Transportation Research Part C 75 (2017) 30–44

Contents lists available at ScienceDirect

Transportation Research Part C

journal homepage: www.elsevier.com/locate/trc

Predicting travel time reliability using mobile phone GPS data

Dawn Woodard [a,*,1], Galina Nogin [a,2], Paul Koch [b], David Racz [c,3], Moises Goldszmidt [d,4], Eric Horvitz [b]

[a] Cornell University, Ithaca, NY, United States
[b] Microsoft Research, Redmond, WA, United States
[c] Microsoft Corp., Palo Alto, CA, United States
[d] Microsoft Research, Mountain View, CA, United States

ARTICLE INFO

Article history:
Received 23 March 2015
Received in revised form 6 September 2016
Accepted 24 October 2016

Keywords:
Location data
Traffic
Travel time
Forecasting
Statistics

ABSTRACT

Estimates of road speeds have become commonplace and central to route planning, but few systems in production provide information about the reliability of the prediction. Probabilistic forecasts of travel time capture reliability and can be used for risk-averse routing, for reporting travel time reliability to a user, or as a component of fleet vehicle decision-support systems. Many of these uses (such as those for mapping services like Bing or Google Maps) require predictions for routes in the road network, at arbitrary times; the highest-volume source of data for this purpose is GPS data from mobile phones. We introduce a method (TRIP) to predict the probability distribution of travel time on an arbitrary route in a road network at an arbitrary time, using GPS data from mobile phones or other probe vehicles. TRIP captures weekly cycles in congestion levels, gives informed predictions for parts of the road network with little data, and is computationally efficient, even for very large road networks and datasets. We apply TRIP to predict travel time on the road network of the Seattle metropolitan region, based on large volumes of GPS data from Windows phones. TRIP provides improved interval predictions (forecast ranges for travel time) relative to Microsoft's engine for travel time prediction as used in Bing Maps. It also provides deterministic predictions that are as accurate as Bing Maps predictions, despite using fewer explanatory variables, and differing from the observed travel times by only 10.1% on average over 35,190 test trips. To our knowledge TRIP is the first method to provide accurate predictions of travel time reliability for complete, large-scale road networks.

1. Introduction

Several mapping services provide predictions of the expected travel time on an arbitrary route in a road network, in real time and using traffic, time of day, day of the week, and other information. They use these predictions to recommend a route or routes with minimum expected travel time. Microsoft's mapping service (Bing Maps) predicts travel time for large-scale

* Corresponding author.
E-mail address: moises.goldszmidt@gmail.com (M. Goldszmidt).
URL: http://people.orie.cornell.edu/woodard (D. Woodard).
[1] Research performed while Visiting Researcher at Microsoft Research.
[2] Research performed while intern at Microsoft Research.
[3] Current address: Google, Mountain View, CA, United States.
[4] Current address: Apple, Cupertino, CA, United States.

http://dx.doi.org/10.1016/j.trc.2016.10.011
0968-090X/

road networks around the world using a method called Clearflow (Microsoft Research, 2012), which employs probabilistic graphical models learned from data to predict flows on arbitrary road segments. The method, which has its roots in the earlier Smartphlow effort on forecasting highway flows and reliability (Horvitz et al., 2005), considers evidence about real-time traffic conditions, road classifications, topology of the road network, speed limits, time of day and day of week, and numerous other variables. With Clearflow, travel time predictions made on all segments across a geographic region are used in route-planning searches (Delling et al., in press).

Beyond expected flows, it is important to consider uncertainty in travel time caused for instance by unpredictable traffic light schedules, accidents, unexpected road conditions, and differences in driver behavior. Such travel time variability (conversely, its *reliability*) also strongly affects the desirability of routes in the road network (Jenelius, 2012; Texas Transportation Institute, 2015). For fleets of delivery vehicles, such as those transporting perishables, decisions including routing need to provide on-time deliveries with high probability. In the case of ambulance fleets, taking into account uncertainty in the travel time of an ambulance to potential emergency scenes leads to improved ambulance positioning decisions, and consequently increases the survival rate of cardiac arrest patients (Erkut et al., 2007). A prediction of the probability distribution of travel time can be more valuable than a deterministic prediction of travel time, by accounting not just for measured traffic congestion and other known conditions, but also for the presence of unmeasured conditions. Distribution predictions of travel time can be used for risk-averse routing, for reporting travel time reliability to a user (e.g. the travel time is predicted to be in the range 10–15 min), and as a component of fleet vehicle decision-support systems (Samaranayake et al., 2012; Westgate et al., 2016).

We introduce a statistical solution to predicting the distribution of travel time on an arbitrary route in the road network, at an arbitrary future time. We call the method TRIP (for travel time reliability inference and prediction). For typical road networks of interest, the number of possible routes is extremely large, and any particular route may have very few or no observed trips in the historical data. For these reasons it is infeasible to apply methods designed for prediction on a particular set of heavily traveled routes, such as Jenelius and Koutsopoulos (2013), Ramezani and Geroliminis (2012), Rahmani et al. (2015). TRIP uses information from all the trips in the historical data to train a model for travel time on routes, learning the characteristics of individual roads and the effect of time of the week, road classification, and speed limit. We model travel time variability both at the trip level and at the level of the individual road network links included in the route. This decomposition is appropriate because some sources of variability affect the entire trip (such as driver habits, vehicle characteristics, or unexpected network-wide traffic conditions), while other sources of variability are localized (such as a delay due to a train crossing or construction). We define a network link to be a directed section of road that is not divided by an intersection, and on which the measured features of the road (road classification, speed limit, number of lanes, etc.) are constant.

TRIP captures important features of the data, including weekly cycles in congestion levels, heavy right skew of travel time distributions, and probabilistic dependence of travel times for different links within the same trip (for example, if the travel speed is high on the first link of the route, the speed is also likely to be high on the other links of the route). We capture the multimodality of travel time distributions using a mixture model where the mixture components correspond to unobserved congestion states, and model the probabilistic dependence of these congestion states across the links of the route using a Markov model. Because we model travel time for individual links, the travel time prediction can be updated en route.

We introduce a computational method for training and prediction based on maximum a posteriori estimation via Expectation Conditional Maximization (Meng and Rubin, 1993). This yields an iterative training process with closed-form update equations that can be computed using parallelization across links and trips. As a result it is computationally efficient even on large road networks and for large datasets.

TRIP uses GPS data from vehicle trips on the road network; we obtain large volumes of such trips using anonymized mobile phone GPS data from Windows phones in the Seattle metropolitan region. We compare the accuracy of our predictions to a variety of alternative approaches, including Clearflow. The GPS location and speed measurements are illustrated in Fig. 1, which shows that they contain valuable information regarding the speed of traffic on individual roads. Unlike other sources of vehicle speed information (Hofleitner et al., 2012b), vehicular GPS data does not require instrumentation on the roadway, and can achieve near-comprehensive coverage of the road network. Additionally, there is increasing evidence that traffic conditions can be estimated accurately using only vehicular GPS data (Work et al., 2010). One challenge of mobile phone GPS data is that it is often sampled at low frequency (typically 1–90 s between measurements). A related data source, GPS data from fleet vehicles, is also often sampled at low frequency (Rahmani et al., 2015), and TRIP can also be applied to such data.

Some existing approaches to predicting the probability distribution of travel time on a road network model exclusively link-level variability, and assume independence of travel time across the links in the route (Westgate et al., 2013; Hunter et al., 2013). This leads to considerable underprediction of the amount of variability (Westgate et al. (2013) and Section 4). Dependence across links is incorporated by Hofleitner et al. (2012a,b), who use a mixture model for travel time on links where the mixture component represents a congestion state (as we do). They allow these congestion states to depend on the link and the time, and model dependence of their congestion states across the road network and across time using a dynamic Bayesian network. This approach is intuitive but computationally demanding (leveraging a high-dimensional particle filter in each iteration of the algorithm), so is unlikely to be efficient enough to use on complete road networks (they apply it to 800 links in the San Francisco arterial network). Additionally, the method still underpredicts the amount of variability in travel time. Motivated by evidence in the data (Section 4), we allow the congestion state to additionally depend on the whole trip. We model dependence of this congestion state across the links of the route, instead of across all links of the