

Learning Traffic Flow Dynamics using Random Fields

Deepthi Mary Dilip¹, DianChao Lin², and Saif Eddin Jabari^{*2,3}

¹BITS Pilani, Dubai Campus, Dubai International Academic City, Dubai, U.A.E.

²New York University Tandon School of Engineering, Brooklyn NY

³New York University Abu Dhabi, Saadiyat Island, P.O. Box 129188, Abu Dhabi, U.A.E.

June 22, 2018

Abstract

This paper presents a mesoscopic stochastic model for the reconstruction of vehicle trajectories from data made available by subsets of (probe) vehicles. Long-range vehicle interactions are applied in a totally asymmetric simple exclusion process to capture information made available to connected and autonomous vehicles. The dynamics are represented by a factor graph, which enables learning of traffic dynamics from historical data using Bayesian belief propagation. Adequate probe penetration levels for faithful reconstruction on single-lane roads is investigated. The estimation technique is tested using a vehicle trajectory dataset generated using an independent microscopic traffic simulator. Although the parameters of the traffic state estimation model are learned from (simulated) historical data, the proposed algorithm is found to be robust to unpredictable conditions. Moreover, by exposing the algorithm to varying traffic conditions with increasingly larger datasets, the probe penetration rates required to capture the traffic dynamics effectively can be substantially reduced. The results also highlight the need to take into account randomness in the spatio-temporal coverage associated with probe data for reliable state estimation algorithms.

Keywords: Automated vehicles; cellular Automata; conditional random fields; stochastic traffic modeling; traffic state estimation; trajectory reconstruction

*Corresponding author, e-mail: sej7@nyu.edu

1 Introduction

Automated vehicle (AV) technologies are beginning to penetrate vehicle fleets in cities throughout the world. It is reasonable to expect that vehicle trajectory data from AVs (e.g., through on-board GPS units) will become a prominent source of high-resolution traffic data. AVs may act as probes in the traffic stream, continuously broadcasting their position and speed in real-time. More importantly, AVs can also provide distance headways (spacing between successive vehicles) using infrared or radio technology [Yuan et al., 2012, Seo and Kusakabe, 2015]. However, privacy issues and technology limitations can limit the ability of traffic management agencies to collect, analyze, and disseminate such information. To overcome this, these data can be fused with data obtained from traditional monitoring devices such as inductive-loop detectors (stationary sensors). As these two data sources complement each other, comprehensive datasets can be obtained for traffic monitoring and state estimation [Hofleitner et al., 2012a]. However, the improvement in accuracy with data fusion, over single sensor applications depends on probe penetration rates and on traffic conditions. In urban road networks, where stationary sensor instrumentation is usually limited and traffic lights play a governing role in the traffic dynamics, a higher number of probes may be necessary to accurately characterize traffic conditions.

A number of modeling techniques have been proposed in the recent years to estimate traffic densities [Herrera and Bayen, 2010, Jabari and Liu, 2012, 2013, Deng et al., 2013], speeds [Work et al., 2008] and travel times [Hellinga et al., 2008, Chen and Rakha, 2012].

Studies have also been carried out to extract patterns from streaming data using data mining techniques [Hunter et al., 2012, Jenelius and Koutsopoulos, 2017, Dilip et al., 2017, Jabari et al., 2018]. To account for the variability in urban traffic, a statistical approach using Coupled Hidden Markov Models was proposed by Herring et al. [2010] to estimate the traffic state from sparse probe data. The limitations of purely statistical approaches were overcome by Hofleitner et al. [2012b], where a hybrid modeling framework combining machine learning with hydrodynamic traffic theory was proposed to predict arterial travel times from streaming GPS probe data. On the other hand, Papathanasopoulou and Antoniou [2015] proposed a data-driven model that captures longitudinal interactions between vehicles.

Research on traffic state estimation from probe data for urban networks has focused on the reconstruction of traffic states at an aggregate level (over an entire intersection-to-intersection road segment) [Furtlehner et al., 2007]. At a finer scale, Herrera and Bayen [2010] reconstructed traffic densities on a freeway section by modifying traditional continuum models with a correction term to nudge the model estimate towards the GPS probe measurements. The techniques proposed did not require the knowledge of on- and off-ramp sensor data for density estimation, and minimal penetration rates required on arterial roads was also analyzed. Vandenberghe et al. [2012] discussed the maximum sampling interval (time between two consecutive probe vehicle samples) required to accurately detect incidents and Jabari and Wynter [2016] study the optimal placement of sensors for

reliable time-to-detection of incidents. Mazaré et al. [2012] perform comparisons of travel time estimates produced using one source of data versus fused data from two data sources (stationary sensors and probe vehicle data). A number of studies have reported probe penetration rates required for traffic state estimation on arterials [Hiribarren and Herrera, 2014, Ban et al., 2009, 2011, Zheng et al., 2018]. While the reliability of probe vehicle data has been investigated and compared against stationary sensor data [Kim and Coifman, 2014, Bar-Gera, 2007], variability in the spatio-temporal coverage of probe vehicles has not been adequately studied, particularly at the fine level required for effective traffic management.

Motivated by the wide spatio-temporal coverage offered by fused traffic data, the question of adequate levels of probe penetration is addressed at a microscopic scale in this paper. The focus of this paper is on the reconstruction of vehicle trajectories over a single roadway, where a stationary sensor captures the arrival times of the vehicles and the speeds of the probe vehicles is used to infer the traffic state over the entire link. A probabilistic approach is proposed for the spatio-temporal reconstruction of dynamic traffic state from sparse probe data, wherein the traffic patterns are learned from historical data using Conditional Random Fields (CRFs). By modeling the vehicle interaction potential to reflect local traffic information (such as spacings between vehicles), our estimation models seamlessly combine traditional car-following theory and simulation with statistical learning techniques to reconstruct microscopic traffic dynamics.

The remainder of the paper is organized as follows: The stochastic traffic flow model along with a simulation algorithm are presented next. The probabilistic inference problem is then stated. This is followed by a presentation of the factor graph approach proposed for solving the inference problem. Model testing and validation experiments are then given. A brief summary and future research discussion conclude the paper.

2 Stochastic Traffic Dynamics Model

2.1 Look-Ahead Dynamics

A discrete (state and time) *mesoscopic* stochastic model is used to represent the traffic dynamics. Vehicle movement in the model is governed by potential functions that describes the (“energy profile” of) local traffic conditions. Similar (but simpler) models have been employed to study interesting traffic phenomena like *synchronized traffic* at ramps and *stop-and-go* regimes [Sopasakis and Katsoulakis, 2006].

The physical roadway is modeled as a one-dimensional uniform lattice L . The spatial co-ordinates of each vehicle α on the roadway is discretized in such a way that each cell can be occupied by at most one vehicle, which is achieved by setting the cell length to an appropriate value, e.g. 7.5 m [Lárraga et al., 2005]. The state of each occupied cell at a discrete time k is completely specified by a discretized speed denoted $v_\alpha^k \in \{0, \dots, v_{\max}\}$, where v_{\max} is the maximum number of cells that can be traversed by a vehicle in one time step. Clearly, v_{\max} depends on the length of the discrete time step, δk . Thus, an order parameter $\sigma^k(l) \in \{-1, 0, \dots, v_{\max}\}$ can

be defined for each cell $l \in L$ at time k to represent the traffic state in the cell, where 0 represents a free cell.

A look-ahead potential is used to capture the response of vehicle α to traffic conditions ahead. Specifically, the state of vehicle α at time-step $k+1$ is a function of their current speed v_α^k and the current speed of their leader $v_{\alpha-1}^k$. Denote the traffic state pertaining to vehicle α at time step k by the vector $Y_\alpha^k = [v_\alpha^k \ v_{\alpha-1}^k]^\top$. More generally, a vehicle's look ahead potential can depend on multiple vehicles ahead. Let M denote the look-ahead distance (in number of vehicles), then $Y_\alpha^k = [v_\alpha^k \ v_{\alpha-1}^k \ \dots \ v_{\alpha-M}^k]^\top$. The look-ahead potential is given, for vehicle α and each $i \in \{0, \dots, v_{\max}\}$ by

$$\varepsilon_\alpha^k(i) = W_i^\top Y_\alpha^{k-1}, \quad (1)$$

where $W_i \in \mathbb{R}^{|Y|}$ is a weight vector that captures the relative importance of each of the state variables in Y_α^k when assessing the energy of vehicle α if it were to assume speed $i \in \{0, \dots, v_{\max}\}$. The parameters can, hence, be encoded into a matrix $W \in \mathbb{R}^{|Y| \times (v_{\max}+1)}$ as $W \equiv [W_0 \ \dots \ W_{v_{\max}}]$. The probability that vehicle α assumes speed i is related to the potential energy as

$$\mathbb{P}(v_\alpha^k = i | Y_\alpha^{k-1}) \propto e^{-\varepsilon_\alpha^{k-1}(i)} = e^{-W_i^\top Y_\alpha^{k-1}}. \quad (2)$$

The position of vehicle α at time step k , denoted by s_α^k , is updated based on its speed v_α^k as follows:

$$s_\alpha^k = \min(s_\alpha^{k-1} + v_\alpha^k, s_{\alpha-1}^{k-1} - 1), \quad (3)$$

where taking the minimum of this quantity and $s_{\alpha-1}^{k-1} - 1$ ensures that vehicles and their leaders do not occupy the same position. Consequently, the state update calculation proceeds in ascending order of α (from the position-wise upstream-most vehicle to the downstream-most vehicle). The traffic state update is carried out in discrete time steps to determine $\sigma^{k+1}(l)$.

2.2 Vehicle Interactions and Coordination

The modeling approach above is, in essence, a totally asymmetric simple exclusion process (TASEP). Such processes are known to have limitations in some contexts, namely, heterogeneous environments with AVs [Jerath et al., 2014]. This is a result of vehicles *reacting* to the traffic conditions in the downstream. Introducing an interaction potential overcomes these limitations and offers enhanced interpretability to the probabilities of vehicles advancing to downstream cells Jerath et al. [2015].

Consider two interacting vehicles, α and β and let $Y_{\alpha,\beta}^k$ denote the traffic state pertaining the interaction between α and β . The interaction between α and β depends on distance between them and their speed difference: $Y_{\alpha,\beta}^k \equiv [g_{\alpha,\beta}^k \ \Delta v_{\alpha,\beta}^k]^\top$, where $g_{\alpha,\beta}^k \equiv s_\alpha^k - s_\beta^k$ and $\Delta v_{\alpha,\beta}^k \equiv v_\alpha^k - v_\beta^k$. For $i, j \in \{1, \dots, v_{\max}\}$, define the interaction potential as

$$\varepsilon_{\alpha,\beta}^k(i, j) = \theta_{i,j}^\top Y_{\alpha,\beta}^k, \quad (4)$$

where $\theta_{i,j} \in \mathbb{R}^2$ is a vector of two parameters which represent the relative importance of each element in $Y_{\alpha,\beta}^k$ and $\varepsilon_{\alpha,\beta}^k(i, j)$ can be interpreted as the potential energy

associated with vehicle α assuming speed i and vehicle β assuming speed j given their present state at time k . This is combined with the look-ahead dynamics by defining a total energy function for $|\mathcal{V}|$ vehicles as

$$E^k(i_1, \dots, i_{|\mathcal{V}|}) = \sum_{\alpha \in \mathcal{V}} \varepsilon_{\alpha}^k(i_{\alpha}) + \sum_{\alpha, \beta: \alpha \neq \beta} \varepsilon_{\alpha, \beta}^k(i_{\alpha}, i_{\beta}) \quad (5)$$

and the probability that the $|\mathcal{V}|$ vehicles assume the speeds $(i_1, \dots, i_{|\mathcal{V}|})$, given the state of the system is given by $\mathbf{Y}^k \equiv \{Y_{\alpha}^k, Y_{\alpha, \beta}^k\}_{\alpha, \beta}$, is related to the total energy as follows:

$$\mathbb{P}(v_1^k = i_1, \dots, v_{|\mathcal{V}|}^k = i_{|\mathcal{V}|} | \mathbf{Y}^{k-1}) \propto e^{-E^{k-1}(i_1, \dots, i_{|\mathcal{V}|})}. \quad (6)$$

When two vehicles α and β do not interact (e.g., vehicles that are too far apart), they are simply omitted from the sum in (5). This can be utilized when, for example, vehicles only interact with other vehicles that are in the immediate vicinity. However, the model is general enough to allow for general interactions allowing communication and coordination between vehicles that are not immediately adjacent.

The steps involved in simulating the traffic dynamics are summarized in Algorithm 1 below, which without loss of generality assumes a free downstream boundary. This can be easily modified to accommodate downstream restrictions in a way that is similar to the upstream state update (see Jabari [2016] for details on boundary treatments).

Algorithm 1: Simulation of Traffic Dynamics

<p>Input: No. lattice sites := N, No. time steps := K, time step := δk, max speed := v_{\max} Look-ahead distance := M, look-ahead matrix := W, interaction tensor := Θ Arrival density := p_1, probability of Slow-down := p_2</p> <p>Initialize: Initial traffic state := $\sigma^0(l)$</p> <p>Iterate: For $k = 1 : K$ do For each α, set $Y_{\alpha}^{k-1} := [v_{\alpha}^{k-1} \ v_{\alpha-1}^{k-1} \ \dots \ v_{\alpha-M}^{k-1} \ g_{\alpha}^{k-1} \ \dots \ g_{\alpha-M+1}^{k-1}]^{\top}$ For each $\beta \neq \alpha$, set $Y_{\alpha, \beta}^k := [g_{\alpha, \beta}^k \ \Delta v_{\alpha, \beta}^k]^{\top}$ For $i \in \{0, \dots, v_{\max}\}$ do Calculate the look-ahead potential $\varepsilon_{\alpha}^k(i) := W_i^{\top} Y_{\alpha}^{k-1}$ For $j \in \{0, \dots, v_{\max}\}$ do Calculate the interaction potential $\varepsilon_{\alpha, \beta}^k(i, j) := \theta_{i, j}^{\top} Y_{\alpha, \beta}^k$ End For End For For each α do Total potential $E_{\alpha}^k(i) := (v_{\max} + 1)\varepsilon_{\alpha}^k(i) + \sum_{j=0}^{v_{\max}} \sum_{\alpha, \beta: \alpha \neq \beta} \varepsilon_{\alpha, \beta}^k(i, j)$ Probability $\pi_{\alpha}^k(i) := E_{\alpha}^k(i) / \sum_{i=0}^{v_{\max}} \sum_{j=0}^{v_{\max}} (\varepsilon_{\alpha}^k(i) + \sum_{\alpha, \beta: \alpha \neq \beta} \varepsilon_{\alpha, \beta}^k(i, j))$ Velocity Update: Sample v_{α}^k from $[\pi_{\alpha}^k(1) \ \dots \ \pi_{\alpha}^k(v_{\max})]^{\top}$</p>
--

```

     $u_1 \sim \text{Uniform}(0, 1)$ 
    If  $u_1 < p_2$  then
         $v_\alpha^k := v_\alpha^k - 1$ 
    End If
    Position Update:
    Compute vehicle positions  $s_\alpha^k$  in ascending order of  $\alpha$ :
    If  $s_\alpha^k := \min(s_\alpha^{k-1} + v_\alpha^k, s_{\alpha-1}^k - 1) > N$  then
         $s_\alpha^k := \infty$ 
    End If
    Traffic State Update:
     $\sigma^k(s_\alpha^k) = v_\alpha^k$ 
End For
Boundary Conditions :
     $u_2 \sim \text{Uniform}(0, 1)$ 
    If  $u_2 < p_1$  and  $\sigma^k(1) = 0$  then
         $y := [v_{\max} \ v_{\underline{\alpha}}^k \ s_{\underline{\alpha}}^k]^\top$  ( $\underline{\alpha}$  is the index of upstream-most vehicle in the system)
        For  $i \in \{0, \dots, v_{\max}\}$ , calculate  $\hat{\pi}_i^k := \exp[-W_i^\top y]$ 
        and normalize:  $\pi_i^k := \hat{\pi}_i^k / \sum_{i'=1}^{v_{\max}} \hat{\pi}_{i'}^k$ 
         $\sigma^k(1) \sim [\pi_1^k \ \dots \ \pi_{v_{\max}}^k]^\top$ 
    End If
End For

```

3 Probabilistic Inference

Assume there are $|\mathcal{V}|$ vehicles in the system and at each time step k , the state (speed and position) of a *subset* of these vehicles is observed. The estimation problem is concerned with determining the state of all vehicles given the partial observations. More accurately, the estimation problem seeks to fit the conditional probability distribution of the state given the observations. Let $v = [i_1 \ \dots \ i_{|\mathcal{V}|}]^\top$ and let $\pi_0(v) \equiv \mathbb{P}(v_1^0 = i_1, \dots, v_{|\mathcal{V}|}^0 = i_{|\mathcal{V}|})$, the prior distribution of speeds (state at time 0), be given. Let $\sigma^{k,\text{obs}}$ denote the observed traffic states (measurements) available at time k . For each time step, the inference problem seeks to determine the conditional probability

$$\pi^k(i_1, \dots, i_{|\mathcal{V}|}) \equiv \mathbb{P}(v_1^k = i_1, \dots, v_{|\mathcal{V}|}^k = i_{|\mathcal{V}|} | \hat{\mathbf{Y}}^{k-1}, \sigma^{k,\text{obs}}), \quad (7)$$

where $\hat{\mathbf{Y}}^{k-1}$ is based on the maximum a posteriori (MAP) estimate of the traffic state at time step $k-1$. Inferring the traffic state at time k is achieved using MAP estimates:

$$\hat{v}_\alpha^k = \arg \max_{i \in \{0, \dots, v_{\max}\}} \pi_\alpha^k(i), \quad (8)$$

where

$$\pi_\alpha^k(i) \equiv \mathbb{P}(v_\alpha^k = i | \hat{\mathbf{Y}}^{k-1}, \sigma^{k,\text{obs}}) \quad (9)$$

is the (conditional) marginal probability density of the state of vehicle α at time k given the state at time $k-1$ and observations at time k . These estimates can

then be used to set up $\widehat{\mathbf{Y}}^k$ as input for the next time step. Although probabilistic inference can be achieved by brute force, this can be computationally tedious. In the next section, it is demonstrated how the dependence structure in the model can be exploited using graphical modeling and how inference in this case can be efficiently performed.

4 Factoring and Factor Graph Representation

Let $V = \{V_\alpha | \alpha \in \mathcal{V}\}$ be a discrete valued random field with the probability mass function (pmf), $\pi(\mathbf{v}) = \pi(i_1, \dots, i_{|\mathcal{V}|})$. The random field V is a Markov random field (MRF) if it satisfies the Markovian property

$$\mathbb{P}(V_\alpha = i | V) = \mathbb{P}(V_\alpha = i | V_{\mathbf{N}_\alpha}) \quad (10)$$

for all $\alpha \in \mathcal{V}$, where \mathbf{N}_α denotes the set of “neighbors” of α . These (conditional) independence assumptions between the variables V_α can be encoded in a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} is indexed by the vertices (also called nodes) \mathcal{V} such that $V = (V_s)_{s \in \mathcal{V}}$ and edges (or arcs) $\mathcal{E} \in \mathcal{V} \times \mathcal{V}$. Note that the vertices coincide with the vehicles themselves; it is for this reason that the same notation is used in both cases. In this study, the edges \mathcal{E} connect the (speeds of) vehicles, rather than the lattice cells, as the random variables of the MRF model. By encoding the spatial dependencies in the speed field through edges \mathcal{E} , the condition in (10) implies that speed of any vehicle is independent of the traffic state given the **local** speed field.

The joint probability distribution over all the variables in the Markov model can be compactly represented by defining a set of cliques $\{\mathcal{C}\}$, which are subsets of vertices of \mathcal{G} such that all vertices in each clique \mathcal{C} are completely connected or mutually adjacent. A clique is said to be maximal if no other vertex in \mathcal{G} can be added without violating the clique property. The joint distribution of the MRF can be expressed over the maximal cliques, as a product of factors as [Koller and Friedman, 2009],

$$\pi(\mathbf{v}) = \frac{1}{Z} \prod_{c \subset \mathcal{G}} \psi_c, \quad (11)$$

where Z is the normalizing constant, \mathbf{v}_c is the restriction of \mathbf{v} to the vertices in the clique c , and $\{\psi_c\}_{c \subset \mathcal{G}}$ is a set of “factors” for each maximal clique $c \subset \mathcal{G}$. Each factor is a non-negative function defined over a clique to represent the (unnormalized) probability distribution between the nodes in the clique. When the potentials are restricted to be strictly positive, the factors can be re-parametrized in the log space, and expressed in terms of the Boltzmann distribution as $\psi_c = e^{-E_c(\mathbf{v}_c)}$, where $\{E_c(\mathbf{v}_c)\}_{c \subset \mathcal{G}}$ are properly defined potential functions over the cliques. Hence,

$$\pi(\mathbf{v}) = \frac{1}{Z} e^{-\sum_{c \subset \mathcal{G}} E_c(\mathbf{v}_c)}. \quad (12)$$

This is analogous to (5): the lower the energy of the clique configuration (of the states), the higher the probability of the configuration. Traditional techniques used to represent the dependency in the traffic dynamics graphically would result in

what is known as loopy Markov model. These are known to not factor uniquely. Instead, the proposed approach defines a partially directed graph, or a Conditional Random Field, parameterized on a set of factors based on the stochastic traffic dynamics presented above. For these dynamics and for each time step, k , the variables Y_α^k and $Y_{\alpha,\beta}^k$ provide a natural means of factoring the probability distribution over the entire graph. Indeed, conditioning on \mathbf{Y}^{k-1} , $\pi^k(i_1, \dots, i_{|\mathcal{V}|}) = \prod_{\alpha \in \mathcal{V}} \pi_\alpha^k(i_\alpha)$; see the definitions in (7) and (9).

Consider the case where probe vehicles are equipped with sensors capable of measuring distances and speeds of other vehicles that are immediately adjacent (their immediate leaders and followers). In this case, the look-ahead potentials are encoded into *node factors*, $\{\psi_\alpha^k\}_{\alpha \in \mathcal{V}}$, and interaction potentials into *edge factors*, $\{\psi_{\alpha,\beta}^k\}_{(\alpha,\beta) \in \mathcal{E}}$. (When two vehicle indices appear in the subscript, the factor is to be implicitly understood as an edge factor.) Fig. 1 presents a factor graph representation of a system with five vehicles, i.e., $\mathcal{V} = \{1, \dots, 5\}$ and interaction between adjacent vehicles. The vehicle speeds are presented by circular nodes in the graph. Factors are represented by square nodes in the graph, where the set $\{\psi_\alpha^k\}_{\alpha \in \mathcal{V}}$ are the factors pertaining to look-ahead dynamics, while $\{\psi_{\alpha,\beta}^k\}_{(\alpha,\beta) \in \mathcal{E}}$ are factors pertaining to vehicle interactions.




Fig. 1: Example factor graph.

The node factors are related to the look-ahead dynamics via the look-ahead potentials:

$$\psi_\alpha^k = [e^{-\varepsilon_\alpha^k(0)} \dots e^{-\varepsilon_\alpha^k(v_{\max})}]^\top \quad (13)$$

and the edge factors are related to vehicle interaction dynamics via the interaction potentials:

$$\psi_{\alpha,\beta}^k = \begin{bmatrix} e^{-\varepsilon_{\alpha,\beta}^k(0,0)} & \dots & e^{-\varepsilon_{\alpha,\beta}^k(0,v_{\max})} \\ & \ddots & \\ e^{-\varepsilon_{\alpha,\beta}^k(v_{\max},0)} & \dots & e^{-\varepsilon_{\alpha,\beta}^k(v_{\max},v_{\max})} \end{bmatrix}. \quad (14)$$

The factor graph framework enables exact inference of the local marginals over nodes or subsets of nodes in tree-structured graphs (such as the one depicted in Fig. 1) using the sum-product algorithm (a.k.a. Bayesian belief propagation). The sum-product algorithm computes the marginal distribution of any random variable(s) as a product of the incoming *messages* from its neighboring factor nodes, as shown in Fig. 1 for V_4 . A message $m_{\psi_c \rightarrow V_\alpha}$ from a factor ψ_c can be interpreted as the information contained in the factor about the variable V_α .

The algorithm is a generalization of the variable elimination algorithm to execute multiple queries on the same tree-structured graph efficiently by guiding the

order of operations. This is achieved by fixing any node as the root of the graph and determining the order in which the messages are propagated from the root to all its leaves by a depth first search algorithm. Messages are sent from all the leaf nodes to the root node during a single forward pass shown as red arrows in Fig. 2. The figure also indicates the order in which the nodes are traversed when the




Fig. 2: Illustration of forward (red arrows) and backward passes (black arrows) in the sum-product algorithm.

node V_7 is the root node. By then sending messages from the root node to the leaf nodes during a single backward pass and storing all the intermediate messages, the marginals over any subset of variables can be computed without recalculation for every inference. More details can be found in standard references on graphical models/machine learning; e.g., [Bishop, 2006, Koller and Friedman, 2009].

When measurements are available for some of the vehicles, given by $\sigma^{k,obs}$, the inferences can be conditioned by *clamping* the corresponding variables to the observed states. As a consequence conditional independence, when a node is observed, it breaks the chain structure into a forest of independent chains. For example, for the five vehicle system in Fig. 1 and Fig. 2, assuming $V_3 = 1$ is given (i.e., the speed of vehicle 3 is known/measured), the independent forest shown in Fig. 3 is obtained.




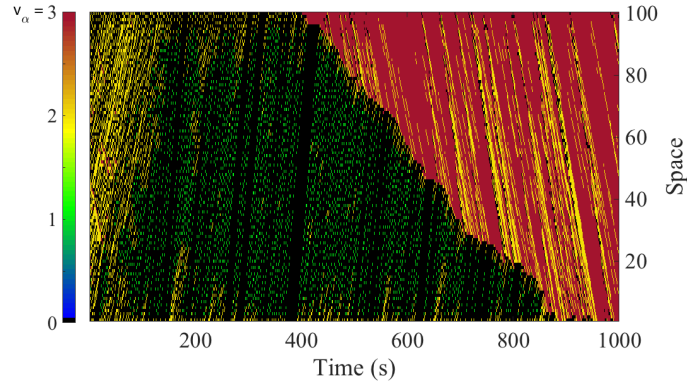
Fig. 3: Decomposition of the factor graph into a forest of independent sub-graphs in the presence of measurements.

5 Model Testing and Validation


Numerical Test: Shockwave

Consider a road of length N cells with $v_{max} = 3$ and assume an arrival density at the upstream end of $p_1 = 0.25$. At the upstream boundary, loop detectors provide information about the occupancy and speed of all upstream vehicles, as well as the entry times of new vehicles into the road. In order to validate the model (with out-of-sample data), an incident is simulated at the downstream boundary of the road section. The simulated traffic dynamics are illustrated in Fig. 4a, which

represents ground truth and used for comparison with the traffic dynamics reconstructed by the CRF model. The interface between the free-flow (red) part of the figure and the congested (green and black) part of the figure represents the propagation of congestion from the source of the incident at the top of the figure into the upstream (against the direction of traffic). The trajectory of the interface represents a shockwave in the traffic stream. A subset of all the simulated vehicles are chosen



(a) Ground truth



(b) Probe penetration=10%

Fig. 4: Validation of CRF model.

randomly to represent a set of probe vehicles. For this study, periodic noise-free updates of the vehicle position s_α^k and speed v_α^k (obtained from successive GPS coordinates) are assumed to be available from the probe vehicles at a (time) cadence of once every $\delta k = 1$ seconds. Consequently, $v_{\max} = 3$ cells and cell lengths of 7.5 meters correspond to a maximum speed of 81 km/hr. The CRF model is used to estimate the speed field sequentially in discrete time-steps, which correspond to the sampling interval of the probe vehicles. The estimated vehicle trajectories are shown in Fig. 4b, indicating that a probe penetration rate of 10% is sufficient to capture the backward propagation of a shockwave generated by the incident located downstream.

5.1 Learning Traffic Dynamics from Historical Data

In this experiment, “historical datasets” are simulated for a road section using a microscopic traffic simulation tool. The simulations are run for 1 hour periods (from 8 am to 9 am) with a 15 minute warm-up period for an arterial link that is 500m long. The vehicle trajectory data, collected from the simulation, is a sequence of spatial co-ordinates sampled every 1 second. This continuous trajectory data is discretized by dividing the roadway into a cell lattice (with cell lengths 7.5m). The free-flow speed along the arterial link is 120 km/hr, and a fixed-time signal is located at the downstream end with a green time of 90s. In order to generate sufficient historical data, datasets are generated for different traffic conditions by varying the random seed as well as by considering different traffic flow conditions. The effect of increasing the probe penetration rate on the estimated traffic states is shown in Fig. 5, which indicates that a penetration rate of 10% is not adequate for learning the traffic dynamics. The results also show that with a single set of historic

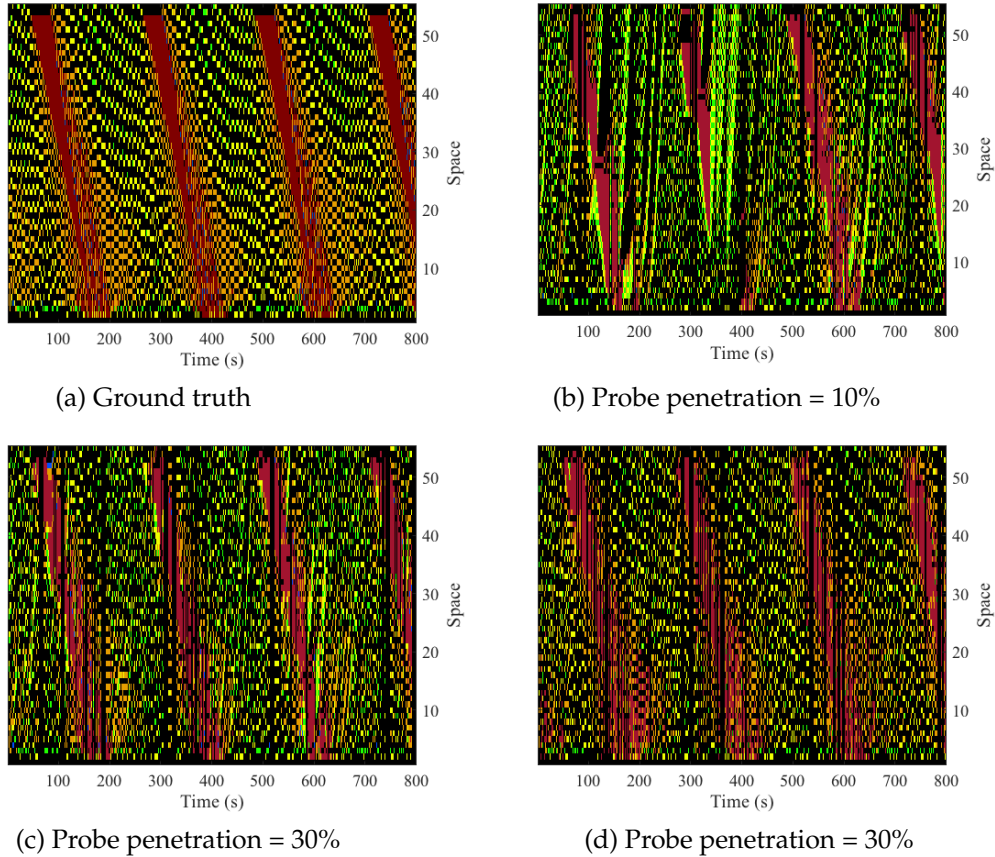


Fig. 5: Spatio-temporal velocity map. Black cells correspond to speeds $i = 0$, green to $i = 1$, yellow to $i = 2$, and red to $i = v_{\max} = 3$.

data to learn from, in order to capture the ground truth conditions with sufficient accuracy, a probe penetration of 20% or more is needed. This is compared to the

case where historical data are used to fit the parameters of the model in Fig. 6. After ten learning phases, it can be observed from Fig. 6 that even with a 10% probe penetration, vehicle trajectories can be reconstructed faithfully (comparable to using standard estimation techniques with a 30% probe penetration rate).





Fig. 6: The effect of utilizing historic data, i.e., learning, on the estimated traffic dynamics.

5.2 Probe Vehicle Distribution


Accuracy of the estimation problem not only depends on penetration rates of probe vehicles, but also on how the probes are distributed in the sample. The estimated traffic states are compared for two random distributions of probes, both representing a penetration rate of 5%. Fig. 7 depicts the distribution of the randomly selected probes in the upper half of the figure with the estimated trajectories in the lower half for a signalized arterial. In this experiment, three signal cycles are simulated with simulation time horizon of length $T = 900$ seconds with a red-time of 100 seconds. The cycles start at time steps $k = 100, 400,$ and 700 . While this information can be easily inferred from a probe level of 5% as shown in Fig. 7(a), when none of the sampled probes pass through the third signal cycle the estimation algorithm fails to capture the build-up and dissipation of queues in the time period from 700 to 900 seconds: as can be seen in Fig. 7(b).

To study the effect of randomness in probe coverage (the distribution of the probes in the sample), a road segment with a on-ramp located at about 300m downstream is considered. The free-flow speed was assumed to be 80km/hr, while the traffic demand was gradually increased in 15 minute intervals from 1200 veh/hr to 2500 veh/hr to capture the build-up and dissipation of the on-ramp queues. The speed ranges correspond to the discrete speed states (the number of lattice cells crossed by a vehicle in a single time step δk).

The traffic state estimated for a time-period of $T = 15$ minutes in congested conditions is depicted in Fig. 8. When compared with the ground truth, the estimate



(a) Probe distribution 1




(b) Probe distribution 2

Fig. 7: Different probe vehicle distributions.

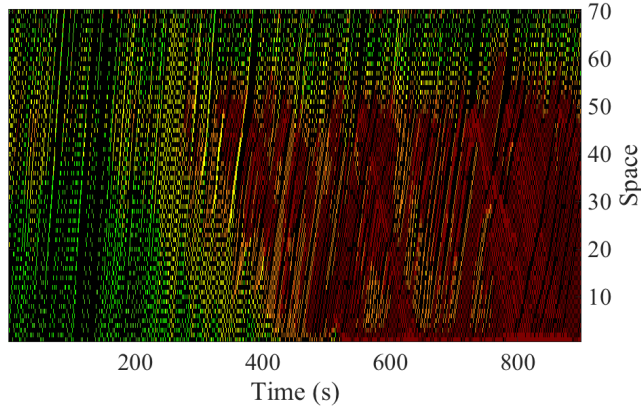
produced with a probe penetration of 20% is sufficient to capture the shockwaves created by the on-ramp.

It should be noted here that although no information regarding the entry times of the on-ramp vehicles was provided to our estimation algorithm, it can be inferred from the output in Fig. 8b. Similar estimation studies have indicated that probe levels of 2% can capture the shockwaves generated by lane-closure on a freeway [Bucknell and Herrera, 2014]. However, in settings where the random arrivals of vehicles govern (e.g., on-ramp vehicles), it is not surprising that a higher probe penetration rate is required for traffic state estimation in congested conditions. Moreover, as Goodall et al. [2016] summarize in their study on the microscopic estimation of freeway vehicles in a connected environment, for traffic signal control problems a minimum penetration rate of 20-30% is required, while for arterial performance measurement the penetration rate ranges from 10-50%.

The spatial distribution of the probe vehicles plays a significant role in the accuracy of estimated results, as observed earlier. To analyze the effect of the random-



(a) Ground truth



(b) Probe penetration = 20%

Fig. 8: Time-space diagram (velocity).

ness introduced by the probe vehicle distribution, $R = 100$ simulations are run for each of the vehicle probe penetration rates, choosing a probe distributions randomly for each penetration rate. The following penetration rates are considered: 5%, 10%, 20%, and 30%. Fig. 9 depicts the frequencies of the MAPEs for each of the penetration rates. The mean value of the MAPE for a probe penetration level of 5% is around 18%, but the high variability observed implies that travel time error can be even higher if the probe distribution is highly random. As the number of probe vehicles increases, this variability in the error decreases.

6 Conclusion

A methodology is presented for traffic state estimation that combines mesoscopic traffic modeling with the statistical power of probabilistic graphical models to learn the traffic patterns from historical data. The modeling approach includes both look-ahead dynamics along with vehicle interaction dynamics.

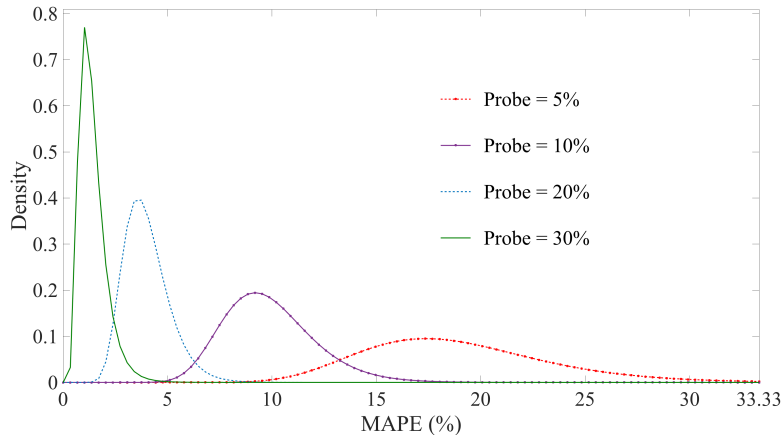


Fig. 9: PDF of the Mean Absolute Percent Error in travel time at different probe levels

A conditional random fields (CRF) approach using a factor graph representation of the dynamics is then proposed for purposes of statistical learning when limited data is available. Coverage of the probe vehicle information can be expected to be highly random as well as sparse in the real-world. The experiments demonstrate that the distribution of probes in a sample can severely impact the estimation results, and hence it is not sufficient to specify adequate penetration levels with a single value.

One of the drawbacks of the proposed approach is: as the traffic state predicted at each time-step becomes input vector for the following estimation model, errors in the state estimates propagate with time. This drawback can be addressed using spatio-temporal graphical representations, but the factor graphs in such setting along with the associated learning models can be very challenging from a computational stand-point. The model can also be extended to multi-lane roads, and the CRF models can be improved by adopting higher order Markov models to capture the influence of vehicles further downstream (ahead of the leader), which could yield higher estimation accuracy with lower probe penetration rates.

Acknowledgment

This work was funded in part by the C2SMART Center, a Tier 1 USDOT University Transportation Center.

References

- X. Ban, R. Herring, P. Hao, and A. Bayen. Delay pattern estimation for signalized intersections using sampled travel times. *Transportation Research Record: Journal of the Transportation Research Board*, (2130):109–119, 2009.

- X. Ban, P. Hao, and Z. Sun. Real time queue length estimation for signalized intersections using travel times from mobile sensors. *Transportation Research Part C: Emerging Technologies*, 19(6):1133–1156, 2011.
- H. Bar-Gera. Evaluation of a cellular phone-based system for measurements of traffic speeds and travel times: A case study from israel. *Transportation Research Part C: Emerging Technologies*, 15(6):380–391, 2007.
- C. Bishop. *Pattern recognition and machine learning*. springer, 2006.
- C. Bucknell and J. Herrera. A trade-off analysis between penetration rate and sampling frequency of mobile sensors in traffic state estimation. *Transportation Research Part C: Emerging Technologies*, 46:132–150, 2014.
- H. Chen and H. Rakha. Prediction of dynamic freeway travel times based on vehicle trajectory construction. In *15th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pages 576–581. IEEE, 2012.
- W. Deng, H. Lei, and X. Zhou. Traffic state estimation and uncertainty quantification based on heterogeneous data sources: A three detector approach. *Transportation Research Part B: Methodological*, 57:132–157, 2013.
- D. Dilip, N. Freris, and S.E. Jabari. Sparse estimation of travel time distributions using Gamma kernels. *Proceedings of the 96th Annual Meeting of the Transportation Research Board (No. 17-02971)*, 2017.
- C. Furtlehner, J. Lasgouttes, and A. de La Fortelle. A belief propagation approach to traffic prediction using probe vehicles. In *IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 1022–1027. IEEE, 2007.
- N. Goodall, B. Smith, and B. Park. Microscopic estimation of freeway vehicle positions from the behavior of connected vehicles. *Journal of Intelligent Transportation Systems*, 20(1):45–54, 2016.
- B. Hellinga, P. Izadpanah, H. Takada, and L. Fu. Decomposing travel times measured by probe-based traffic monitoring systems to individual road segments. *Transportation Research Part C: Emerging Technologies*, 16(6):768–782, 2008.
- J. Herrera and A. Bayen. Incorporation of Lagrangian measurements in freeway traffic state estimation. *Transportation Research Part B: Methodological*, 44(4):460–481, 2010.
- R. Herring, A. Hofleitner, P. Abbeel, and A. Bayen. Estimating arterial traffic conditions using sparse probe data. In *2010 13th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pages 929–936. IEEE, 2010.
- G. Hiribarren and J. Herrera. Real time traffic states estimation on arterials based on trajectory data. *Transportation Research Part B: Methodological*, 69:19–30, 2014.

- A. Hofleitner, R. Herring, P. Abbeel, and A. Bayen. Learning the dynamics of arterial traffic from probe data using a dynamic Bayesian network. *IEEE Transactions on Intelligent Transportation Systems*, 13(4):1679–1693, 2012a.
- A. Hofleitner, R. Herring, and A. Bayen. Arterial travel time forecast with streaming data: A hybrid approach of flow modeling and machine learning. *Transportation Research Part B: Methodological*, 46(9):1097–1122, 2012b.
- T. Hunter, T. Das, M. Zaharia, P. Abbeel, and A. Bayen. Large scale estimation in cyberphysical systems using streaming data: A case study with smartphone traces. *arXiv preprint arXiv:1212.3393*, pages 1–13, 2012.
- S.E. Jabari. Node modeling for congested urban road networks. *Transportation Research Part B: Methodological*, 91:229–249, 2016.
- S.E. Jabari and H. Liu. A stochastic model of traffic flow: Theoretical foundations. *Transportation Research Part B: Methodological*, 46(1):156–174, 2012.
- S.E. Jabari and H. Liu. A stochastic model of traffic flow: Gaussian approximation and estimation. *Transportation Research Part B: Methodological*, 47:15–41, 2013.
- S.E. Jabari and L. Wynter. Sensor placement with time-to-detection guarantees. *EURO Journal on Transportation and Logistics*, 5(4):415–433, 2016.
- S.E. Jabari, N. Freris, and D. Dilip. Sparse travel time estimation from streaming data. *arXiv:1804.08130*, 2018.
- E. Jenelius and H. Koutsopoulos. Urban network travel time prediction based on a probabilistic principal component analysis model of probe data. *IEEE Transactions on Intelligent Transportation Systems*, pages 436–445, 2017.
- K. Jerath, A. Ray, S. Brennan, and V. Gayah. Statistical mechanics-inspired framework for studying the effects of mixed traffic flows on highway congestion. In *American Control Conference (ACC)*, pages 5402–5407. IEEE, 2014.
- K. Jerath, A. Ray, S. Brennan, and V. Gayah. Dynamic prediction of vehicle cluster distribution in mixed traffic: A statistical mechanics-inspired method. *IEEE Transactions on Intelligent Transportation Systems*, 16(5):2424–2434, 2015.
- S. Kim and B. Coifman. Comparing INRIX speed data against concurrent loop detector stations over several months. *Transportation Research Part C: Emerging Technologies*, 49:59–72, 2014.
- D. Koller and N. Friedman. *Probabilistic graphical models: Principles and techniques*. MIT press, 2009.
- M. Lárraga, J. Del Rio, and L. Alvarez-Lcaza. Cellular automata for one-lane traffic flow modeling. *Transportation Research Part C: Emerging Technologies*, 13(1):63–74, 2005.

- P. Mazaré, O. Tossavainen, A. Bayen, and D. Work. Trade-offs between inductive loops and GPS probe vehicles for travel time estimation: A Mobile Century case study. *Proceedings of the 91st Annual Meeting of the Transportation Research Board (No. 12-2746)*, 2012.
- V. Papathanasopoulou and C. Antoniou. Towards data-driven car-following models. *Transportation Research Part C: Emerging Technologies*, 55:496–509, 2015.
- T. Seo and T. Kusakabe. Probe vehicle-based traffic state estimation method with spacing information and conservation law. *Transportation Research Part C: Emerging Technologies*, 59:391–403, 2015.
- A. Sopasakis and M. Katsoulakis. Stochastic modeling and simulation of traffic flow: Asymmetric single exclusion process with Arrhenius look-ahead dynamics. *SIAM Journal on Applied Mathematics*, 66(3):921–944, 2006.
- W. Vandenberghe, E. Vanhauwaert, S. Verbrugge, I. Moerman, and P. Demeester. Feasibility of expanding traffic monitoring systems with floating car data technology. *IET Intelligent Transport Systems*, 6(4):347–354, 2012.
- D. Work, O. Tossavainen, S. Blandin, A. Bayen, T. Iwuchukwu, and K. Tracton. An ensemble Kalman filtering approach to highway traffic estimation using GPS enabled mobile devices. In *47th IEEE Conference on Decision and Control (CDC)*, pages 5062–5068. IEEE, 2008.
- Y. Yuan, J. Van Lint, E. Wilson, F. van Wageningen-Kessels, and S. Hoogendoorn. Real-time Lagrangian traffic state estimator for freeways. *IEEE Transactions on Intelligent Transportation Systems*, 13(1):59–70, 2012.
- F. Zheng, S.E. Jabari, H. Liu, and D. Lin. Traffic state estimation using stochastic Lagrangian dynamics. *arXiv preprint arXiv:1806.02692*, 2018.