

Spatio-temporal Coupled Bayesian Robust Principal Component Analysis for Road Traffic Event Detection

Shiming Yang¹ and Konstantinos Kalpakis¹ and Alain Biem²

Abstract—Road traffic sensors provide us with rich multi-variable datastreams about the current traffic conditions. Occasionally, there are unusual traffic events (such as accidents, jams, severe weather, etc) that disrupt the expected road traffic conditions. Detecting the occurrence of such events in an online and real-time manner is useful to drivers in planning their routes and in the management of the transportation infrastructure.

We propose a new method for detecting traffic events that impact road traffic conditions by extending the Bayesian Robust Principal Component Analysis (RPCA) approach. Our method couples multiple traffic datastreams so that they share a certain sparse structure. This sparse structure is used to localize traffic events in space and time. The traffic datastreams are measurements of different physical quantities (e.g. traffic flow, road occupancy) by different nearby sensors. Our proposed method process datastreams in an incremental way with little computational cost, and hence it is suitable to detect events in an online and real-time manner.

We experimentally analyze the detection performance of the proposed coupled Bayesian RPCA using real data from loop detectors on the Minnesota I-494. We find that our method significantly improves the detection accuracy when compared with the traditional PCA and non-coupled Bayesian RPCA.


I. INTRODUCTION

A large number of traffic sensors are continuously deployed to collect data for traffic conditions. Federal and State transportation agencies carry out various programs to collect traffic data by means of inductive loop detectors, video surveillance systems, and microwave radar sensors [13]. Collected traffic data include traffic volume, velocity, density, and vehicle classification. These data serve different purposes of study, such as alerting drivers about congestion and accidents, planning new road pavements to accommodate predicted traffic loads, and so on.

There are two major ways to utilize the automatically and continuously collected sensor data. First, we can detect traffic events in their early stage, and send early warnings to drivers for decision making. Second, we can use data assimilation methods to fuse real observations into forecasting models and produce more accurate near-future traffic condition prediction, which also helps drivers avoid traffic jams and plan better travel routes. In this study, we focus on the first way of using traffic observations. On busy highways, accidents

may cause severe and quick accumulation of vehicles. By the time police receives report and issues a warning to other drivers, it may have past several minutes or even longer. In this situation, many drivers may miss the chance to take an exit before entering the traffic jam trap. However, if we utilize continuous, automated road sensor observations, we can promptly detect the incidents and traffic jams in their early stage of formation. Knowing the development of incidents can help choose appropriate prediction models for issuing early warnings, and drivers will have a chance to adapt their routes.

Traffic observations demonstrate strong spatial and temporal patterns, showing periodicity and strong correlation between adjacent upstream and downstream observations. These patterns may vary depending on time in a day, day of the week, seasons, or locations. Figure 1 shows annual average of traffic flow on each weekday with 15-minute time resolution. The traffic flow data were collected from 42 sensors along the southbound I-494 in Minnesota by the Minnesota Department of Transportation. The figure shows different flow patterns in each weekday, e.g. Saturdays and Sundays have smaller flow amount than weekdays. At different hours within a day, traffic flow also shows different patterns, while some of these patterns persist at the same time but different days. Besides, we can observe that neighboring sensors present similar patterns. Occasional incidents or events of long duration show as abnormal events on the background of normal traffic behaviors.



(a) Weekday patterns

Fig. 1. Annual average of traffic flows on each weekday with 15-minute time resolution for 38 loop detector sensors on I-494 southbound and eastbound.

To find “abnormal” events, one simple yet feasible way is to define a normal range of traffic measurements based on experience, and use thresholds to identify such special events. However, threshold based methods are not reliable

¹S. Yang is with Department of Computer Science and Electrical Engineering, University of Maryland, Baltimore County, 1000 Hilltop Circle, USA shiming1@umbc.edu

¹K. Kalpakis is with Faculty of Computer Science and Electrical Engineering, University of Maryland, Baltimore County, 1000 Hilltop Circle, USA kalpakis@csee.umbc.edu

²A. Biem is with IBM T.J. Watson Research Center, NY, USA. biem@us.ibm.com

nor adaptive to changing environment. One approach, namely robust PCA, which is widely used in detecting moving objects in video frames shows promising connections to the traffic event detection. In robust PCA, the normal background data dwell in a low dimensional space, due to the strong correlation among normal observations, and special events dwell in a sparse subspace, which allows an observation set to be decomposed into two subspaces, a normal background subspace specified by a low rank matrix, and an anomaly subspace specified by a sparse matrix [4], [5].

A traffic time series can be viewed as images, if we collect the daily observations as one column of pixels in that ‘image’.matrix. Special events can be viewed as anomalies in the image. Thus, a 1-dimensional datastream of traffic observations can be converted into a matrix format.

In this paper, we convert 1-dimensional time series into matrix format, and then decompose that matrix into superposition of a low rank, a sparse, and a noise matrix. We use a non-parametric Bayesian method, the Bayesian RPCA to learn the latent low rank and sparse matrices from the observation. Furthermore, we extend the Bayesian RPCA to multiple multiple variables/timeseries/datastreams, which may correspond to different observations that are aligned in time and space or aligned in time but collected by nearby sensors, by coupling them to share sparsity pattern/structure. Our hypothesis is that by sharing a sparsity structure among multiple datastreams affected by the same events, we may be able to improve the detection accuracy of such events from those observation datastreams. We call this approach the coupled Bayesian RPCA. Using real traffic data, we experimentally show that our proposed approach achieves higher accuracy in detecting different types of traffic events.

The remainder of this paper is organized as follows. In Section II, we briefly review related work on traffic event detection, the robust PCA and its typical applications. In Section III, we describe the characteristics of traffic data streams, and show the spatial correlations between sensors. We also demonstrate different types of events of interest. In Section IV, we describe our proposed coupled Bayesian Robust PCA method. The results of the experimental evaluation of our method are given in Section V. Conclusions are in Section VI.

II. RELATED WORK

Detecting interesting events is a pervasive problem in many applications, such as video security surveillance, text mining, etc. In road traffic event detection, a large volume of work is based on surveillance video. In some methods, pixel-level features are extracted to represent interesting spatial and/or temporal events. Jiang et al. [10] proposed a dynamic hierarchical clustering method to detect abnormalities in traffic video. Morris and Trivedi [11] compute traffic flow parameters from live video streams, and use speed profiles to categorize traffic motion states.

Ihler et al. [9] use Markov Modulated Poisson Processes (MMPP) to discover events from timeseries of count data. In a probabilistic model framework, they separate

the observed timeseries as a superposition of regular and aperiodic processes. They use a non-homogeneous Poisson process to model regular count data. For rare events, they use a Markov chain to model the transition between states including increased, decreased, or unchanged activities. The MMPP method was applied to both highway traffic and building pedestrian counts, and achieved significantly higher accuracy compared to a baseline threshold model.

Principal components analysis (PCA) is well known technique used for dimensionality reduction. PCA also has successful applications in finding outliers. It can be viewed in probabilistic Bayesian way by using a latent variable [3], [14]. With Bayesian PCA, the latent variable can be inferred [3]. Conventional PCA’s major drawback is that it is not robust to outliers. Recently, the robust PCA [4], [6], [7] (a special case of matrix completion) has attracted much attention in decomposing matrices into low rank and sparse components matrices. This method finds wide applications in many engineering and statistical modeling problems, where order, dimensionality, or complexity of a model can be evaluated by the rank of an appropriate matrix. The low rank component can be viewed as a “denoised” version of the data. The sparse component is useful in detecting the outliers. For example, in detecting moving objects in a sequence of movie frames with a varying background, one can use the sparse components to locate the moving objects. Becker et al [2] demonstrated an image processing system that recover images from noisy observations using matrix completion. Candès et al. [5], [6] provided examples of matrix completion in recovering signals from the magnitude of their Fourier transform, as well a collaborative filtering in online recommendation systems.

Ding et al. [8] proposed a hierarchical Bayesian model to decompose sequential video image matrix into a low rank and a sparse component. They considered that an object will have larger probability to show in the next frame, if it shows up in the previous frame. Hence, they introduce Markov dependence into the Bayesian RPCA. This is a good example of utilizing known structure of sparsity to guide the learning scheme. In our study, we experimentally find that the Bayesian RPCA has good performance for detecting traffic events. Then, we explore how to couple multiple traffic time series by taking advantage of any shared sparsity structure.

This work provides a partial answer to Ihler’s question [9] “*how those covarying data streams should be combined, and to what degree their parameters can be shared*”.

III. TRAFFIC DATA AND EVENTS

A. Traffic data

A variety of sensors are used to measure traffic conditions, such as inductive loop detectors, video surveillance cameras, microwave radars and probe vehicles. Nowadays, using GPS-enabled smartphones to collect traffic data is attracting an increasing interest and attention. Typically measured quantities include traffic density, vehicle speed, traffic flow, etc. The traffic flow measures the number of vehicles that pass a single point in a given time interval. Road occupancy is

the percentage of time that a detector is active due to the presence of vehicles during a time period. It can be used as a proxy for road density [13]. Given a distribution of vehicle lengths, the average vehicle speed can be estimated from traffic flow rate and road occupancy.

Traffic data have some important characteristics. First, daily traffic data have strong temporal correlation, especially when grouped by weekday.

Second, they contain sparse interesting features. Compared to daily repeated traffic patterns, traffic events are rare and hence manifested as random sparse features in the data that persist for some time. As result, we get structured sparsity with random occurrences in the traffic data streams.

Third, traffic data also present strong spatial correlation. Downstream traffic measurements are influenced by nearby upstream traffic measurements, and vice versa. For example, Figure 2 illustrates strong linear correlation (≥ 0.76) between upstream and downstream sensor measurements (except for sensors at the flyovers) for the Minnesota real dataset described in Section V.

Fourth, using different types of traffic measurements allows to better capture the characteristics and effects of traffic events. Hence, we seek to take advantage of the shared sparsity structure between temporarily and spatially correlated variables/measurements of different types.




Fig. 2. Linear correlation between upstream and downstream sensor flow rate measurements in a section of I-494.

B. Traffic events

Information about current traffic events is critical to drivers when deciding their traveling routes and schedules, as well as to the management of the transportation infrastructure. Government transportation agencies publish timely traffic incident information, such as the type (crash, roadwork, stall, and hazard), location, start time, and duration of incidents.

Typically, an incident may cause the downstream traffic flow and road density to reduce sharply, while the upstream traffic may temporarily reduce flow and increase road density. However, sensor measurements indicate exceptions are possible. For example, late night roadwork that is announced way ahead of time may have limited impact on traffic measurements, and hence it will not be detectable from the sensor measurements alone.

Traffic events as announced have low spatial and temporal resolution: we often only get the closest intersection where an event happened, while there is latency between the occurrence and reported time of events. Therefore, it is desirable to

increase the accuracy of the location and happening time of events from traffic data streams. Moreover, it is also desirable to detect events before they are reported to the relevant transportation agencies, so that it reduces the latencies and improves asset utilization when managing incidents.

Furthermore, due to limited transportation management assets, many events go unreported, despite the fact that have noticeable impact on traffic and would be useful to drivers. For example, traffic jams caused by non-accidents can produce low flow rate and high road density; social activities, such as football games and music concerts, can produce unexpected high traffic volumes; reduced visibility due to bad weather (heavy snow, fog, or rain) leads into reduced road capacity, increased road density, and reduced vehicle flow rates. In the experiments presented below, we will study how the algorithm performs on finding events like those as well as reported incidents.

Figure 3 illustrates four events of different types, which are indicated by green shadow areas. Event 1 is a roadwork event lasting 47 minutes and blocking a road lane ¹. We observe reduced flow rate and road occupancy during that time, which normally has higher vehicle flux.

Event 2 corresponds to reduced traffic speed due to limited visibility ². Traffic remained free flow, but had decreased flow rate and increased road occupancy. This suggests that vehicles proceeded with caution at lower speeds while the road density is still low.

Event 3 demonstrates an unusual traffic increase at a Saturday night. Such a peak is generally not observed during weekdays but only on some weekends. We hypothesize that it was caused by social activities that usually happen in weekends.

Event 4 presents a traffic jam in early morning on a weekday. There was no incident reported from the transportation agency. However, there was heavy snow and low visibility during that time period, and taking into account that it was during morning rush hour on a Wednesday, we infer the severe weather caused that traffic jam.

IV. COUPLED BAYESIAN RPCA

In this section, we introduce the Bayesian RPCA, and describe our extension of this method to multiple coupled time series. In our extension, the coupling of multiple timeseries generated by dynamical systems is achieved through sharing the same sparsity structure in the Bayesian RPCA. Bayesian RPCA was formally presented in [8] and used in a shopping mall human video surveillance system.

Given a measurement matrix Y , we would like to separate Y into a superposition of three components $Y = L + S + E$, where L is a low rank matrix for normal background, S is a sparse matrix for rare events, and E is a noise matrix with small magnitude elements for noise. The Bayesian RPCA has strong connections to factor analysis, and has deep roots in linear Gaussian systems [3], [12].

¹Data source: MNDOT, 511MN.org

²Date source:weathersource.com

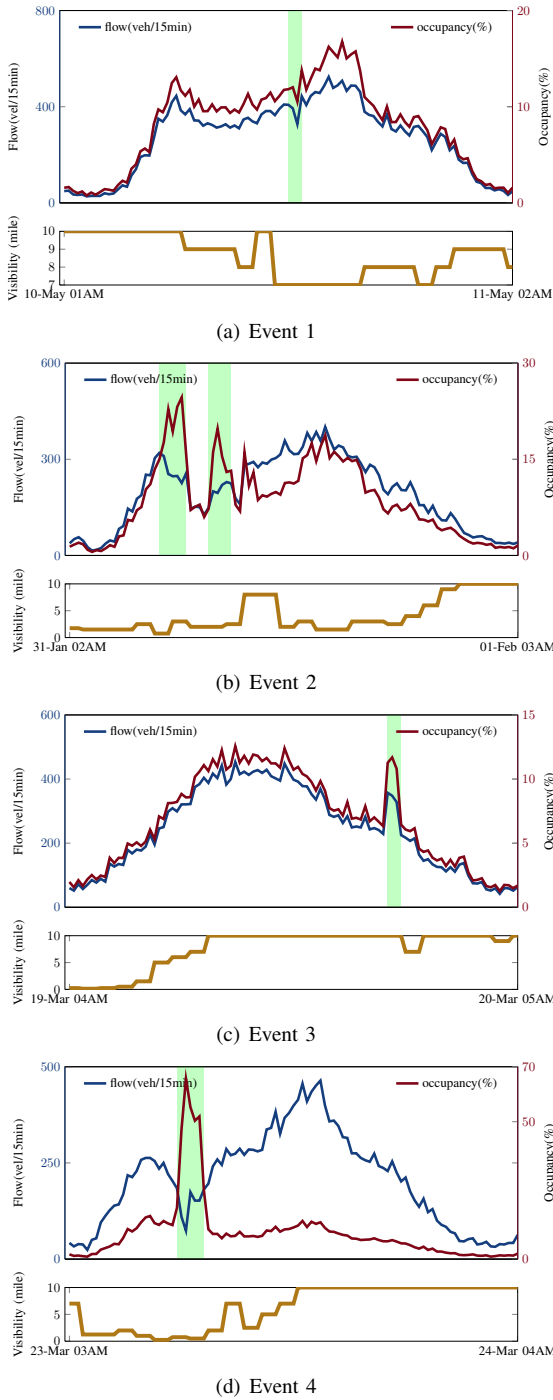


Fig. 3. Examples of different event types and their impact on sensor readings.

The low rank component is expressed through SVD as $L = D\Lambda W^T$, where $\Lambda \in \mathbb{R}^{r \times r}$ is a diagonal matrix with main diagonal entries being the singular values of L , $D \in \mathbb{R}^{n \times r}$ and $W \in \mathbb{R}^{d \times r}$ are constructed from the left and right singular vectors of L . Since the rank r is unknown, non-parametric Bayesian approach treats it as a latent variable, and learns it from the data. In [8], the singular values matrix is further decomposed into $Z\Lambda$, such

that $L = D(Z\Lambda)W^T$, where $D, W \in \mathbb{R}^{n \times k}$ are drawn from multivariate Gaussian distributions, and $Z, \Lambda \in \mathbb{R}^{k \times k}$ are drawn from multivariate Beta-Bernoulli and Gaussian-Gamma random processes. This method is similar to dictionary learning, where k is the dictionary size and D is the dictionary [8]. A Bernoulli distribution with hyperparameters having a Beta prior distribution is used to learn the binary values of the diagonal elements of Z , while a multivariate normal distribution with hyperparameters having a Gamma prior distribution is used to learn the real values of the diagonal elements in Λ .

This Bayesian non-parametric approach provides a flexible way to estimate the rank of the low-rank component matrix without a preset parameter.

The sparse component S is decomposed as $S = B \circ X$, where \circ denotes the element-wise multiplication of two matrices, $s_{ij} = b_{ij}x_{ij}$, where B is a binary matrix and X is a real matrix. The binary elements of B are learned from a Bernoulli process with hyperparameters having a Beta prior distribution. The columns of X are learned from a multivariate Gaussian distribution with hyperparameters having a Gamma prior distribution. In [8], Markov spatio-temporal dependencies in the sparse component are considered, since a moving object in one video frame is highly likely to show up in the next frame at a nearby location. Therefore, a Markov dependency along the columns (corresponding to video frames) is enforced in the sparse component.

In our traffic problem, an event at one time in a given day is unlikely to happen again the next day at the same time and location. However, an event will persist for some time after it happens. So when a traffic event happens, it is likely that it will persist to the next time point.

When measuring many variables of a dynamical system, it is natural to consider use multiple of them to learn the system, instead of using only one variable. Besides, those variables may share some internal structure, given that they are generated from a single system/process. We extend the Bayesian RPCA method by coupling the sparsity structure of multiple variables that exhibit spatio-temporal dependencies. For example, given two data matrices Y_1 and Y_2 observing different physical quantities/variables at the same time and location, we have

$$\begin{aligned} Y_1 &= D_1(Z_1\Lambda_1)W_1^T + B \circ X_1 + E_1 \\ Y_2 &= D_2(Z_2\Lambda_2)W_2^T + B \circ X_2 + E_2, \end{aligned}$$

where the binary matrix B , capturing sparsity, is shared by both data matrices. Such sparsity structure sharing can be further extended to measurements from multiple sensors in a neighborhood, which may be affected by the same event in time. Figure 4 shows the graphical model of our extension to Bayesian RPCA, with coupled variables sharing the same sparsity structure.

In traffic applications, loop detector sensors measure/collect multiple variables. The flow rate is a measurement of number of vehicles passing through the sensor during a unit time (e.g. 30 seconds). Usually, small flow rate means

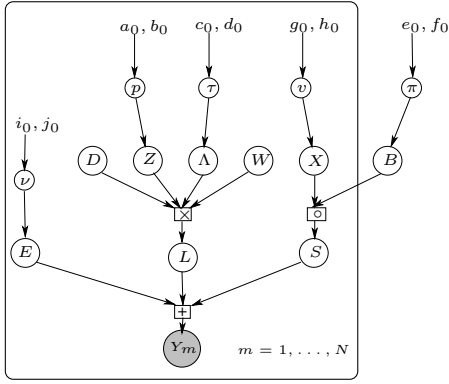


Fig. 4. Graphical model of Bayesian RPCA method with shared sparsity structure between multiple measured variables. There is one “plate” for each measured variable.

free flow on the road, while larger flow rate means higher density on the road. However, when there is a traffic jam and all vehicles move slowly or even stop, the flow rate is low, but the traffic is not free flowing. An extra variable, the road occupancy, is helpful in distinguishing between these two event types. Also, adjacent sensors may record the same events; hence they are very likely to share the same structure of sparsity of B .

V. EVENTS DETECTION EXPERIMENTS

A. Experimental data.

We obtained a dataset of vehicle flow and road occupancy measured by 38 inductive loop detectors on highway I-494 from the Minnesota Traffic Observatory (MTO) for the whole year of 2011. The 38 loop detector sensors measure traffic flow and occupancy on the southern and eastern bound I-494. The dataset has time resolution of 30 seconds, which accounts for 2880 data points per day for each measurement. Inductive loop detectors fail sometimes, and the missing measurements are flagged as -1 in the dataset ($\approx 8.9\%$ values are missing).

In addition to the readings from the loop detectors on the Minnesota I-494 highway, we also gather incident reports, and weather information. We gather weather data from weathersource.com. We gather incident reports for the segment of road under study for 2011 from the RSS feeds published by the Minnesota Department of Transportation at 511MN.org. Each incident is described by start time, duration, location, camera ID, event type, and blocking. For example, on March 23, 2011 at 11:41AM in cross-section of I-494 EB and Flying Cloud Drive, there was a crash, which lasted for 45 minutes, and blocked at least one lane.

We preprocess the sensor readings by aggregating them into 15-minute resolution, which is the typical time interval used in vehicle volume counts [1]. Considering the strong daily patterns presented by the sensor data, we construct a matrix for each sensor/variable, whose columns contain the daily aggregated data. Thus, an $n \times d$ matrix collects one sensor’s data for d days with n points on each day. We also

remove an entire day, if the sensor was flagged as failed for the whole day. However, for short duration failures, failure, where some sensor readings during the day are flagged as -1, we keep that day’s column.

To evaluate our method of detecting events, we need to identity/annotate the sensor measurements with information from ground-truth traffic events. To this end, we use the gathered incident reports with some changes, as described below. The gathered collection of incidents has some limitations. First, the collection contains only incident events, such as roadwork, crash, hazard, which often reduce the traffic flow. As discussed earlier, other events (e.g. due to social activities or severe weather) may also affect traffic flow yet they are not included in that collection. Second, the collection may contain incidents that do not impact the sensor measurements and thus can not be detected from those sensor measurements.

To mitigate these limitations, we proceed as follows. We plot the whole year’s worth of flow, occupancy, their annual average for each 15 minutes in a day, reported incidents, and visibility, aligned by time. Then, we have asked three people to individually review those plots and annotate time segments that they consider fall into the following three types of traffic events: (1) events that reduce flow rate or speed, (2) free flow with unusual high traffic volume, and (3) sensor failure. Annotations where the humans originally disagree, the reviewers were asked to reach a consensus annotation. During this annotation process, the reviewers had no access to the output (detected events) of our proposed approach.

In what follows, we consider the measurements from loop detectors with IDs 196 and 197 on the eastern-bound I-494, which are within 0.5 miles of each other. The ground-truth consists of 232 events total for the two sensors: 148 type-1 events, 83 type-2 events, and 1 type-3 event.

B. Coupling multiple variables

We compare three methods for detecting events from measurements from a single sensor: our proposed method using sparsity coupling for the traffic flow and road occupancy, the traditional PCA on traffic flow as a baseline, and the Bayesian RPCA with traffic flow. We use the same hyperparameters and priors for the proposed method and the original Bayesian RPCA.

For the PCA method, we use a simple threshold on the noise (residual not captured by the principal components). We used 25 principal components (out of 96 possible components) capturing over 95% of the energy of the data (Euclidean norm). The noise threshold was set to 30 vehicles/15 mins. We decided on this threshold value so that PCA finds approximately the same number of false-positive events as the other methods. With this setting, the number of false-positives for PCA (flow/occupancy), Bayesian RPCA (flow/occupancy), and our proposed coupled Bayesian RPCA were 214/247, 281/290, and 257 respectively.

For each method, we count the number of detected events, detected but unlabeled events, undetected but labeled events. For this Minnesota dataset,

Method	Variable	Overall	Type 1	Type 2
PCA	flow	46.1%	49.3%	39.8%
PCA	occupancy	54.3%	59.5%	45.8%
BRPCA	flow	80.2%	73.6%	91.6%
BRPCA	occupancy	81.5%	86.5%	73.5%
Coupled BRPCA	flow& occupancy	83.2%	87.2%	77.1%

TABLE I
DETECTION ACCURACY COMPARISON OF THREE METHODS.

Table V-B shows the accuracy (ratio of the number of detected events over ground-truth events) of the three methods. Our proposed coupled Bayesian RPCA has the best accuracy (83.2%) among the three, especially in finding events that have negative impacts on traffic conditions (events of type 1). For type 2 events, Bayesian RPCA using only traffic flow has better accuracy than coupling flow and occupancy in our method. Note however that type 2 events (free flow with unusual high traffic volume) have significantly less negative impact than type 1 events,

Figure 5(a) shows an incident (green dots in the *Events* axes) reported by 511MN.org, and is also annotated by the human reviewers. The incident report indicates a car crash near sensor 196 with lane blocking lasting 20 minutes. We observe reduced vehicle flow and road occupancy. The event is correctly detected by our method.

Figure 5(b) illustrate a case of reduced traffic speed due to low visibility. During the time period indicated by the red bars, there was heavy snow in the I-494 area with visibility lower than 2 miles. The flow rate is significantly lower than average, while the road occupancy is not reduced accordingly, suggesting that drivers proceeded with caution at low speed during that time. No incidents were reported by 511.org near that time/location, but the event is annotated by the human reviewers. The event is correctly detected by our method.

Figure 5(c) suggests a lane blocking event. There is no reported incident at the time/location, but the event is annotated by the human reviewers. The event is correctly detected by our method.

Figure 5(d) illustrates a case where there is a reported incident, which does not impact the traffic flow or occupancy. No event in the nearby time/location is annotated by the human reviewers, nor detected by our method.

C. Coupling measurements from neighboring sensors

In light of the strong correlation between measurements from neighboring sensors we discussed in Section III, we couple single variable measurements (either traffic flow or road occupancy) from adjacent sensors. Table V-C shows that each coupling road occupancy from neighboring sensors gives the best accuracy overall and for type 1 events. Coupling traffic flow gives the best accuracy for type 2 events.

The (coupled) Bayesian RPCA can be implemented for online processing purpose. Its running time also promises real-time processing. Above experiments were running on

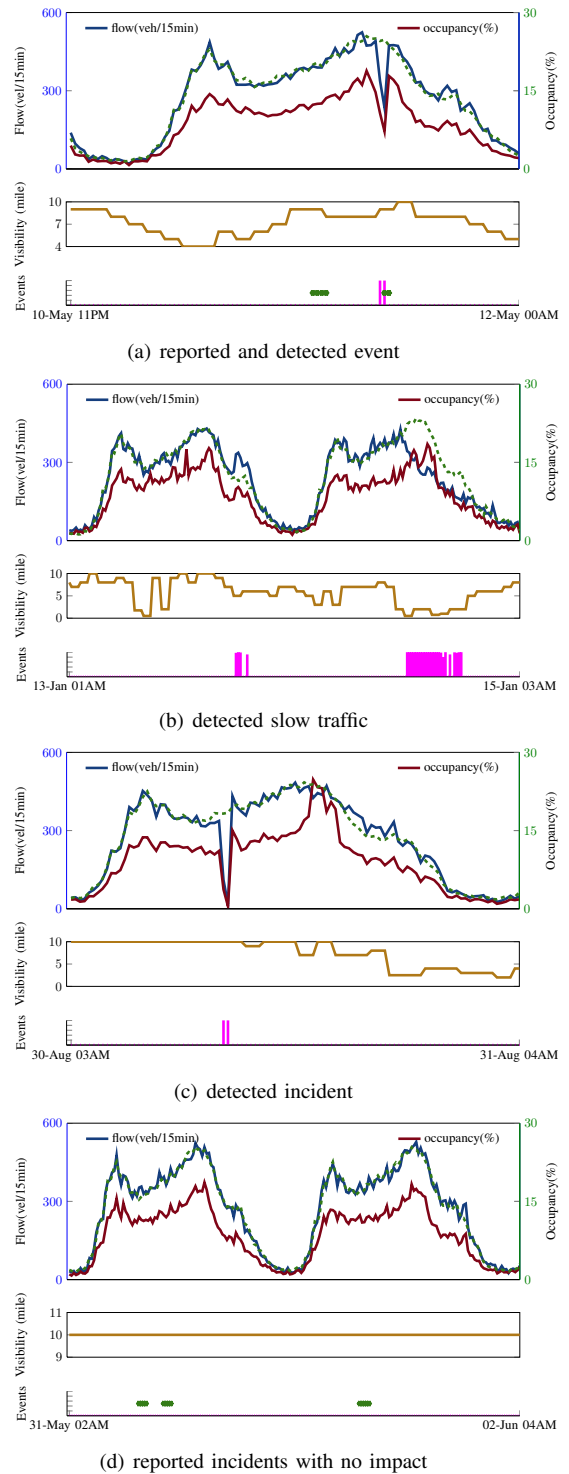


Fig. 5. Comparison of detected events and ground-truth events.

Spatial Coupling	Overall	Type 1	Type 2
flow	82.5%	74.7%	100%
occupancy	93.7%	94.9%	90.9%

TABLE II
DETECTION ACCURACY COMPARISON OF THREE METHODS.

a 64-bit machine with 8GB memory, Matlab 2011b environment. Average running time using one variable for an entire year is about 2 minutes, with 1000 burn-in and 1000 iterations in Gibbs sampling. Average running time for two coupled variables takes 4 minutes. However, in real use, the matrix size can be dramatically reduced by using only recent a few months data, or only use the same weekday's data for a specific weekday.

VI. CONCLUSION

We consider the problem of detecting traffic events from continuously collected measurements from road traffic condition sensors. We focus on traffic slow down, unexpected high traffic volume, and traffic jams events. Detecting such traffic events can help drivers make early decisions to choose travel routes, saving time and energy.

We propose an extension to Bayesian Robust PCA for detecting events in road traffic datastreams. Our extension couples multiple streams of sensor measurements so that they share the sparsity pattern which is inherent in the robust PCA approach. The sparsity sharing can be used both in the temporal dimension (by using different variables measured at the same location) and in the spatial dimension (by using the same variable measured at nearby locations). We experimentally demonstrate using a real dataset that our proposed method significantly improves upon the event detection accuracy of the traditional PCA and Bayesian Robust PCA methods. Furthermore, the Bayesian nature of the method facilitates a probabilistic interpretation of the detected events, while removing the tension in deciding the number of principal components to be used. Moreover, the proposed method can process traffic datastreams incrementally and thus can detect events in an online and real-time fashion.

Though we focus on event detection, in the future we plan to use the low-rank component of the traffic datastream computed by the proposed method, to derive initial and boundary conditions for macroscopic road traffic models. For example, the low-rank component can be used to estimate "normal" incoming/outgoing traffic flow at the entry/exit ramps (by removing the effects of outliers/incidents), which will allow to improve the accuracy of short-term forecasting of road traffic conditions.

ACKNOWLEDGMENT

We would like to thank IBM for partially supporting this research effort. We would also like to thank the Minnesota Traffic Observatory for providing us with of raw loop detector dataset, and thank Ms. Yulei Wang and Ms. Yao Li for their help on annotating the dataset with traffic events.

REFERENCES

- [1] Guide for interpreting short duration traffic count reports. Washington State Department of Transportation, www.wsdot.wa.gov, 2010.
- [2] S. Becker, E. J. Candès, and M. Grant. Templates for convex cone problems with applications to sparse signal recovery. *Mathematical Programming Computation*, 3(3):165–218, 2011.
- [3] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

- [4] E. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(2), 2011.
- [5] E. J. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- [6] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.
- [7] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- [8] X. Ding, L. He, and L. Carin. Bayesian robust principal component analysis. *Image Processing, IEEE Transactions on*, 20:3419–3430, 2011.
- [9] A. Ihler, J. Hutchins, and P. Smyth. Learning to detect events with markov-modulated poisson processes. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1, 2007.
- [10] F. Jiang, Y. Wu, and A. Katsaggelos. Abnormal event detection from surveillance video by dynamic hierarchical clustering. *IEEE International Conference on Image Processing*, 5:V–145–V–148, 2007.
- [11] B. Morris and M. Trivedi. Real-time video based highway traffic measurement and performance monitoring. *Proceedings of the IEEE Intelligent Transportation System Conference*, pages 59–64, 2007.
- [12] K. P. Murphy. *Machine Learning A Probabilistic Perspective*. MIT Press, 2012.
- [13] S. Olariu and M. Weigle, editors. *Vehicular Networks From Theory to Practice*. CRC Press, 2009.
- [14] M. E. Tipping and C. M. Bishop. Mixtures of probabilistic principal component analysis. *Neural Computation*, 11:443–482, 1999.