Vehicle Routing with an Environment Model

Tomáš Šingliar
Boeing Research and Technology
Bellevue, WA
tomas.singliar@boeing.com

Miloš Hauskrecht
Department of Computer Science
University of Pittsburgh
milos@cs.pitt.edu

Abstract

We consider the problem of vehicle routing in the fully general setting of a transportation network with stochastic, time-dependent travel times. To forecast future system state, we adapt a traffic flow model to provide probabilistic predictions. This becomes a generative model for the Markov decision process formalizing the planning problem. We propose an efficient planning algorithm to take advantage of such a predictive asset. The evaluation on actual traffic flow data shows that the largest improvement is to be had from forecasting the traffic flow better, followed by a smaller improvement from the new planning algorithm.

1 Introduction

Ideally, expected shortest time planning would be based on the exact time cost of future actions. However, travel time (action cost) in transportation networks varies with the traffic conditions and is not completely predictable. As the driver executes a plan, changing conditions and/or improved traffic prediction may change the optimal course of action. Algorithmically, we thus have two interlocking challenges: a prediction task, and a planning problem.

A fully general Markov decision process (MDP) formulation is not tractable with a prohibitively large state space encompassing the traffic state. A simple observation that one driver's actions have no measurable effect on traffic flow is a powerful approximate assumption. In such a case, the two separate tasks decouple into predicting traffic, and planning efficiently with such prediction.

Civil engineers have several theories of traffic flow [Schadschneider *et al.*, 2005]. We modify a cell transmission model to serve as a generative model for the (Semi-)MDP formalizing the planning task. Advantages of our dynamic model include that unlike most flow models, it is natural to parameterize – learn from observed data; and that its parameters have a clear interpretation. Because of the complex non-linear interactions involved, we implement a particle filter for approximating the predictive distributions of traffic quantities.

We propose a planning algorithm (the k-robust planner) that works by sampling several traffic state trajectories. Then it finds the shortest path for each of these trajectories and finally selects the best candidate according to a prior decision criterion.

2 The model

We represent the state of traffic flow vectors of state variables defining the flow speed $\mathbf{v}$, volume $\mathbf{y}$ and demand $\mathbf{q}$ (number of vehicles present). The components of the vectors correspond to road segments. The dynamics reflect the interactions among traffic variables over time and space. The joint distribution of traffic quantities is assumed to factorize along the segments at each time. The "entanglement" by temporal interaction is sufficient to model the spatial dependencies. For this to hold, the model time step δt needs to be sufficiently fine.

Interactions occur mostly between demands and volumes, while speeds v_i^t are thought of as a "dependent variable". Hence the one-step conditional distribution decomposes to:

$$p(\mathbf{s}^{t+\delta t}|\mathbf{s}) = \prod_{i=1}^{N} p(q_i^{t+\delta t}|\mathbf{s}^t)p(y_i^{t+\delta t}|\mathbf{s}^t)p(v_i^{t+\delta t}|y_i^{t+\delta t}) \quad (1)$$

The hidden state (demand $\mathbf{q}$) and observed (volume $\mathbf{v}$) dynamics is derived from the law of flow conservation. The distribution of the demand on a segment depends on the previous state and the volume shifting to and from neighboring segments. The inflow/outflow volume to/from a segment is an unobserved quantity (due to lack of instrumentation on entry/exit ramps) and is the major source of uncertainty in the model.

The third term of the traffic model in Equation 1, $p(v_i|y_i)$, captures how speed on a particular segment of the road depends on its traffic volume. In our data, this relationship varies widely depending on physical properties of the road infrastructure. To model this, we estimate the joint $p(v_i, y_i)$, separately for each segment i, from historical data by a kernel estimate. A small Gaussian kernel is centered on each data point $(v_i^{(n)}, y_i^{(n)})$, with axis-aligned covariance equal to the empirically observed variance in the respective covariates. Figure 1 shows an example of the volume-speed diagram implemented in the model.

Inference. Let us denote the time of last observation by T. Our prediction goal corresponds to the probabilistic query

$$p(\mathbf{v}^{T+1:H}, \mathbf{y}^{T+1:H}|\mathbf{v}^{1:T}, \mathbf{y}^{1:T}) \quad (2)$$

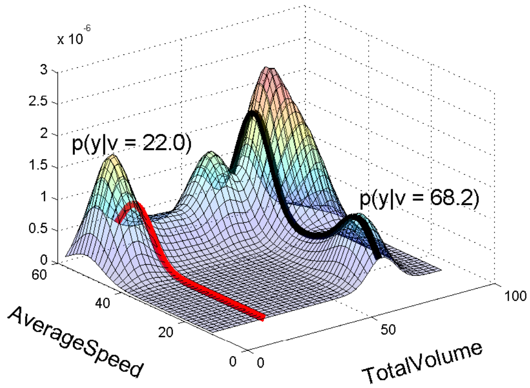

Figure 1: A fundamental diagram and the resulting conditional distributions (thick lines) for two flow volume values.

We need to track the hidden state $\mathbf{q}$ until time T (with $\mathbf{v}$ and $\mathbf{y}$ observed) and predict from there on, with $\mathbf{v}$ and $\mathbf{y}$ unobserved. An approximate Monte Carlo scheme [Doucet *et al.*, 2001] needs to be used due to the arbitrary form of input distributions and non-linear interactions (Figure 1).

3 Planning algorithm

The proposed algorithm is inspired by the particle filtering and somewhat reminiscent of Pegasus [Ng and Jordan, 2000]. First, we sample from the dynamic model k traffic state trajectories $T^{(1)}, \ldots, T^{(k)}$. A trajectory is a mapping from $T : Time \to S$, where S is the space state of traffic variables $\mathbf{s} = (\mathbf{q}, \mathbf{y}, \mathbf{v})$. In practice, we use a finely-grained piecewise linear approximation. Each state trajectory $T^{(i)}$ begins in the start state $\mathbf{s}_0$ which coincides with the time of plan execution: $T^{(i)}(t_0) = \mathbf{s}_0$. *The trajectory $T^{(i)}$ now represents a fixed deterministic evolution of the system*; thus, the link traversal times are now fixed and we can plan "optimally" using plain A^*, obtaining a plan (path) $p^{(i)}$.

Thus we have k paths $p^{(1)}, \ldots, p^{(k)}$, each of them optimal under a possible future evolution of the system. To pick one as the result of planning, we evaluate how well each performs on the remaining trajectories; e.g. evaluate $p^{(1)}$ against $T^{(2)}, \ldots, T^{(k)}$, thus obtaining $k-1$ samples from the distribution of $p^{(1)}$'s cost. The path with the best mean travel time (taken over the $k-1$ remaining trajectories) is selected as the result. Any other selection criterion may be used as well. We call this algorithm k-robust planning. Clearly, the complexity of this algorithm will be dominated by running the k instances of plain A^*.

4 Experimental evaluation

The data comes from Pittsburgh metro area, which has approximately 150 sensor locations that collect the volume (number of vehicles) and the average speed of travel in 5-minute intervals. The underlying topology and speed limits are provided by the ArcGIS geographic information system and the StreetMap database.

There are often very few sensible alternatives for travel between two points in a city[1]. Even when one encounters an adverse traffic situation, it rarely justifies changing one's route. Therefore, for our comparison we manually selected origin-destination pairs where at least two routes very close in expected travel time exist.

Four combinations of planning methods and prediction models are executed:

1. the baseline A^* with speeds determined by historical average for the space and time of day
2. A^* that takes for its cost the mean of the model predictive distribution
3. the k-robust method for two different values of k
4. Upper bound is established by a clairvoyant method that always picks the optimal path

The results show that in all cases, using the dynamic model to predict traffic state and planning with that prediction yields significantly shorter travel times (as much as 5 minutes on a half-hour trip) than the baseline approach, which uses observed historical average speed to determine travel time.

The proposed k-robust algorithm (for $k = 5$) also results in an improvement, albeit about 7 times smaller than the one realized by using the model's prediction. A further small improvement (half again) is observed for $k = 20$. This suggests that inter-temporal correlations of travel cost indeed play a role, even though a weaker one compared to the importance of predicting the evolution of traffic better on average. The best method closes about 75% of the gap between the baseline and the unachievable clairvoyant solution.

5 Conclusions

We have shown that vehicular traffic prediction can be successfully exploited in order to plan faster routes, giving a model and a planning algorithm to use it. There is of course a catch for practical applications: we have assumed that the driver's actions do not influence the evolution of traffic. The predictions are no longer valid if a non-negligible fraction of the traffic network users execute plans based on such predictions. This is a much more complex topic we deliberately avoided in this paper.

References

[Doucet *et al.*, 2001] A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer Verlag, New York, 2001.

[Ng and Jordan, 2000] A. Ng and M. Jordan. PEGASUS: A policy search method for large MDPs and POMDPs. In *Proceedings of UAI-00*, 2000.

[Schadschneider *et al.*, 2005] Andreas Schadschneider, Thorsten Pöschel, Reinhart Kühne, Michael Schreckenberg, and Dietrich E. Wolf, editors. *Traffic and Granular Flow*, 2005.

[1]Here we are limited to roads major enough to have sensor instrumentation and thus data for learning and evaluation.