

Large scale estimation of arterial traffic and structural analysis of traffic patterns using probe vehicles

Aude Hofleitner* Ryan Herring[†] Alexandre Bayen[‡]
Yufei Han[§] Fabien Moutarde[§] Arnaud de La Fortelle[§]

**91st Annual Meeting of the Transportation Research Board
January 22-26, 2012, Washington D.C**

Word Count:

Number of words:	6811
Number of figures:	5 (250 words each)
Number of tables:	0 (250 words each)
Total:	8061

*Corresponding Author, Department of Electrical Engineering and Computer Science, University of California, Berkeley and UPE/IFSTTAR/GRETTIA, France, aude.hofleitner@polytechnique.edu

[†]Apple Inc. Affiliation during redaction of the paper: California Center for Innovative Transportation, Berkeley CA, ryanherring@berkeley.edu

[‡]Department of Electrical Engineering and Computer Science and Department of Civil and Environmental Engineering, Systems Engineering, University of California, Berkeley, bayen@berkeley.edu

[§]Robotics Lab (CAOR), Mines ParisTech, Paris, France, Yufei.Han@mines-paristech.fr
Fabien.Moutarde@mines-paristech.fr

Abstract

Estimating and analyzing traffic conditions on large arterial networks is an inherently difficult task. The first goal of this article is to demonstrate how arterial traffic conditions can be estimated using sparsely sampled GPS probe vehicle data provided by a small percentage of vehicles. Traffic signals, stop signs, and other flow inhibitors make estimating arterial traffic conditions significantly more difficult than estimating highway traffic conditions. To address these challenges, we propose a statistical modeling framework that leverages a large historical database and relies on the fact that traffic conditions tend to follow distinct patterns over the course of a week. This model is operational in North California, as part of the *Mobile Millennium* traffic estimation platform. The second goal of the article is to provide a global network-level analysis of traffic patterns using matrix factorization and clustering methods. These techniques allow us to characterize spatial traffic patterns in the network and to analyze traffic dynamics at a network scale. We identify traffic patterns that indicate intrinsic spatio-temporal characteristics over the entire network and give insight into the traffic dynamics of an entire city. By integrating our estimation technique with our analysis method, we achieve a general framework for extracting, processing and interpreting traffic information using GPS probe vehicle data.

1 Introduction and related work

Traffic congestion has a significant impact on economic activity throughout much of the world. Accurate, reliable traffic monitoring systems, leveraging the latest advances in technology and research are essential for active congestion control. They can also be used to study large scale traffic patterns and to understand specific travel behavior, network bottlenecks, to design long term infrastructure planning and to optimize mobility.

Until recently, traffic monitoring systems have relied exclusively on data feeds from dedicated sensing infrastructure (loop detectors and radars in particular). For highway networks covered by such infrastructure systems, it has become common practice to perform estimation of flow, density or speed at a very fine spatio-temporal scale [3], using traffic flow models developed in the last decades [32, 7, 41]. Probe vehicle data has also been successfully integrated into these models [45, 43, 24, 30]. For arterials, traffic monitoring is substantially more difficult: probe vehicle data is the only significant data source available today with the prospect of global coverage in the future. The lack of ubiquity and reliability, the variety of data types and specifications, and the randomness of its spatio-temporal coverage encourage the use of both historical and real-time data to provide accurate estimates of traffic conditions on large transportation networks. The *Mobile Millennium* project [26] receives probe vehicle data from a dozen of different sources. In Figure 1, we illustrate one of the data source of the *Mobile Millennium* project: it shows a snapshot of probe measurements from San Francisco taxis collected on an arbitrary day from midnight to 7:00am (small dots) as well as a snapshot of the probe locations at 7:00am (large dots). This figure illustrates both the breadth of coverage when aggregating data over long periods of time and the limited information available at a given point in time, limiting the direct estimation of the macroscopic state of traffic at a fine spatio-temporal scale. Note that filtering algorithms are designed to limit the bias of the different sources of data. For example, we filter the measurements during which the hired status of the taxis changes. Aside from less abundant sensing compared to existing highway traffic monitoring systems, the arterial network presents additional modeling and estimation challenges. The underlying flow physics is more complex because of traffic lights (often with unknown cycles), intersections, stop signs, parallel queues, and other phenomena.

We introduce a statistical approach for real-time arterial traffic estimation from probe vehicle data, leveraging massive amounts of historical data. Statistical approaches have been proposed that rely on either a single measurement per time interval or aggregated measurements per time interval [19, 10], neither of which is appropriate in our setting since probe data on arterials is available at random times and random locations. Some researchers have examined the processing of high-frequency probe data (one measurement approximately every 20 seconds or less) [43], which allows for reliable calculation of short distance speeds and travel times. In this article, we specifically address the processing of sparse probe data where this level of granularity is not available. Finally, other approaches based on regression [36], optimization [1], neural networks and pattern matching [8] have all been proposed. None of these approaches addresses the issue of processing sparse probe data on a dense arterial network.

Besides the ability of providing real time traffic estimation, the results produced by the model can be further analyzed to provide a large scale understanding of traffic dynamics both in time and in space. Most of previous research in traffic data analysis focus on temporal dynamics of individual links (either on arterial or highways) using data-driven approaches: in [39, 44, 35], Kalman filter and its extensions, originated from the theory framework of state space linear dynamic model, are used for modeling and tracking temporal variations of traffic flows; [46, 37] use neural networks to achieve short-term non-linear prediction of traffic flows based on historic observations; finally, [40] proposes to perform traffic prediction on individual links based on clustering of temporal patterns of traffic flows, while [11] adopts a time-series



Figure 1: San Francisco taxi measurement locations, observed at a rate of once per minute. Each small dot represents the measurement of the location of a taxi, received between midnight and 7:00am, on March 29th, 2010. The large dots represent the location of taxis visible in the system at 7:00am on that day.

51 analysis (Autoregressive Moving Average) on traffic flows in order to forecast traffic states. Very
 52 little progress has been made in analyzing the temporal dynamics of *global* traffic states of an
 53 *entire large-scale road network*. We call *global* traffic state, the aggregation of the congestion
 54 states of all the link of the network. Traffic states of neighboring individual roads are often highly
 55 correlated (both spatially and temporally) and the identification of specific traffic patterns
 56 or traffic configurations is very informative. They can be used to better understand global
 57 network-level traffic dynamics and serve as prior knowledge or constraints for the design of traffic
 58 estimation and prediction platforms. The analysis of traffic patterns is also useful for traffic
 59 management centers and public entities to plan infrastructure developments and to improve the
 60 performances of the available network using large-scale control strategies.

61 This article proposes an algorithm to identify spatial configurations of traffic states over the
 62 entire network and analyze large-scale traffic dynamics from traffic state estimates produced
 63 and collected over long periods of time. We define the *network-level traffic state* as the vector
 64 of traffic states for each link of the network at a given moment in time. It is represented in the
 65 form of multi-variate data, where its dimension is proportional to the amount of links in the
 66 transportation network. In large networks, this data structure quickly becomes too big to handle,
 67 limiting the analysis in the original high-dimensional space. In machine learning, this issue is
 68 commonly addressed using dimension reduction techniques (feature extraction) to simplify the
 69 representation of the data, remove redundancies and improve the efficiency of analysis techniques
 70 such as classification. Important applications of these algorithms include image processing
 71 and natural language processing [12, 9]. In this work, we propose to use a dimensionality
 72 reduction matrix factorization technique known as *Non-negative Matrix Factorization* (NMF) [6,
 73 25] to obtain a low dimensional representation of network-level traffic states. Both the well-
 74 known *Principal Component Analysis* (PCA) method, and *Locality Preserving Projection* (LPP)
 75 technique are other examples of matrix factorization [28, 15]. However, in contrast to PCA or
 76 LPP, the NMF algorithm imposes strict non-negativity constraints on the decomposition result.
 77 This allows NMF to approximate the n -dimensional data vector by an *additive* combination of
 78 a set of learned bases. This property also leads to a part-based representation of the original
 79 data. The learned bases correspond to *latent components* of the original data so that the
 80 original data is approximated by a linear *positive* superposition of the latent components. The
 81 properties of the NMF have already been exploited for various applications. In text analysis,
 82 the learned bases are used to label different latent topics contained in text documents. In face
 83 image representation, the NMF bases indicate important localized components of the face, such
 84 as the eyes, the mouth or the cheeks. We expect that the distinctive characteristics of NMF
 85 will lead to a low-dimensional representation of network-level traffic states that exhibits global
 86 configurations of local traffic states and reflects intrinsic traffic patterns of network-level traffic

87 states.

88 The rest of this article is organized as follows. In Section 2 we present the real-time traffic
89 estimation algorithm implemented in the *Mobile Millennium* system. It processes sparsely
90 sampled probe vehicles sending their location at random places and random times and leverages
91 historical data using a Bayesian update. In Section 3, we introduce the NMF algorithm, used
92 in the remainder of the article to perform large scale analysis of the dynamics of traffic. In
93 Section 4 and 5, we illustrate and provide a detailed analysis of typical spatial configuration
94 patterns of network-level traffic states found by NMF projections. Section 6 further analyzes
95 temporal dynamic patterns of the network-level traffic state, which describe evolutions of traffic
96 states in the whole network. In Section 7, we conclude our work and discuss our future plans.

97 **2 Large scale statistical model for arterial traffic estima-** 98 **tion**

99 We propose a parametric statistical model for large scale traffic estimation from sparsely sampled
100 probe vehicles. The parameters of the model represent traffic patterns that are learned from
101 massive amounts of probe vehicle travel times collected over long periods of time (section 2.1).
102 The historic patterns are used as prior information in a Bayesian real-time estimation algorithm
103 with streaming data (section 2.2). The statistical model is based on assumptions aimed at
104 limiting the computational complexity of the algorithm while providing an adapted framework
105 for arterial traffic estimation when little data is available in real-time but large quantities of
106 historical data are collected over time.

- 107 1. The travel time on a link is a *random variable* (RV) and we assume that travel times
108 on different links are independent RVs. In this article, we assume that travel times are
109 normally distributed. Other distributions can also be used (e.g. Gamma, log-normal,
110 and so on) without modifying the basic concepts of the algorithm. We will specify the
111 equations that require modification under a non-normality assumption.
- 112 2. Any given moment in time belongs to exactly one *historic time period*, characterized by a
113 day of the week, a start time and an end time. The set of historic time periods is denoted
114 \mathcal{T} .
- 115 3. All travel time observations from a specific link l are independent and identically dis-
116 tributed within a given (historic) time period, $t \in \mathcal{T}$.
- 117 4. Probe vehicles send their location periodically (typically every minute). A trajectory re-
118 construction algorithm [27] provides the most likely path p of the vehicle between successive
119 location reports. The path p is defined as a set of consecutive links, L_p , along with the
120 fraction of the first and last link traversed and the total travel time associated with the
121 entire path, y_p (time between successive location reports). The fraction of link l traversed
122 for the p th observation on link l is denoted $w_{p,l}$. The time spent on link l is denoted $x_{p,l}$
123 and have the constraint that $\sum_{l \in L_p} x_{p,l} = y_p$. The set of path observations for time period
124 t is denoted \mathbf{P}_t . We assume that these path observations constitute the only data available
125 to the model.
- 126 5. Each link of the network has a known minimum travel time b_l , found by considering the
127 travel time that results from driving some percentage over the speed limit for the entire
128 link. Note that the maximum travel time is harder to determine as travel times increase
129 with congestion. A statistical analysis of available measurements may provide information
130 about maximum travel times.

2.1 Learning historic traffic patterns

The historical model of arterial traffic estimates the parameters $Q_{l,t}$ of the travel time distribution on each link l for each historic time period t . The corresponding *probability density function* (PDF) of travel times is denoted $g_{l,t}(\cdot)$. In the case of Gaussian distributions, the parameters $Q_{l,t}$ are written $Q_{l,t} = (\mu_{l,t}, \sigma_{l,t})$, where $\mu_{l,t}$ and $\sigma_{l,t}$ represent the mean and the standard deviation of the travel times on link l for the period t .

For the p th path observation, the path travel time distribution is denoted $g_{L_p,t}(\cdot)$. Under assumption 1, the PDF of travel time on a path is computed as the convolution of the PDF of travel times of the links that make up the path.

The historic algorithm determines the values of $Q_{l,t}$ for each link and time period that are most consistent with the probe data received. This is achieved by maximizing the log-likelihood of the data given the parameters, which is written as

$$\arg \max_{\mathbf{Q}_t} \sum_{p \in \mathbf{P}_t} \ln(g_{L_p,t}(y_p)), \quad (1)$$

where \mathbf{Q}_t is the set of $Q_{l,t}$ for all links l of the network. This optimization problem may be challenging due to the high number of variables (number of links times number of parameters per link travel time distribution), coupled through the PDF of path travel times $g_{L_p,t}$. In the case of Gaussian link travel times, solving (1) amounts to simultaneously estimating the mean and the variance of every link in the network and is not formulated as a convex problem. To face this difficulty, we decouple the optimization into two separate subproblems (*travel time allocation* [16, 20] and *parameter optimization*), each of which is easier to solve on its own, and then iterate between these subproblems until converging to an (locally) optimal solution.

If we knew how much time each probe vehicle drove on each link of its path (instead of just the total travel time), it would be easy to estimate the mean and standard deviation for each link in the network (sample mean and standard deviation of the link travel time observations). Since the sampling scheme only provides the total travel time on the path, we determine the most likely amount of time spent on each link (travel time allocation). Unfortunately, the most likely link travel times depend upon the link travel time parameters (μ and σ) that need to be estimated. This would appear to a chicken-and-egg problem, but there is a sound mathematical justification (hard EM) for iterating between these two steps. The link parameters are used to determine the most likely travel times and then the most likely travel times are used to update the parameters.

Travel Time Allocation: To solve the travel time allocation problem, we assume that estimates of the link parameters $Q_{l,t}$ are available (and fixed). We specify *lower bounds* b_l on the travel time allocated for each link l of the network to model the bounded speed of vehicles and ensure that a sensible solution is returned. For each observation $p \in \mathbf{P}_t$, we maximize the log-likelihood of the travel times $x_{l,p}$ spent on each link l of the path, with the constraint that they sum to the path travel time y_p . The optimization problem reads

$$\begin{aligned} \arg \max_x \quad & \sum_{l \in L_p} \log \mathcal{N}(x_{p,l}; w_{p,l} \mu_{l,t}, w_{p,l} \sigma_{l,t}) \\ \text{s.t.} \quad & \sum_{l \in L_p} x_{p,l} = y_p \\ & x_{p,l} \geq w_{p,l} b_l, \forall l \in L_p, \end{aligned} \quad (2)$$

where the minimum travel time $w_{p,l} b_l$ is found using the minimum travel time b_l on link l , scaled by the fraction of link traveled $w_{p,l}$, as introduced in item 5. The notation $\mathcal{N}(x; \mu, \sigma)$ represents the PDF of a Gaussian variable with mean μ and standard deviation σ , evaluated at x :

$$\mathcal{N}(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right).$$

168 Problem (2) is a (small scale) quadratic program (QP) [4] which can be solved analytically (see
 169 algorithm 1 for details). Note that if a vehicle travels faster than the maximum speed, the
 170 allocation problem is infeasible. The vehicle is considered as an outlier and the observation is
 171 discarded from the set of observations. We call $\mathbf{X}_{l,t}$, the vector of allocated travel times for link
 172 l during time period t . Note that the travel times $x_{p,l}$ are scaled by the proportion of the link
 173 traveled, $w_{p,l}$, before being added to the set of allocated travel times $\mathbf{X}_{l,t}$.

174 **Parameter Optimization:** Given $\mathbf{X}_{l,t}$, the computation of the parameters $Q_{l,t}$ depends
 175 on the choice of the class of distribution chosen. In the case of Gaussian distributions, this com-
 176 putation is straightforward as $\mu_{l,t}$ and $\sigma_{l,t}$ respectively represent the sample mean and standard
 177 deviation of $\mathbf{X}_{l,t}$.

178 **Full Historic Arterial Traffic Algorithm:** After initializing the parameters $Q_{l,t}$ for
 179 each link of the network, the algorithm iterates between allocating the travel times for each
 180 path in \mathbf{P}_t and optimizing the link parameters given the allocated travel times in $\mathbf{X}_{l,t}$. The
 181 convergence of the algorithm is checked by computing the log-likelihood after each iteration,
 182 which is guaranteed to increase at each iteration until convergence.

185 2.2 Bayesian Real-time Traffic Estimation

186 The parameters $Q_{l,t}$ learned by the historic model are used as prior information to estimate
 187 current traffic conditions via a *Bayesian update* (see [38] for more details on Bayesian statistics).
 188 In Bayesian statistics, parameters are considered as RVs and thus have a probability distribution.
 189 Here, we compute the probability distribution (known as *posterior distribution*) of the mean
 190 travel time (seen as a RV) given the allocated link travel times and a prior distribution on the
 191 mean travel time denoted f_0 .

192 Let Δ_t represent the duration between successive real time estimates. We run the algorithm
 193 at time t_2 using the path data available for the *current time window* $[t_1, t_2]$, with $t_1 = t_2 - \Delta_t$.
 194 The duration Δ_t between successive updates depends upon the amount of data available in
 195 real-time and should remain inferior to the duration of the historical time intervals. If the data
 196 volume is large, the model can be run up to every 5 minutes. Running the model more frequently
 197 will likely not increase the performance and may lead to estimates that fluctuate too much due
 198 in particular to the periodic dynamics associated with the presence of traffic signals [22, 23].

199 For generic RVs X and Y with realization x and y , the notation $f(x|y)$ is read “probability
 200 that X has the realization x given that Y has the realization y ” and denotes the conditional
 201 probability of RV X given the observation of the RV Y . Let y_{l,t_2} denote the set of travel times
 202 allocated to link l between t_1 and t_2 . Using Bayes theorem, the posterior probability on the
 203 mean travel time $\hat{\mu}_{l,t_2}$ is proportional to the likelihood of the data times the prior:

$$f(\hat{\mu}_{l,t_2}|y_{l,t_2}, \sigma_{l,t}) \propto f(y_{l,t_2}|\hat{\mu}_{l,t_2}, \sigma_{l,t})f_0(\hat{\mu}_{l,t_2}), \quad (3)$$

204 The symbol \propto is read “is proportional to”. The proportionality constant is chosen such that
 205 the integral of $\hat{\mu}_{l,t_2} \mapsto f(\hat{\mu}_{l,t_2}|y_{l,t_2}, \sigma_{l,t})$ on \mathbb{R} is equal to one. At time t_2 , the Bayesian update
 206 determines the value of mean travel times $\hat{\mu}_{l,t_2}$ that maximizes the posterior probability.

207 Assuming Gaussian link travel times, a natural choice for f_0 is a Gaussian distribution
 208 (conjugate prior [38]). Since f_0 represents prior information on the mean travel time, its mean
 209 is set to the historical mean $\mu_{l,t}$ and its standard deviation $\sigma_{0;l,t}$ is chosen to represent how
 210 much real-time condition can deviate from the historical values. Typically $\sigma_{0;l,t}$ is large to give
 211 more weight to real-time data as soon as they are in sufficient quantity. Because of the Gaussian
 212 prior, the allocated link travel times y_{l,t_2} and the mean travel time $\hat{\mu}_{l,t_2}$ are jointly Gaussian
 213 and we compute the parameters of the posterior (Gaussian) distribution of $\hat{\mu}_{l,t}$. In particular,
 214 we update the mean link travel times as the mean of the posterior distribution:

Algorithm 1 Travel time allocation algorithm. The core of the algorithm is contained in lines 11-15, which computes the total expected path variance (V) and the difference between expected and actual travel times (Z). With these two quantities, each link is allocated the expected link travel time adjusted by some proportion of Z , where this proportion is computed using the link variance divided by the total path variance. This procedure can lead to some links being allocated a travel time below the minimum for that link. The set \mathbf{J} is introduced to track the links with initial allocated travel times below the lower bound and the main procedure is repeated by setting the travel times for these links to the lower bound and optimizing with respect to the remaining links. Note that the travel times are scaled by the proportion of the link traveled (line 19) before being added to the set of allocated travel times $\mathbf{X}_{l,t}$.

Require: $t \in \mathcal{T}$ is fixed to some particular time period.

```

1: for  $l \in \mathcal{L}$  do
2:    $\mathbf{X}_{l,t} = \emptyset$  {Initialize allocated travel time sets to be empty.}
3: end for
4: for  $p \in \mathbf{P}_t$  do {For all probe path observations.}
5:   if  $\sum_{l \in L_p} w_{p,l} b_l > y_p$  then
6:     Travel time allocation infeasible for this path. This means that the observation represented
       travel that is considered faster than realistically possible, so the observation is considered
       an outlier. Remove  $p$  from  $\mathbf{P}_t$ .
7:   else
8:      $\mathbf{J} = \emptyset$  { $\mathbf{J}$  contains all links for which the travel time allocation is fixed to be equal to the
       lower bound.}
9:     repeat
10:       $x_{p,l} = w_{p,l} b_l, \forall l \in \mathbf{J}$  {For all links that had an infeasible allocation in the previous pass
        through this loop, set the allocation to the lower bound.}
11:       $V = \sum_{l \in L_p \setminus \mathbf{J}} w_{p,l} \sigma_{l,t}^2$  {Calculate the path variance for the links not fixed to the lower
        bound.}
12:       $Z = y_p - \sum_{l \in \mathbf{J}} w_{p,l} b_l - \sum_{l \in L_p \setminus \mathbf{J}} w_{p,l} \mu_{l,t}$  {Calculate the difference between expected and actual
        travel time for the links not fixed to the lower bound.}
13:      for  $l \in L_p$  do {Allocate excess travel time in proportion of link variance to path vari-
        ance.}
14:         $x_{p,l} = w_{p,l} \mu_{l,t} + \frac{w_{p,l} \sigma_{l,t}^2}{V} Z$ 
15:      end for
16:       $\mathbf{J} = \mathbf{J} \cup \{l \in L_p : x_{p,l} < w_{p,l} b_l\}$  {Find all links violating the lower bound.}
17:    until  $x_{p,l} \geq w_{p,l} b_l, \forall l \in L_p$ 
18:    for  $l \in L_p$  do
19:       $\mathbf{X}_{l,t} = \mathbf{X}_{l,t} \cup \left( \frac{x_{p,l}}{w_{p,l}} \right)$  {Add the allocated travel time to  $\mathbf{X}_{l,t}$ .}
20:    end for
21:  end if
22: end for
23: return  $\mathbf{X}_{l,t}, \forall l \in \mathcal{L}$ 

```

$$\hat{\mu}_{l,t} = \frac{\sigma_{0;l,t}^2}{\frac{\sigma_{l,t}^2}{N_{l,t_2}} + \sigma_{0;l,t}^2} \bar{x} + \frac{\sigma_{l,t}^2}{\frac{\sigma_{l,t}^2}{N_{l,t_2}} + \sigma_{0;l,t}^2} \mu_{l,t}$$

where N_{l,t_2} is the number of travel times allocated to link l during the current time interval (t_1, t_2) and \bar{x} is the sample mean of the allocated travel times y_{l,t_2} .

To summarize, the real-time estimation algorithm performs the travel time allocation on each probe observation and then uses the allocated travel times and the historical traffic parameters to perform a Bayesian update of the link parameters.

The precise analysis of the performance of this model is out of the scope of this article. We refer the reader to the following references assessing the results of the *Mobile Millennium* project for more details [2, 17].

3 Non-negative matrix factorization (NMF)

In this section, we present *Non-negative Matrix Factorization* (NMF), which is used for approximating network-level traffic states as positive sums of a limited number of global traffic configurations. NMF [31, 6, 33, 25, 9] is a particular type of matrix factorization, in the same domain as the well-known *Principal Component Analysis* (PCA) method and *Locality Preserving Projection* (LPP). In all cases, given a set of multivariate n -dimensional data vectors placed in m columns of a $n \times m$ matrix X , matrix factorization decomposes the matrix into a product of a $n \times s$ loading matrix M and a $s \times m$ score matrix V , where s represents the dimensionality of the subspace to which we project the original data. Through this matrix decomposition, each n -dimensional data vector is approximated by a linear combination of the s columns of M , weighted by the components in the corresponding column of V . We can regard all s column vectors in loading matrix M as a group of projection bases that are learned optimally to represent the original data. The variable s is typically chosen to be significantly smaller than both n and m so that the obtained score matrix V forms a low-dimensional subspace projection of the network-level traffic states, on which we can perform further data analysis. The specificity of NMF is the enforced positivity of both the weights in V , and of the columns of M forming the NMF decomposition basis. This non-negativity therefore provides an approximation of the n -dimensional data vector by an *additive* combination of a set of learned bases. Furthermore, the NMF components forming the basis tend to be sparse, which leads to a part-based representation of the original data.

A network-level traffic state is a vector of size equal to the number of links in the network, where the i^{th} entry corresponds to the traffic state on the i^{th} link of the network. Arterial networks are typically dense (numerous links and intersections) and the number of links in any decent size network is often over a thousand links. Assuming that k samples of n -dimensional network-level traffic states are stored as an $n \times k$ matrix X , NMF factorizes X as a product of a non-negative $n \times s$ matrix M and a non-negative $s \times k$ matrix V which minimizes the Frobenius norm of the reconstruction error between X and its factorized approximation MV . We recall that the Frobenius norm of a matrix $A \in \mathbb{R}^{n \times m}$ with entry on column i and line j denoted $A_{i,j}$ is defined as

$$\|A\|_F = \sqrt{\sum_{j=1}^n \sum_{i=1}^m |a_{i,j}|^2}$$

and is equal to the sum of the singular values of A . The matrix factorization problem reads:

$$\arg \min_{(M,V)} \|\mathbf{X} - \mathbf{M}\mathbf{V}\|_F \quad \text{s. t.} \quad M \geq 0, V \geq 0, \quad (4)$$

where the inequalities $M \geq 0$, $V \geq 0$ represent the non-negativity constraints (each element of the matrices are non-negative). Training of NMF is implemented using multiplicative updates [31], fixing either M or V and updating the left following the KKT condition. The NMF cost function shown in equation (4) is not convex. However, fixing either M or V leads to a convex subproblem to solve. Multiplicative updates and other gradient based optimization procedure can not guarantee the global optimum of the NMF solution. Nevertheless, in data mining, local minimum is still enough to be useful. Given fixed M or V , the NMF objective is a convex optimization issue. The NMF projects the high-dimensional network-level traffic states on a s -dimensional subspace, which is spanned by the columns of M . According to equation (4), the column space of V corresponds to coordinates of network-level traffic states with respect to the learned set of bases in M . The column space of V forms a low-dimensional representation of the network-level traffic states. As mentioned in the introduction, each network-level traffic state $X_j \in \mathbb{R}^n$ is approximated by an additive linear superposition of the column space of M due to the non-negative constraint. The approximation of X_j is written

$$X_j \approx \sum_{i=1}^k M_i V_{i,j}, \quad (5)$$

where M_i denotes the i^{th} column of M and $V_{i,j}$ is the element at the i^{th} column and j^{th} row of V . It is important to interpret what the matrices M and V represent in terms of traffic analysis. The column space of M represents typical elements of the spatial configuration patterns with respect to the network-level traffic states. Based on the columns of M , we represent complex spatial arrangements of local traffic states over the entire network. As for V , equation (4) indicates that each element $V_{i,j}$ represents to which degree the j^{th} network-level traffic state observation is associated with the i^{th} expanding basis in matrix M (i^{th} spatial configuration). For example, if the spatial configuration formed by the i^{th} column of M is the best representation of the j^{th} network-level traffic state, then $V_{i,j}$ will take the largest value in the j^{th} row of V [6]. As a result, the derived low-dimensional representation formed by the columns space of V are intuitively consistent with information about spatial distribution patterns of local traffic states. By contrast, the PCA and LPP based projections only aim at best reconstruction of traffic observations with either maximizing data variances or preserving neighboring structures. The projection results of PCA and LPP are thus less likely to be associated with interpretable latent traffic configuration patterns than the NMF. Therefore, we choose NMF to analyze the network-level traffic states in our case.

In this article, the traffic states used for the clustering analysis are *fluidity indices*. A fluidity index is a value in $[0, 1]$ computed as the ratio between the free flow and the estimated travel times. They are provided by the estimation algorithm described in Section 2 and operational in the *Mobile Millennium* [26] traffic platform, which receives data from a dozen of feeds totalling several millions of data points per day for Northern California traffic. The *Mobile Millennium* platform has been operational since November 2008 and has been storing historical data since then, providing a rich database of historical traffic dynamics in the Bay Area. In real-time, the model estimates travel times and fluidity indices from the streaming data and leverages the historical data using the Bayesian update presented in Section 2.2. The estimates are updated on each link of the network every five minutes. We focus our study on a network consisting of 2626 links for a duration of 184 days, from 00:00 May 1st 2010 to 23:55 October 31st 2010, totaling 52292 estimates per link ($12 \times 24 \times 184$). We store the fluidity index of each link at each time sampling step in a matrix X containing 2626 rows and 52292 columns. Our clustering results includes two parts, firstly we perform clustering on network-level traffic states, in order to find some typical spatial configurations of network-level traffic states, as described in Section 4. Secondly, we perform clustering on temporal trajectories of network-level traffic states, from which we can study traffic dynamics, shown in Section 6.

4 Congestion patterns: spatial configurations of global traffic states

An important outcome of dimensionality reduction is to identify typical spatial congestion patterns (i.e. spatial configurations of congestion). While doing this on the original 2626 dimensional data would be rather sloppy and computer intensive, it is much more feasible in the low dimensional space obtained by NMF.

NMF has one essential parameter: the number s of components over which decomposition is done. The parameter s also corresponds to the dimension of the target subspace where we perform clustering. The choice of s is empirical (s is called a *meta parameter*) and is done by analyzing results obtained for increasing values of s from 3 to 30. Our analysis focuses on the reconstruction error (value of the objective function (4) at optimum) and the clustering results. The reconstruction error continually decreases as the dimension s increases. This result is expected as the optimization problem (4) is performed on a larger set and thus the factorization models with higher complexity always leads to better fitting to the original data. Our clustering of global traffic states consists of clustering the traffic data projected in the s -dimension subspace using a *k-means* algorithm [34, 29]. The *k-means* algorithm is a widely used unsupervised clustering algorithm. It partitions observations into k clusters in which each observation belongs to the cluster with the nearest mean. We represent the clusters obtained in the s -dimensional space in three dimensions, limiting the number of NMF components to three but keeping the clustering results obtained in the s -dimensional space. We notice that values of s inferior to eleven lead to clustering results which seem visually inadequate: the 3D representation of the clusters shows important overlap between the clusters. The clusters become separated for values of s greater than fifteen. Increasing s over 15 does not seem to bring any improvement in the clustering results, while it significantly increases the NMF computation and memory usage costs. Therefore, we set the number of NMF components to $s = 15$ for all subsequent analysis presented in this article. This value achieves a balance between the descriptive power of NMF projection and the computational efficiency.

In clustering analysis, we also need to choose the number of clusters in *k-means*, denoted by k . The challenge is different from that for the choice of s : the choice of k does not influence the computational costs significantly but changes the interpretability of the results. The number of clusters represents the number of “global congestion patterns” that may arise. Too low values of k may not represent the different congestion patterns whereas too high values of k may decrease the possibilities of interpretation by separating similar congestion states into different clusters. After analyzing the results obtained for increasing values of k , it seems that the most insightful clustering is obtained with $k = 5$ clusters. The average fluidity index value (obtained by averaging index values on all links) are shown for each of the five clusters in the table at the top of figure 2. It appears that two clusters (cyan squares and green stars) correspond to different types of “mostly fluid” states, whereas the remaining three clusters (blue circles, yellow diamonds and red stars) represent “congested states”. We study the physical significance of each cluster by constructing histograms of the fluidity index values, counting occurrence frequencies of fluidity index values in each cluster. We find that fluidity index values in the *night and early morning Free-Flow* (NFF) and *Evening Free-Flow* (EFF) cluster are higher as a whole than those in the clusters corresponding to occurrences of congestion (*Morning Increasing Congestion*, MIC, *Mid-Day Congestion*, MDC and *Afternoon Decreasing Congestion* ADC clusters).

Figure 2 shows that the significance of the distributional patterns with respect to evaluating global traffic states is generally consistent with that of average fluidity index values, which implies that the average fluidity index value could also be used as an easy-to-use and efficient indicator of global traffic states in our case.

Marker symbol	Average fluidity	Cluster name
Green stars	0.7757	Night + early morning Free-Flow (NFF)
Blue circles	0.7185	Morning Increasing Congestion (MIC)
Red stars	0.6393	Mid-Day Congestion (MDC)
Yellow diamonds	0.6730	Afternoon Decreasing Congestion (ADC)
Cyan squares	0.7420	Evening Free-Flow (EFF)

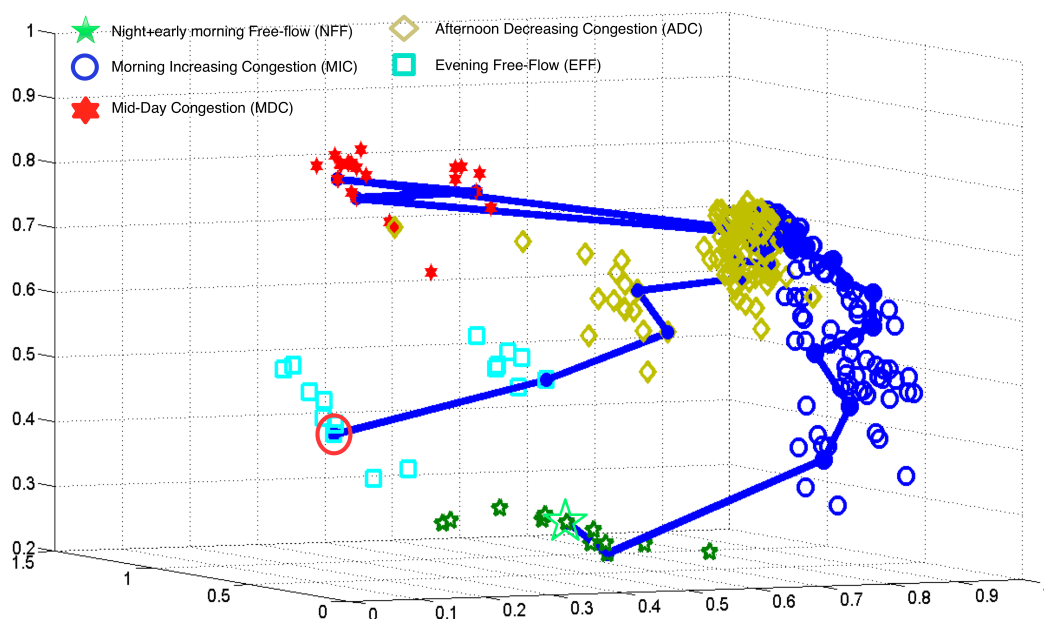





Figure 2: The clustering shows an organization of global congestion states per time of the day. The table shows the average fluidity values of each of the global state clusters. The figure shows the temporal evolution of global congestion states, projected in the 3D-NMF space using different colors and symbols to represent the five different clusters. The first and the last network estimates of the day are represented with a large star and a large circle respectively.




(a) Night Free-Flow (NFF)



(b) Morning Increasing Congestion (MIC)



(c) Evening Free-Flow cluster (EFF)



(d) Afternoon Decreasing Congestion (ADC)



(e) Mid Day Congestion (MDC)

Figure 3: Typical spatial configurations of traffic states for each of the five clusters. On each figure, we display the links with fluid index values less than 0.7 (congested links). Most of the congestion occurs within the region highlighted by the dashed circle, which is the downtown region of San Francisco. The NFF and EFF clusters have a smaller number of links highlighted than the MIC, ADC and MDC clusters indicating the difference in congestion levels.

343 As done in the primary analysis for the choice of s and for visualization purposes, we illustrate
344 spatial layouts of the global traffic state distribution in 3D-NMF space (obtained by requesting
345 3 components only instead of 15), but we apply the k -means clustering algorithm in the larger
346 15-D NMF space. The physical interpretation of the five clusters is clear in Figure 2 in which
347 we show all states projected in 3D-NMF, together with a typical temporal evolution trajectory
348 of a single day. The whole trajectory is indicated by the blue line in Figure 2. The green star
349 and red circle are the starting point and ending point of the trajectory, corresponding to traffic
350 observations at 00:00 and 23:55 respectively. The temporal arrangements of the network-level
351 traffic states along the trajectory underline the temporal interpretation of the five clusters: the
352 green-star cluster corresponds to night and early morning free-flowing, from which typical day
353 evolution goes into morning intermediate states (before 10:00) corresponding to the blue-circle
354 cluster; mid-day congestion (red-star cluster) generally occurs between 10:05 and 15:00, and
355 represents a clearly different congestion state in 3D-NMF space, with a sudden jump of traffic
356 states from the blue-circle cluster to the red-star one, and sudden jump back into the after-
357 noon intermediate state (yellow-diamond cluster) around 15:00. The traffic settles to a specific
358 evening near-free-flow state from 18:00 to 23:55 (cyan-square cluster). Interestingly, both the
359 projection of the global congestion states in 3D-NMF space and the clustering results in 15D-
360 NMF show a clear distinction between morning and afternoon intermediate congestion states,
361 and also between late evening and night/early-morning near-free-flow states.

362
363 In Figure 3, we show traffic patterns corresponding to spatial configurations of congestion
364 for centers of each of the five identified clusters. Each cluster center is derived by averaging all
365 elements of the corresponding cluster, so as to indicate a representative spatial configuration of
366 traffic states of each cluster. In this figure, we display the links with fluidity index values less
367 than 0.7 (congested links) on the Google Map screenshots. Generally, most of congestion occurs
368 within the regions highlighted by the dashed circle in figure 3(e). This region corresponds to
369 the downtown region of San Francisco. Compared with the downtown region, the western and
370 southern region of San Francisco are less likely to suffer from congestion (left and bottom part
371 in the San Francisco road network). This analysis is very useful for traffic management cen-
372 ters and public entities to understand the most important bottlenecks that cause heavy traffic
373 conditions. Moreover, our results show that some of the major bottlenecks remain constant
374 throughout the day whereas others evolve with the different traffic patterns of the day. This
375 dynamical analysis can lead to specific management strategies to address this recurring conges-
376 tion. As a matter of fact, in [13], we constructed a regression model to predict the global traffic
377 dynamics based on the analysis results of the spatial congestion patterns. This work indicates
378 the promising potentials of spatial congestion patterns in forecasting congestion and improving
379 traffic management.

380 5 Spatial decomposition of the road network

381 Another motivation for using NMF in dimensionality reduction is its property to approximate
382 original data by an additive linear combination of a limited set of “components” (a.k.a. NMF
383 “basis”). Due to the non-negative constraints, spatial arrangements of the components are usu-
384 ally sparse, which means that values in most regions of each basis are (close to) zero except
385 several localized regions. These localized regions with large values correspond to typical pat-
386 terns or representative components of the original signals (the global congestion states), and
387 typically correspond to independent “parts” of the data. Therefore, NMF is often used to ex-
388 tract part-based representation or latent semantic topics from the data in image processing or
389 text classification. For example, when NMF is applied to image datasets, it automatically ex-
390 tracts some part-based representation of the type of objects present in the images [31, 25, 9].
391 We study this “part-based” representation of global congestion states to analyze the physical

392 significance of NMF components obtained on traffic data.

393 For arterial traffic, the localized regions with distinctively large values in each NMF basis
394 correspond to a group of links with highly correlated traffic states. In this section, we construct
395 the localized components by selecting the links which represent the top 20% largest values in
396 each basis and indicate their spatial locations using red legends in the road network. Figure 4
397 shows several typical spatial arrangements of localized components, out of the fifteen arrange-
398 ments learned during the NMF training. We notice that a component corresponds to streets in
399 a localized West region (Figure 4(a)), and another to streets in the central region (Figure 4(b)),
400 which could indicate that the traffic within each of these regions is highly correlated with each
401 other whereas the traffic between distinct regions exhibits relatively independent behaviors.
402 Such a characterization of independent regions of traffic dynamics is important to significantly
403 reduce the computational costs of a large variety of estimation models, in particular estimation
404 models based on graphical models [10, 18]. We could leverage this characterization in approx-
405 imate inference algorithms to reduce the computational costs while maintaining an accurate
406 representation of traffic dynamics and limiting the estimation error [5].


407 Other NMF components highlight correlations of traffic in parallel directions: in Figure 4(c),
408 a majority of the links of the NMF component are horizontally-oriented, whereas in Figure 4(d) a
409 majority of the links are vertically-oriented. As we highlight in Figure 4(c) and Figure 4(d), the
410 links concentrated within the downtown tend to be more consistent with the orientational pat-
411 terns. These links with similar orientations are likely to have correlated traffic dynamical be-
412 haviors, whereas traffic flows with orthogonal orientations have a less important impact on each
413 other. These correlations properties can be used to learn the structure of the graphical model
414 representing conditional independences between traffic states on the network (both spatially and
415 temporarily).

416
417 According to the physical representation of the NMF components, it seems that different
418 NMF bases focus on different localized connected regions of the network. This could imply
419 that NMF detects both strong correlation of traffic dynamics within each localized region and
420 relative independence between these regions. However, this connectivity and localization of the
421 components could be improved. Standard NMF does not guarantee connected nor localized
422 components and the above promising results motivate us to investigate this physical represen-
423 tation of spatial configuration of traffic states further. A possible approach is to modify the
424 NMF algorithm in order to favor *localized* sparsity, which should help to unveil more distinct
425 part-based network decomposition.


426 6 Temporal analysis of global traffic states

427 In this section, we analyze the daily dynamics of network-level congestion states projected in
428 the NMF space. This analysis is important to understand how congestion forms and dissipates
429 throughout the day. For each day in the studied period, we represent the trajectory of the
430 network-level traffic states in the NMF space as the projection of the temporal sequence of the
431 network-level traffic states from the beginning to the end of the day. The projections are linked
432 together to form a solid curve representing the trajectory and we notice that trajectories are
433 nearly closed in the NMF space. Note that for visualization purposes, the projection is done
434 on the 3D-NMF space. Figure 5 (top) shows a typical day trajectory with successive temporal
435 intervals along the trajectory plotted using different colors, to give an idea of the dynamics along
436 the curve.


437
438 It is noteworthy that over the 184 days of reconstructed traffic data, there are only, in
439 3D-NMF projection, exactly seven different typical trajectories, as shown in figure 5 (center).
440 Furthermore, our analysis shows that each one of these seven typical trajectories corresponds to




(a) “West Part” NMF component



(b) “Central” NMF component



(c) “East-West transit” NMF component



(d) “North-South transit” NMF component

Figure 4: Examples of interesting NMF components, either highlighting localized behavior (a and b), or flow-direction correlations (c and d).

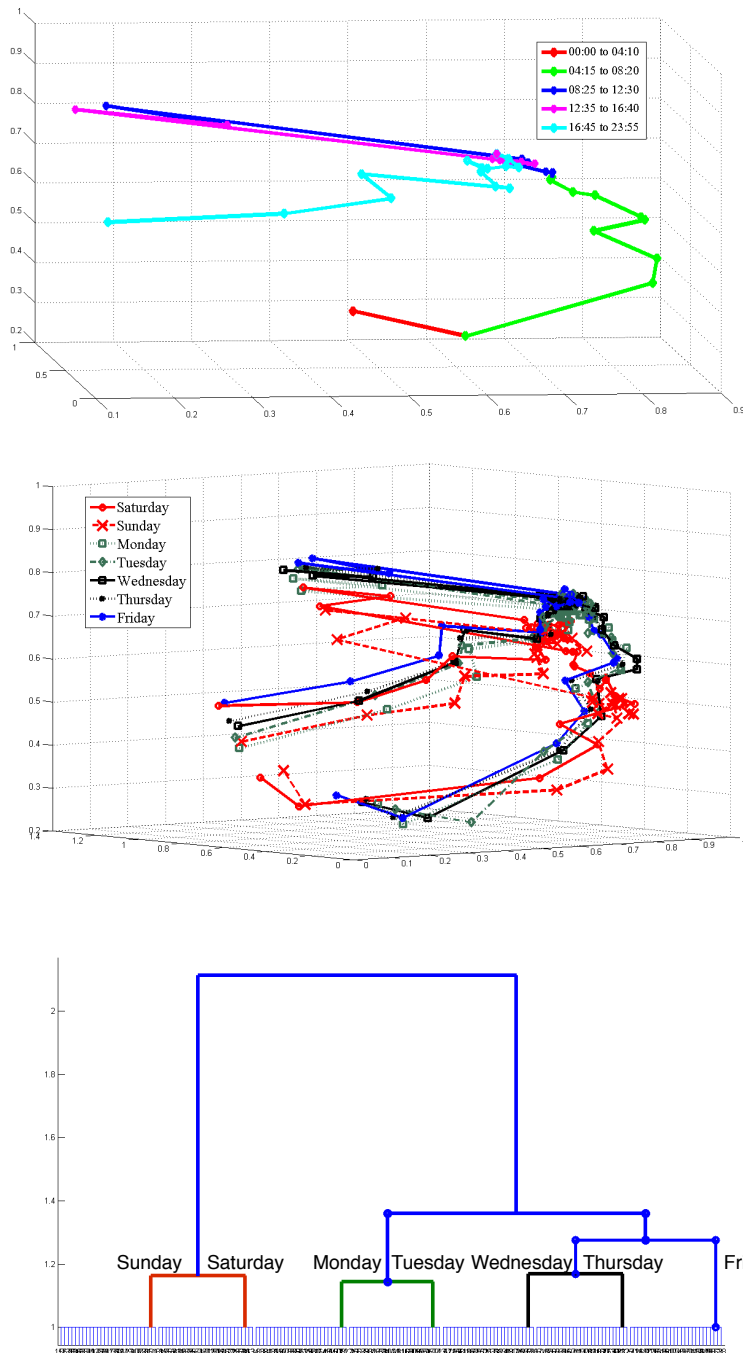


Figure 5: Daily trajectories of network fluidity indices projected in 3D-NMF space exhibit seven different typical trajectories, representing the days of the week. **Top:** Example of a daily trajectory with coloring representing the five different times of the day. **Center:** The seven different trajectories, representing a typical daily dynamic for each day of the week. **Bottom:** Dendrogram representing the hierarchical clustering analysis of the daily trajectories.

441 a particular day of the week and are thus called *day trajectory patterns*. Note that individual
 442 day trajectories for same day-of-week, although superposed in 3D-NMF, are slightly different
 443 one from another in 15D-NMF space, in which we perform clustering. Also, there are only two
 444 weekday holidays within the analyzed period, and they exhibit only small deviation from the
 445 ordinary same day-of-week. This may be a consequence of the estimation algorithms which does
 446 not use a holiday-specific model to process the historic data and fuse it with the real-time data.

447 Differences between the different day trajectory patterns concentrate within the time inter-
 448 val corresponding to transitions between congestion states, in particular between the morning
 449 increasing congestion and the mid-day congestion and between the mid-day congestion and the
 450 evening decreasing congestion. Characterizing these specific time intervals that represent the
 451 differences in daily dynamics allows us to identify and/or predict different evolution patterns of
 452 traffic states and to develop mid-term or long-term traffic forecast [14, 13].

453
 454 In this data set, one complete evolution trajectory contains 288 sampling steps (estima-
 455 tions are performed every five minutes), which is represented by a 2626×288 matrix (the
 456 network has 2626 links). As for the previous sections, our analysis is done in 15-D NMF
 457 space (3-D space is only used for visualization purposes). Each trajectory is represented by
 458 a sequence of 288 network-level traffic state projected on the 15-D NMF space and denoted
 459 $\{h_1, h_2, \dots, h_{288}\}$, where $h_i \in \mathbb{R}^{15}$. To measure similarity between trajectories $\{h_1^a, h_2^a, \dots, h_{288}^a\}$
 460 and $\{h_1^b, h_2^b, \dots, h_{288}^b\}$, representing days a and b respectively, we calculate *cosine distances* be-
 461 tween the NMF projections at corresponding times of the day and sum the cosine distances over
 462 the different estimation times $k = 1 \dots 288$:

$$D = \sum_{k=1}^{288} \text{cosdis}(h_k^a, h_k^b), \quad (6)$$

$$\text{where } \text{cosdis}(h_k^a, h_k^b) = 1 - \frac{h_k^a \cdot h_k^b}{\|h_k^a\| \|h_k^b\|}. \quad (7)$$

463
 464 The function *cosdis* is the cosine distance between two vectors and is defined in (7). It evalu-
 465 ates the cosine value of the angle between the two data vectors h_k^a and h_k^b in the 15-D NMF
 466 projection space. Larger cosine distance values indicate more important differences between
 467 the two vectors. Due to the mathematical definition of the cosine function, the derived cosine
 468 distance is normalized into the range $[0,1]$. Based on the defined distance measure between
 469 sequences, we can perform hierarchical clustering of daily traffic observation sequences in 15D-
 470 NMF space [29, 42]. The successive similarity-based groupings are shown on the dendrogram in
 471 Figure 5 (bottom) following the same color legends as in the middle figure. In the dendrogram,
 472 daily sequences of network-level traffic states are grouped gradually into clusters in the form of
 473 U-shaped trees. The height of each U-shaped tree (vertical axis) represents the distance between
 474 the sets of daily sequences being connected. Leaf nodes along the horizontal axis correspond
 475 to all daily sequences of network-level traffic states. We notice that at the bottom level of the
 476 hierarchical tree, daily sequences are first aggregated with respect to each day of the week. It
 477 underlines the intuition that each day of the week has a particular temporal dynamic pattern in
 478 terms of network-level traffic states. By increasing thresholds of distance settings, we trace back
 479 along the U-shaped trees towards its root. The seven days of the week are further clustered into
 480 four different groups indicating the days that tend to follow similar dynamic patterns. Weekend
 481 (Saturday and Sunday) are clustered together. As for the week days, Monday and Tuesday,
 482 representing the beginning of a week, appear to have a different temporal dynamic pattern from
 483 Wednesday and Thursday (middle of the week). Traffic dynamics on Fridays also tend to de-
 484 viate slightly from that of the other days and is assigned to a separate group. As the distance
 485 threshold increases, Friday is added to the Wednesday and Thursday cluster. Therefore, we can
 486 say that there are generally three kinds of temporal dynamic patterns of network-level traffic
 487 states in the data, corresponding to the beginning of the week (Monday and Tuesday), the end

488 of the week (Wednesday, Thursday and Friday) and weekends (Saturday and Sunday). If we
489 increase the threshold even more, the two clusters of week days merge leading to two clusters
490 representing the week-end days on one side and the week days on the other side. The distance
491 thresholds need to be increased significantly more for these two clusters to merge, which indi-
492 cates the importance in the differences in daily dynamics between week days and weekends. It
493 is expected for Monday and Friday to have different dynamics (coming back or leaving for the
494 week-end). However, it is slightly surprising that Monday and Tuesday are clustered together
495 while Wednesday and Thursday (and then Friday) form another cluster. The data seems to
496 indicate a beginning of the week vs. end of the week clustering, with Friday being the most
497 different of the other days.

498 7 Conclusion and discussion

499 In this article, we have proposed and presented: (1) a probabilistic modeling framework for
500 efficient estimation of arterial traffic conditions from sparse probe data; (2) a novel traffic data
501 mining approach to analyze large-scale traffic patterns and dynamics.

502 The proposed estimation method leverages massive amounts of historical data to learn sta-
503 tistical distributions of travel times and fuses them with streaming data to produce real-time
504 estimates of traffic conditions using a Bayesian update. The Bayesian framework allows us to
505 properly weight the relative importance of the real time and the historical data (depending on
506 the amount of data available in real time) to produce robust estimates, even when little data is
507 available in real-time. This model is operational in the *Mobile Millennium* system and has been
508 producing travel time and fluidity indices since March 2010 [2].

509 The output of this estimation model is used as a first real-world platform for a new traffic
510 data mining method using Non-negative Matrix Factorization to allow large-scale analysis of
511 spatial and temporal traffic patterns. The principle is to perform dimensionality reduction,
512 which allows for clustering of spatial congestion patterns, and easy analysis/categorization of
513 temporal daily dynamics. Furthermore, the part-based decomposition feature of Non-negative
514 Matrix Factorization automatically unveils areas of the road network with strong correlations.

515 Current and future research focus on: (1) integrating traffic flow theory and statistical mod-
516 els to have a more accurate modeling of traffic dynamics, both at the link [22, 23] and at the
517 network [21] level in order to improve the estimation capabilities of the system; (2) modifica-
518 tions of Non-negative Matrix Factorization sparsity constraint to favor geographically-localized
519 components; (3) taking advantage of low-dimensional Non-negative Matrix Factorization repre-
520 sentation for performing long-term traffic prediction [13].

521 Acknowledgements

522 The authors wish to thank Timothy Hunter from UC Berkeley for providing filtered probe tra-
523 jectories from the raw measurements of the probe vehicles. We thank the *California Center for*
524 *Innovative Transportation* (CCIT) staff for their contributions to develop, build, and deploy the
525 system infrastructure of *Mobile Millennium* on which this article relies. This research was sup-
526 ported by the Federal and California DOTs, Nokia, the Center for Information Technology Re-
527 search in the Interest of Society (CITRIS), the French National Research funding Agency ANR
528 (part of this work was supported by grant ANR-08-SYSC-017 for the “TRAVESTI” project),
529 and French Department of Sustainable Development and Transports MEEDDM. This collabo-
530 ration between CAOR of Mines ParisTech and Berkeley was initiated partly thanks to funding
531 MEEDDM-09-SDI-003 granted by French authorities within the CalFrance franco-californian
532 cooperation framework.

References

- 534 [1] S.J. Agbolosu-Amison, B. Park, and I. Yun. Comparative evaluation of heuristic optimization
535 methods in urban arterial network optimization. In *12th Intelligent Transportation
536 Systems Conference (ITSC '09)*, 2009.
- 537 [2] A. Bayen, J. Butler, and A. Patire et al. Mobile Millennium final report. Technical
538 report, University of California, Berkeley, CCIT Research Report UCB-ITS-CWP-2011-
539 6, To appear in 2011.
- 540 [3] P. Bickel, C. Chen, J. Kwon, J. Rice, E. Van Zwet, and P. Varaiya. Measuring traffic.
541 *Statistical Science*, 22(4):581–597, 2007.
- 542 [4] S.P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- 543 [5] X. Boyen and D. Koller. Tractable inference for complex stochastic processes. In *Proc.
544 UAI*, volume 98, 1998.
- 545 [6] D. Cai, X. He, X. Wu, and J. Han. Non-negative matrix factorization on manifold. In *The
546 proceedings of ICDM 2008*, 2008.
- 547 [7] C. Daganzo. The cell transmission model: A dynamic representation of highway traffic
548 consistent with the hydrodynamic theory. *Transportation Research B*, 28(4):269–287, 1994.
- 549 [8] C. de Fabritiis, R. Ragona, and G. Valenti. Traffic estimation and prediction based on
550 real time floating car data. In *Intelligent Transportation Systems, 2008. ITSC 2008. 11th
551 International IEEE Conference on*, pages 197–203, 2008.
- 552 [9] T. Feng, S.Z. Li, H.Y. Shum, and H.J. Zhang. Local nonnegative matrix factorization as a
553 visual representation. In *Proceedings of the 2nd International Conference On Development
554 and Learning*, 2002.
- 555 [10] C. Furtlehner, J. Lasgouttes, and A. de la Fortelle. A belief propagation approach to traffic
556 prediction using probe vehicles. In *10th Intelligent Transportation Systems Conference
557 (ITSC '07)*, pages 1022–1027, 2007.
- 558 [11] B. Ghosh, B. Basu, and M. O'Mahony. Multivariate short-term traffic flow forecasting using
559 time-series analysis. *IEEE Transaction on Intelligence Transportation Systems*, 10(2):246–
560 254, 2009.
- 561 [12] I. Guyon and A. Elisseeff. An introduction to variable and feature selection. *The Journal
562 of Machine Learning Research*, 3:1157–1182, 2003.
- 563 [13] Y. Han and F. Moutarde. Analysis of network-level traffic states using locality preservative
564 non-negative matrix factorization. In *14th IEEE Intelligent Transport Systems Conference
565 (ITSC'2011)*, 2011.
- 566 [14] Y. Han and F. Moutarde. Clustering and modeling of network-level traffic states based on
567 locality preservative non-negative matrix factorization. In *8th Intelligent Transport Systems
568 (ITS) European Congress*, 2011.
- 569 [15] X. He and P. Niyogi. Locality preserving projections. In *Proceedings of 17th Neural Infor-
570 mation Processing Systems*, 2003.
- 571 [16] B. Hellinga, P. Izadpanah, H. Takada, and L. Fu. Decomposing travel times measured
572 by probe-based traffic monitoring systems to individual road segments. *Transportation
573 Research Part C*, 16(6):768 – 782, 2008.
- 574 [17] R. Herring. *Real-Time Traffic Modeling and Estimation with Streaming Probe Data using
575 Machine Learning*. PhD thesis, UC Berkeley, Departement of Industrial Engineering and
576 Operations Research, 2010.

- 577 [18] R. Herring, A. Hofleitner, P. Abbeel, and A. Bayen. Estimating arterial traffic conditions
578 using sparse probe data. In *Proceedings of the 13th International IEEE Conference on*
579 *Intelligent Transportation Systems*, Madeira, Portugal, September 2010.
- 580 [19] R. Herring, A. Hofleitner, S. Amin, T. Abou Nasr, A. Abdel Khalek, P. Abbeel, and
581 A. Bayen. Using mobile phones to forecast arterial traffic through statistical learning. In
582 *Proceedings of the 89th Annual Meeting of the Transportation Research Board*, Washington
583 D.C., 2010.
- 584 [20] A. Hofleitner and A. Bayen. Optimal decomposition of travel times measured by probe
585 vehicles using a statistical traffic flow model. *IEEE Intelligent Transportation System*
586 *Conference (IEEE ITSC '11)*, 2011.
- 587 [21] A. Hofleitner, R. Herring, and A. Bayen. Arterial travel time forecast with streaming data:
588 a hybrid flow model - machine learning approach. *submitted, Transportation Research Part*
589 *B*, 2011.
- 590 [22] A. Hofleitner, R. Herring, and A. Bayen. A hydrodynamic theory based statistical model
591 of arterial traffic. *Technical Report UC Berkeley, UCB-ITS-CWP-2011-2*, January 2011.
- 592 [23] A. Hofleitner, R. Herring, and A. Bayen. Probability distributions of travel times on arte-
593 rial networks: a traffic flow and horizontal queuing theory approach. *91st Transportation*
594 *Research Board Annual Meeting*, January 2012.
- 595 [24] E. Horvitz, J. Apacible, R. Sarin, and L. Liao. Prediction, expectation, and surprise:
596 Methods, designs, and study of a deployed traffic forecasting service. In *Twenty-First*
597 *Conference on Uncertainty in Artificial Intelligence*, 2005.
- 598 [25] P.O. Hoyer. Non-negative matrix factorization with sparseness constraints. *Journal of*
599 *Machine Learning Research*, vol.5:1457–1469, 2004.
- 600 [26] The *Mobile Millennium* Project. <http://traffic.berkeley.edu>.
- 601 [27] T. Hunter, R. Herring, A. Hofleitner, A. Bayen, and P. Abbeel. Trajectory reconstruction
602 of noisy GPS probe vehicles in arterial traffic. In *preparation for IEEE Transactions on*
603 *Intelligent Transport Systems*.
- 604 [28] I.T. Jolliffe. *Principal component analysis*. Springer, 2002.
- 605 [29] T. Kanungo, D.M. Mount, N.S. Netanyahu, C.D. Piatko, and A.Y. Wu. An efficient k-
606 means clustering algorithm: Analysis and implementation. *IEEE Trans Pattern Analysis*
607 *and Machine Intelligence*, vol.24:881–892, 2002.
- 608 [30] A. Krause, E. Horvitz, A. Kansal, and F. Zhao. Toward community sensing. In *Proceedings*
609 *of ACM/IEEE International Conference on Information Processing in Sensor Networks*
610 *(IPSN)*, St. Louis, MO, April 2008.
- 611 [31] D.D. Lee and H.S. Seung. Algorithms for non-negative matrix factorization. In *Proceedings*
612 *of NIPS 2000*, pp.556-562, 2000.
- 613 [32] M. Lighthill and G. Whitham. On kinematic waves. II. a theory of traffic flow on long
614 crowded roads. *Proceedings of the Royal Society of London. Series A, Mathematical and*
615 *Physical Sciences*, 229(1178):317–345, May 1955.
- 616 [33] C.J. Lin. Projected gradient methods for non-negative matrix factorization. *Neural Com-*
617 *putation*, vol.19,no.10:2756–2779, 2007.
- 618 [34] J. MacQueen. Some methods for classification and analysis of multivariate observations.
619 In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*,
620 volume 1, page 14. California, USA, 1967.
- 621 [35] W. Min, L. Wynter, and Y. Amemiya. Road traffic prediction with spatio-temporal corre-
622 lations. Technical report, IBM Watson Research Center, 2007.

- 623 [36] X. Min, J. Hu, Q. Chen, T. Zhang, and Y. Zhang. Short-term traffic flow forecasting of
624 urban network based on dynamic STARIMA model. In *Proceedings of the 12th International*
625 *IEEE Conference on Intelligent Transportation Systems (ITSC '09)*, 2009.
- 626 [37] C. Quek, M. Pasquier, and B. Lim. POP-TRAFFIC:A Novel Fuzzy Neural Approach to
627 Road Traffic Analysis and Prediction. *IEEE Transactions on Intelligent Transportation*
628 *Systems*, 7(2):133–146, 2006.
- 629 [38] C. Robert. *The Bayesian choice: a decision-theoretic motivation*. Springer-Verlag, 1994.
- 630 [39] A. Statthopoulos and M. G. Karlaftis. A multivariate state space approach for urban traffic
631 flow modeling and predicting. *Journal of Transportation Research Part C*, 11:121–135, 2003.
- 632 [40] C. Stutz and T.A. Runkler. Classification and Prediction of Road Traffic Using Application-
633 Specific Fuzzy Clustering. *IEEE Transactions on Fuzzy Systems*, 10(3):297–308, 2002.
- 634 [41] X. Sun, L. Munoz, and R. Horowitz. Mixture Kalman filter based highway congestion mode
635 and vehicle density estimator and its application. In *Proceedings of the 2004 American*
636 *Control Conference*, pages 2098–2103, Boston, MA, 2004.
- 637 [42] G.J. Szekely and M.L. Rizzo. Hierarchical clustering via joint between-within dis-
638 tances:extending ward’s minimum variance method. *Journal of Classification*, vol.22:151–
639 183, 2005.
- 640 [43] A. Thiagarajan, L. Sivalingam, K. LaCurts, S. Toledo, J. Eriksson, S. Madden, and
641 H. Balakrishnan. VTrack: Accurate, Energy-Aware Traffic Delay Estimation Using Mo-
642 bile Phones. In *7th ACM Conference on Embedded Networked Sensor Systems (SenSys)*,
643 Berkeley, CA, November 2009.
- 644 [44] Y. Wang and M. Papageorgiou. Real-time freeway traffic state estimation based on extended
645 kalman filter: a general approach. *Transportation Research Part B*, 39:141–167, 2005.
- 646 [45] D. Work, S. Blandin, O. Tossavainen, B. Piccoli, and A. Bayen. A traffic model for velocity
647 data assimilation. *Applied Research Mathematics eXpress (ARMX)*, April 2010.
- 648 [46] H. Yin, S.C. Wong, J. Xu, and C.K. Wong. Urban traffic flow prediction using a fuzzy-
649 neural approach. *Transportation Research Part C: Emerging Technologies*, 10(2):85–98,
650 2002.