperative to find novel ways to improve efficiency of the existing infrastructure. A major obstacle is the inability to enforce socially optimal routes upon the commuters. The goal of this research project is to improve routing efficiency by leveraging heterogeneity in commuter preferences to help bridge the gap between the de facto user equilibrium solution and the utopic socially optimal solution.

Problem Statement

We consider the problem of assigning drivers to routes to minimize overall congestion in a traffic network. The vast majority of the literature on the topic either: *a)* implicitly assumes that drivers will follow recommended routes, yielding a so-called *socially optimal solution*; or, *b)* accepts that users are self-interested and offers users a *user equilibrium solution* which optimizes their personal preferences. Both of these approaches suffer from considerable drawbacks. On the one hand, the socially optimal solution is utopic since it is not possible to realize in practice in the absence of incentives due to users being unwilling to employ routes that do not meet their personal preferences. On the other hand, the user equilibrium solution, being user- rather than system-centric, does not minimize overall congestion. In this project, we aim to address the following research question:

How can we bridge the gap between the (utopic) socially optimal solution and the (de facto) user optimal solution? Can we formulate and solve a problem where the trade-off between social and user optimality can be controlled by means of a single design parameter?

Research Methodology

We consider a traffic network in which nodes and links represent intersections and road segments, respectively. Contrary to the vast majority of the literature, we associate with each road and intersection their intrinsic characteristics such as the travel time distribution for the segment, the expected wait time at the intersection, the expected number of accidents, etc. Then, for any given route between a given origin and a given destination, we are able to compute a set of characteristics for that route, e.g., length, average travel time, minimum/maximum travel times. These characteristics are deemed useful to help users choose between all possible routes from an origin to a destination.

First, given the aforementioned traffic network, we propose a method for learning user preferences over routes. Given a finite set of origin-destination pairs with two possible routes connecting them and with each route having different characteristics, we ask users to select which route they prefer or to state that they are indifferent between the two. Subsequently, we cluster drivers based on their answers to the questions and learn, for each cluster, a single (linear) utility function which represents the preferences of all drivers assigned to that cluster. To learn the utility function for each cluster, we formulate a mixed-integer linear program which seeks to determine the utility vector that maximizes the number of drivers in the cluster whose preferences we predict correctly (in the training data). We also propose a formulation in which the weights for the different types of prediction errors are different and tailored to the application at hand.

Second, having clustered drivers based on their preferences and having learned user preferences for each cluster, we propose a novel traffic assignment problem that bridges the gap between the utopic system optimal solution and the de facto user optimal solution. Specifically, we build upon the socially optimal traffic assignment problem and augment it with feasibility constraints which ensure that users are likely to accept the recommended routes. In particular, these constraints require a single input parameter γ and stipulate that the system can only recommend routes for which the utility derived by each user in the system is within γ of the utility derived from their preferred route. The number of clusters employed controls the quality of the learned utilities and the complexity of the personalized socially optimal traffic assignment problem. The parameter γ controls the trade-off between the level of personalization and social optimality: the higher (lower) γ, the more are socially (user) optimal routes are favored with the understanding.

Results

We evaluate our utility learning approach on data collected from Amazon Mechanical Turk and using historical traffic data from the L.A. Metro area. We compare with the Logistic Regression method. Figure 1 illustrates the average accuracy rates using our approach and the Logistic Regression method as the number of clusters is varied from 1 to 250 (total number of users). From the Figure, we observe that our approach performs better out of sample even for small numbers of clusters. Figure 2 shows that using different penalty values for different types of incorrect/correct predictions influences the number of correct predictions.

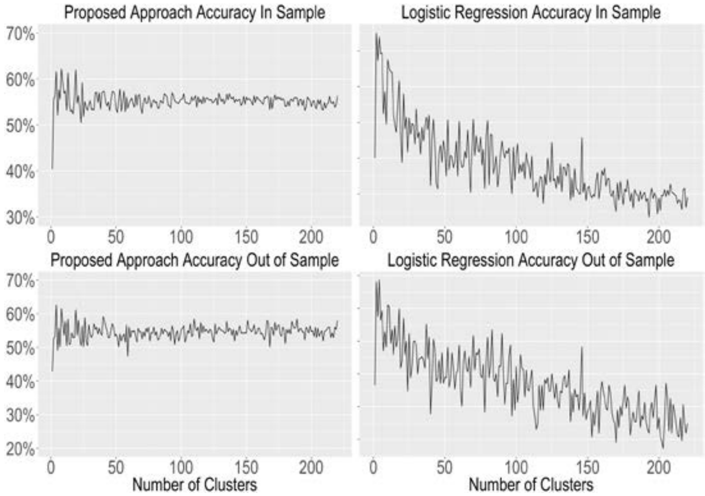

Figure 1

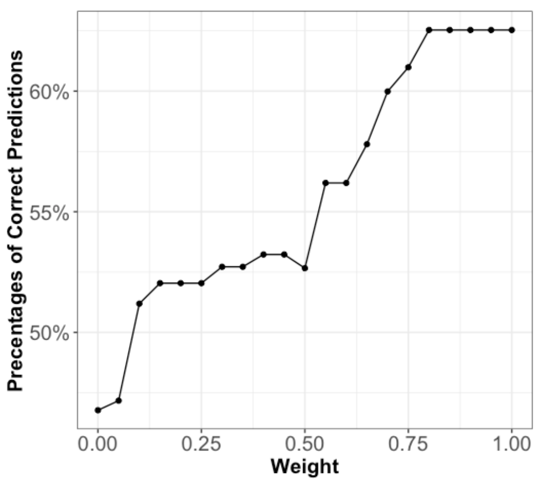

Figure 2