Integrated Infrastructure and Vehicle Based Monitoring for Enhanced Efficiency and Safety

Brendan Morris and Mohan Trivedi
Computer Vision and Robotics Research Laboratory
University of California, San Diego
La Jolla, California 92093-0434
Email: {b1morris, mtrivedi}@ucsd.edu

Abstract

In a world that is increasingly reliant on automobiles, efficient management of the transportation network for increased satisfaction and safety is a key functionality of intelligent transportation systems (ITS). A wide range of these ITS applications rely on the extraction of current data as well as historical information in order to perform appropriately. This paper presents a general framework for integration of infrastructure and vehicle based sensing. Maps link the coarse but wide area coverage offered by infrastructure with detailed local measurements made at the vehicle. The integration of these two sources provides better understanding of all aspects of driving situations. A motion trajectory learning paradigm is used to extract characteristic driving behaviors which is used to describe traffic conditions and improve safety. Extensive experimental studies have proven the potential of the proposed framework.

1 Introduction

The last 100 years has brought great advancements and developments in personal transportation transforming the horse drawn world into one dominated by automobiles. The emergence of the automobile has opened up the world, providing almost unlimited access and mobility. In the United States alone, a staggering 240 million vehicles travel over 12 million miles annually on a network consisting of 4 million miles of road whose maintenance costs $40 million [1]. Such a large infrastructure has immediate social, economic, energy, and environmental impact. Motor vehicle taxes generate $30 billion annually. Americans use 175 billion gallons of fuel for highway travel releasing a number of emissions into the air. Between 1995 and 2001 there was 10% increase in average commute time as people experienced slower speeds and increased delay while stuck in congestion [2]. Perhaps most alarming were the 2.5 million injury accidents and 41 thousand motor vehicle related fatalaties in 2008 [3]. These numbers represent just a small portion of worldwide dependence on automobiles and emerging countries such as India and China will dramatically effect the future of transportation.

In order to manage such a vast transportation network, it is essential to invest in intelligent transportation systems (ITS) technologies. ITS solutions provide the means to extract and manage information necessary for continued advancement. Without their use it would be impossible to continually monitor our roadways, as no human could process such large amounts of data. The key requirements for successful ITS systems are the ability to extract and process data in real-time, provide robustness to a wide range of operating conditions, and technologies that can adapt to changes in the environment which is essential for long term deployment. The ITS community has the power to improve the quality of modern life by providing greener transportation solutions, greater satisfaction through smoother commutes, and ultimately a safer driving experience.

In this paper we present a general framework for the integration of infrastructure and vehicle sensing for intelligent transportation monitoring and active driver safety, shown in Fig. 1. The key goal of this scheme is an improved understanding of all aspects of driving situations. This requires a deep understanding driving behavior. This behavior can be learned by examining motion trajectories in the surround of driver (vehicle sensing) and providing observation context through scene knowledge (infrastructure sensing).

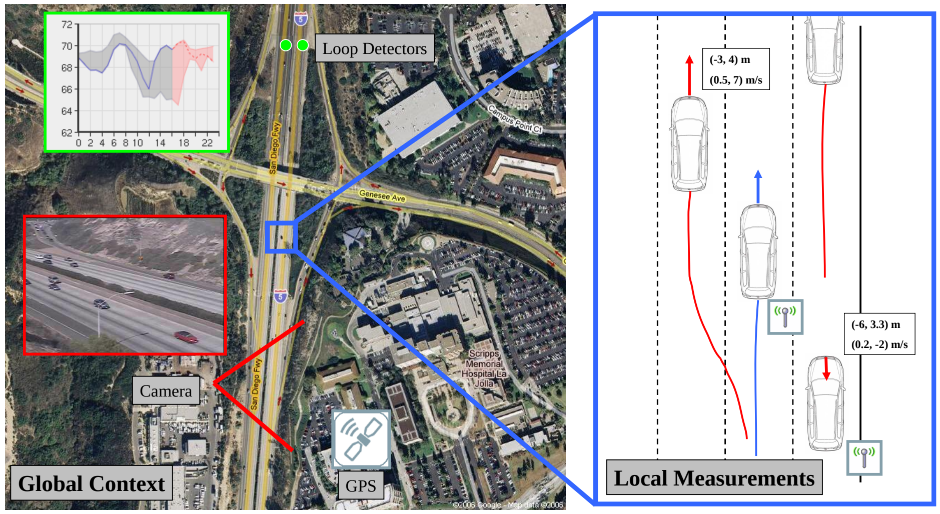

Figure 1: Active safety with trajectory analysis. Local measurements made while driving at the vehicle level and put into global context by fusion with infrastructure sensors by GPS registration.

2 Trajectory Analysis

The most critical safety concerns during driving are objects in motion (with relative motion) which can collide and cause damage. Motion also indicates when a lane change occurs or if someone is breaking to avoid a collision. The motion encountered while driving is not random, it has some underlying structure imbued by traffic laws. By leveraging this structure, driving behavior can be analyzed and studied [4]. The diagram in Fig. 2 shows a general framework for understanding behavior based on motion. Motion is represented by trajectories because it compactly encodes behavior and provides a historical summary both spatially and temporally. The temporal history is essential for understanding behavior because it is defined not by just current actions but sequences over time. By observing motion and collecting trajectories over time, a database is formed that contains examples of typical behaviors. These typical behaviors, while unknown, can be teased out of the trajectories by clustering. After clustering trajectories, the typical behaviors can be modeled probabilistically to give a spatio-temporal representation. Finally, in an online (live) analysis

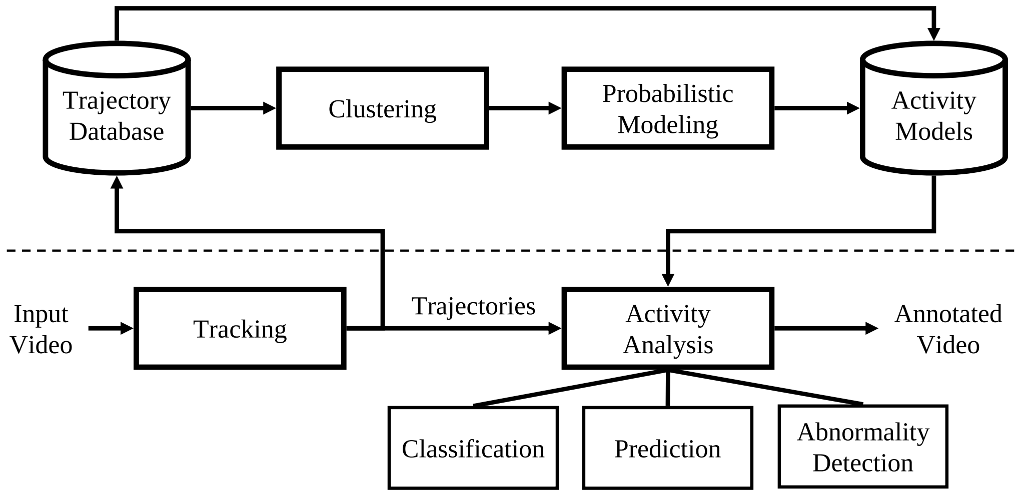

Figure 2: Trajectory learning framework for behavior analysis. Vehicles are detected and tracked, motion describes behavior, can be clustered to learn typical behaviors, which are modeled probabilistically for use in live analysis.

mode new trajectories from observed obstacles can be examined and categorized using the learned behavior models. The behavior of every detected object can be categorized as typical or abnormality and finally future behavior can be predicted.

2.1 Learning Typical Patterns

A key observation for trajectory analysis is that typical actions are repetitive while the unusual do not occur often therefore with sufficient observation it is possible to learn these prototypical behaviors. In order to learn typical patterns, a training database of trajectories is accumulated and condensed into a small set of behaviors through clustering. Unfortunately, the number of typical behaviors is not usually known a priori and must be estimated. The true number of activities is estimated by first clustering the training trajectories into a large number of groups. This rough grouping is refined into the tight compact clusters representing typical behaviors by agglomerative merging. Further details on the cluster procedure can be found in [4, 5]. In Fig. 3, the typical patterns are shown for a number of different

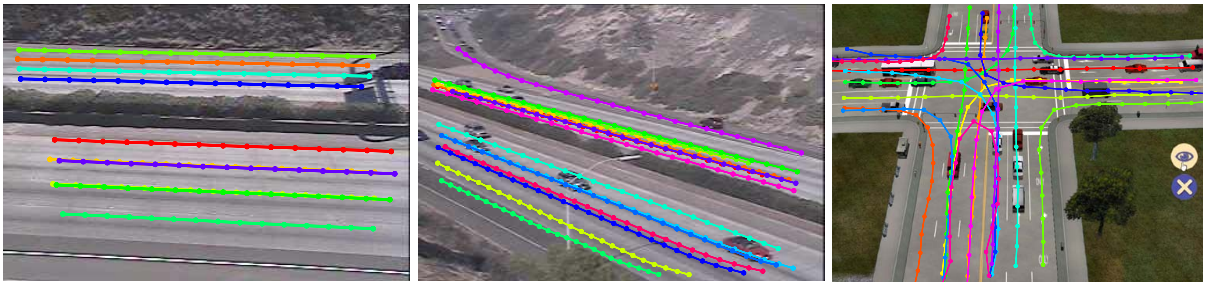

(a) (b) (c)

Figure 3: Typical trajectory patterns. (a) Profile view of highway traffic on I5 with the lanes clearly demarcated. (b) 3/4 view of I5 obtained by adjusting the PTZ camera parameters. (c) Example of typical intersection patterns corresponding to the acceptable intersection maneuvers.

(a) (b) (c) (d)

Figure 4: Abnormal tracks not adhering to prototypical paths

transportation scenes with prototype patterns corresponding to the lanes of the highway or to the allowable maneuvers at an intersection.

2.2 Behavior Analysis

A vocabulary to describe behavior is established by learning the typical scene patterns. During live video analysis, object activity is matched with an element of this learned set. The reference behavior set explains what are expected actions and can be used to predict future events. This type of prediction is more powerful than simple one step prediction associated with Kalman or particle filters because the prediction looks further into the future and is conditioned on actual observed behaviors. Fig. 5 shows the evolution of a left turn. The probability of a behavior is shown in the white box over the associated pattern. Early

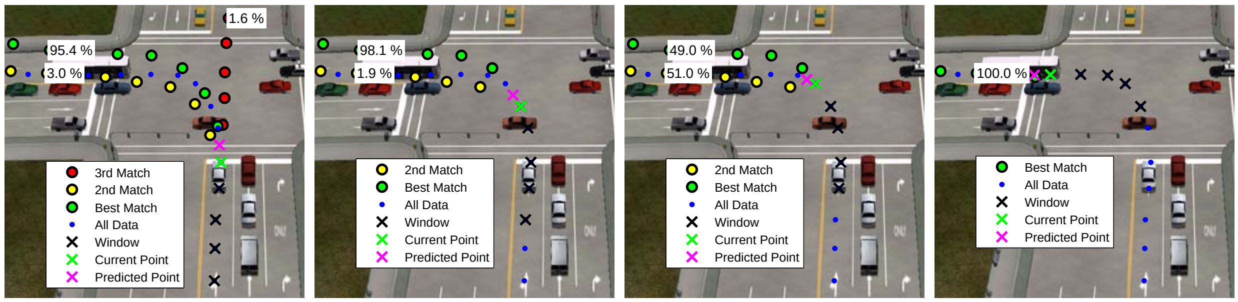

Figure 5: Left turn prediction behaves as expected. As more data points are collected the prediction better matches the true lane.

into the left turn maneuver it is unclear which lane the vehicle will turn into but as more information is obtained the estimate improves. This type of analysis is of great importance for intersection safety [6, 7] where the predicted paths are used to assess the probability and time to collision.

It might be more important to detect not the typical behaviors but the abnormal because these indicate when the out of ordinary occurs. A driver must be aware of these anomalous events because it is not possible to predict well what will happen in the future and may lead to dangerous situations. These abnormalities are detected when none of the reference set adequately explain the observed motion. A collection of abnormal trajectories are presented in Fig. 4.

3 Global Context

Global context can be maintained through mapping. Maps provide a way to position oneself in the larger transportation network with well defined landmarks. A portion of San Diego is displayed in Fig. 6 with red markers indicating the location of highway inductive loop sensors. This information is maintained in California through UC Berkeley by PeMS [8] and is used for live traffic updates and history statewide. By linking external infrastructure sensors to the map, a user could retrieve pertinent information such as current speed, traffic

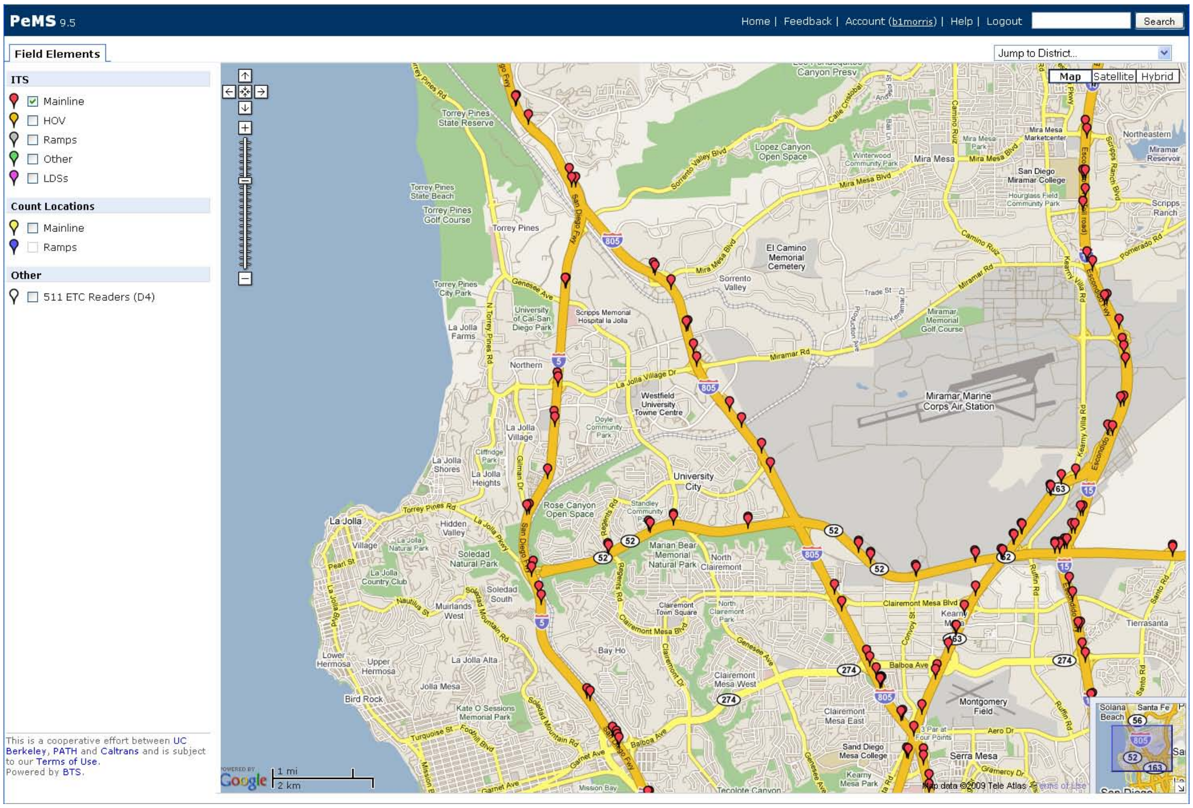

Figure 6: Maps provide context during driving. PeMS maintains historic traffic measurements inductive loop sensors spread all across California.

conditions, and speed, based on location.

Video cameras are an attractive infrastructure sensor because of an active research community, very high informational content, and widespread deployment. Traffic cameras as well as read light and speeding cameras have been widely deployed by a number of agencies, such as local government and news broadcasts, because they provide a quick view of traffic conditions with relative ease and minimal intrusion. These cameras provide valuable video signals that can be reused for alternative computer vision based analysis such as highway congestion performance measurements similar to loops [9], detection of stalled vehicles [10, 11], vehicle classification which is necessary for real-time fleet composition used for highway management functions such as estimating emissions or infrastructure load assessment [12], as well as intersection safety [6, 7].

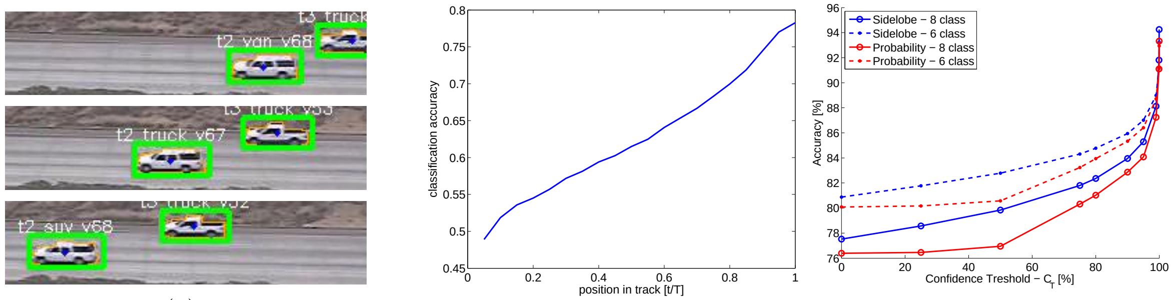

(a) (b) (c)

Figure 7: (a) Vehicle correctly classified as SUV after initially estimated as VAN and TRUCK. (b) Classification accuracy is improved by leveraging evidence accumulated by tracking. (c) Accuracy can be improved by focusing on highly confident examples (those with large amounts of evidence).

The VECTOR system described in the following section was designed for real-time highway traffic management and analysis and generates loop detector type measurements as well as determines the types of vehicles on the road over time. Accumulated traffic statistics are used to build a traffic model useful for online traffic flow analysis, such as detection of speeding vehicles.

3.1 Vehicle Type Classification

The VECTOR system identifies 8 of the most often occurring vehicle classes based on the 2001 National Household Travel Survey conducted by the U.S. Department of Transportation [2]. Knowledge of vehicle type is essential for a wide range of highway management functions such as estimating emissions [13] or infrastructure load assessment [14]. This detailed real-time fleet composition is currently a missing component is most emission studies.

Information redundancy obtained by the number of different visible views of a vehicle during visual tracking is exploited to improve vehicle type classification. While tracking a vehicle, a number of measurements are taken to describe its shape and appearance and provides a unique signature. This signature is refined with each new video frame to over-

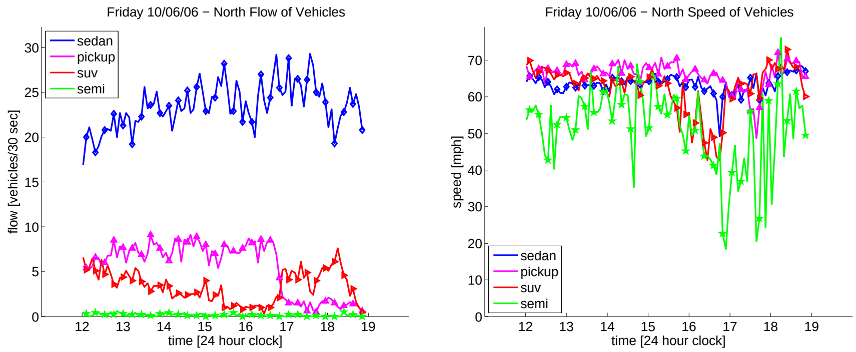

(a) (b)

Figure 8: Traffic statistics (a) flow and (b) speed separated by vehicle type.

come noisy measurements and matched with a vehicle database for improved results. Fig. 7 demonstrates how the classification is improved in consecutive frames as more data is accumulated. Full details of the track based classification scheme can be found in [15, 16].

3.2 Traffic Flow Analysis

The VECTOR system delivers flow ($\frac{\#vehicles}{time}$) and density ($\frac{\#vehicles}{distance}$), similar to loop sensors, as well as speed (MPH) estimates in 30 second intervals. The speed of each measured for each and every vehicle, which is difficult to do with loop detectors. These camera based statistics have a high degree correlation with both the true manually counted flow and PeMS measurements [9]. Since VECTOR does vehicle type classification, it is possible to extract more rich traffic measurements. The flow and speed are compiled for each type of vehicle and show in Fig. 8. These parameters are essential for understanding the differing effects of commercial or private vehicles on highway control, for the study of environmental impact from emissions, and for the estimation of infrastructure and road wear and tear. The efficiency of the highway link is also estimated by analyzing flow and speed characteristics over time to characterize congestion [17].

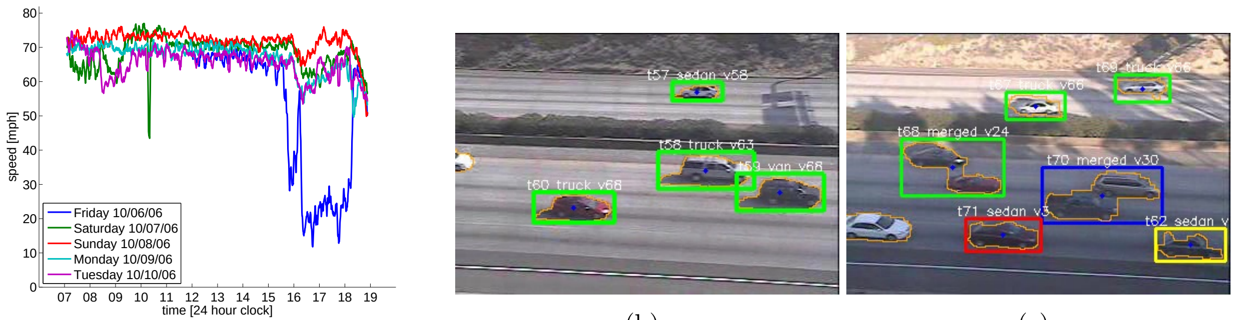

(a) (b) (c)

Figure 9: Speed profiling based on daily models {speeding, normal, slow, stopped} = {blue, green, yellow, red}. (a) Speed characteristics for different days of the week. (b) Normal free flowing traffic. (c) Commuter congestion causes differing characteristics in either highway direction. The normal speed at this hour southbound is significantly slower than northbound.

3.3 Speed Compliance

Daily speed variations can be tracked using historical measurements, Fig. 9(a) shows the speed fluctuations over the course of a week. Notice the significant slowdown during the Friday evening commute not seen on other days. VECTOR uses these daily speed profiles to indicate the motion state of vehicle during online tracking by the color a bounding box {speeding, normal, slow, stopped} = {blue, green, yellow, red}. Sample output frames are shown in Fig. 9. Rather than relying on posted speed limits, speeding vehicles are recognized based on the current conditions. When there is congestion dangerous speeds are significantly lower than the posted limit.

4 Local Measurements

Another key tool for ITS highway monitoring are the vehicles on the road themselves. Everyday cars travel freely along the country's road networks providing greater coverage than any loop or camera sensor could hope to achieve. The situations and interactions encountered by every driver on a daily basis can be leveraged to provide a more complete view of traffic.

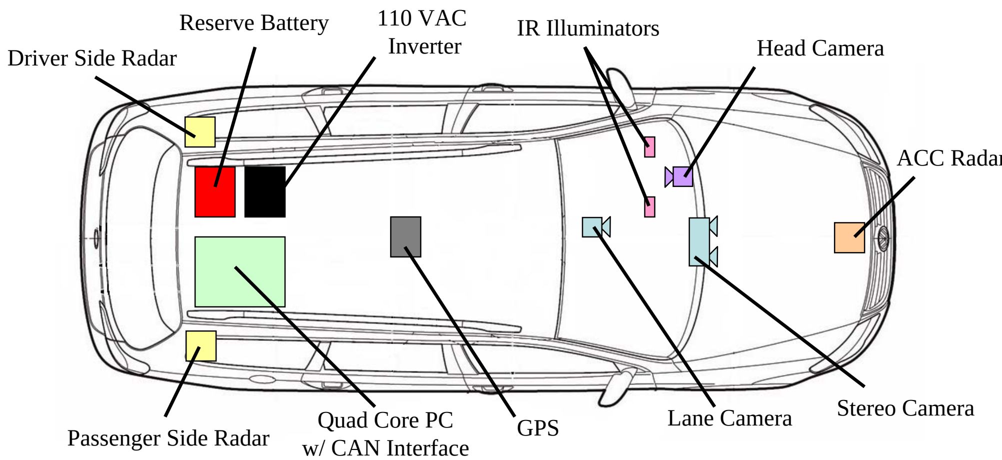

Figure 10: LISAX instrumented vehicle uses radar and camera systems to obtain surround trajectories and road localization information.

Auto companies are equipping vehicles with advanced sensors which can provide a wealth of information about both the current vehicle state as well as local driving conditions. By utilizing these sophisticated sensing devices, new insights can be gleaned and more effective control and safety strategies can be implemented.

4.1 LISA testbeds

The laboratory for intelligent and safe automobiles (LISA) provides unique testbeds to study the local surrounding environment of a vehicle. Three vehicles have been outfitted for synchronous data capture. Each is a mobile computing platform that is able to collect and analyze, in real-time, measurements from the on-board vehicle sensors via CAN (Fig. 12(a)), GPS (Fig. 12(b)), and a host of sensors, *e.g.* video cameras (Fig. 11) and radar, designed to capture the interior and exterior of the car. LISAX, shown in Fig. 10, is the most recent testbed based upon a Volkswagen Passat. There are a number of sensors including a camera system for looking inside at the driver, radar and camera systems to look outside the vehicle, vehicle state sensors sent along the CAN bus, as well as GPS for positioning. The main goal of this setup is to provide enough sensory input to completely understand the vehicle

Figure 11: LISAQ collects video from 8 different cameras which capture the interior and surround of the vehicle.

surround as well as the intentions of the driver. Fig. 11 shows sample video collected from 7 cameras in LISAQ which is based on an Infinity Q45. The video allows monitoring of the vehicle surround as well as the actions of the driver (head, hands, feet).

4.2 Looking Out

One thing to notice in Fig. 6 is the gap between consecutive sensors. Some locations have high loop counts while others are much less dense. Because infrastructure sensors provide such wide coverage, generally of important landmarks, their resolution is coarse. By utilizing road vehicles as high-tech mobile data capture sources, finer resolution and better coverage of the roads is possible. In addition, road safety can be greatly improved by increasing vehicle (and driver) awareness to the surroundings.

Fig. 13 provides a visualization of data collected from LISA. Three cameras provide a view of the driver, out the front windshield, and a of the rear of the car. To the right a selection

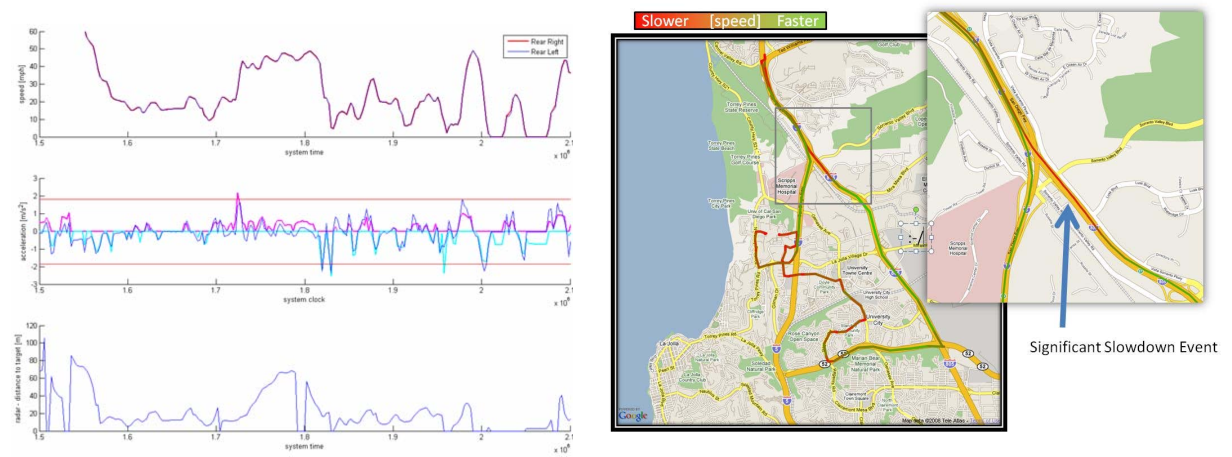

(a) (b)

Figure 12: (a) Vehicle sensor measurements from CAN. (b) GPS based mapping provides complete network coverage. Gaps between infrastructure sensors can be filled with vehicle measurements with finer resolution in traffic speed conditions.

of CAN measurements are displayed which indicate the current state of the vehicle. The two upper panes show the vehicle surround and denotes other vehicles in color. Each obstacle is tracked (see the pink trajectory in the surround zoom pane) and the resulting trajectory is used to assess the driving situation. In this figure, the LISA testbed is being overtaken from behind. Notice the clear distinction between the trajectories of the lead vehicle in red and the overtaking vehicle in pink. Similar to the infrastructure mounted cameras, the trajectories from this moving platform give an indication of behavior. Prototypical driving behaviors can be learned to explain longitudinal and lateral motion which includes acts such as acceleration, braking, lane change, and turns. The future behavior of the surround obstacles as well as the ego-vehicle, since its motion provides another track, can be predicted with the prototype motion models. This prediction allows assessment of the safety of the traffic configuration. Further, the criticality of ego-maneuvers can be evaluated and a driver can be warned when a desired move would result in unsafe outcome.

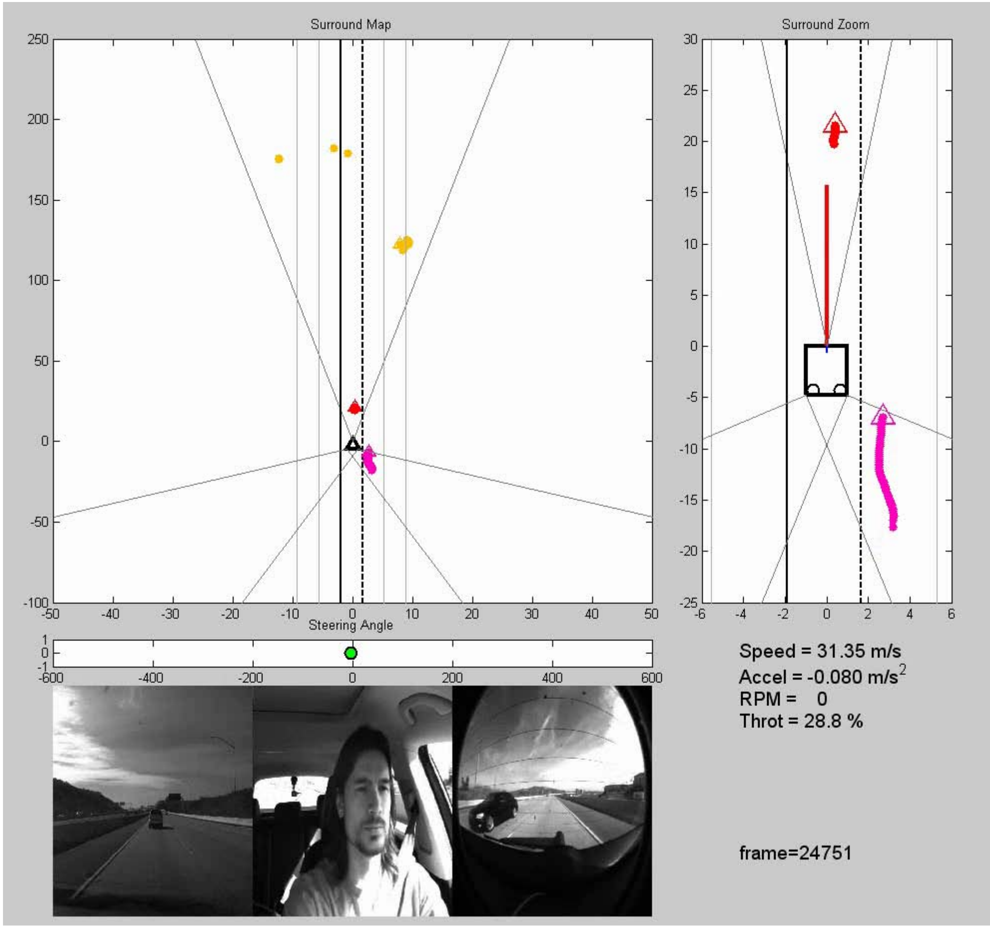

Figure 13: There are clear differences between an overtaking car in pink and the lead vehicle in red. These trajectory differences are used to categorize different driving behaviors.

5 Future Directions

ITS technologies ample opportunities for exciting new research. Continual improvement of systems to ensure real-time, robust, and adaptive solutions is necessary. This incorporates more general computer vision issues such as robust tracking through harsh elements, shadows, and heavy occlusion. Safety can be further improved by defining safety performance measures to asses the the effectiveness of the driver assistance systems that are now being developed [18]. Infrastructure based safety systems will need to convey information to road drivers and in a way that is informative but not distracting necessitating work in communications and human factors. But the future of ITS does not rely solely on computer vision. True ITS systems will require elegant fusion of a wide variety of data sources (integration of VECTOR, loop sensors, and vehicle based sensors). The future of ITS will be shaped by urban and highway planners, scientist, and policy makers all working together to improve travel.

6 Concluding Remarks

In this paper we presented a general framework for integrating infrastructure and vehicle sensing to improve transportation monitoring and driver safety. By analyzing motion, increased awareness and understanding of driving situations is achieved. The behavior of drivers is studied in a large scale through the use of infrastructure sensors and in the local setting by using advanced sensor equipped vehicles and are tied together through maps. We conducted systematic and extensive experimental studies to prove the feasibility an promise of the proposed framework. The future of ITS depends on the integration of these types of systems as the front end data collection tools for higher level understanding of transportation issues.

Acknowledgments

We would like to thank ITS America for the opportunity to take part in the student essay competition. This work has been supported by wide range of agencies and industry partners, National Science Foundation, Caltrans, Department of Homeland Security, Technical Support Working Group , UC Discovery Grants - DIMI, Volkswagen-Audi, Nissan, and Daimler Chrysler. We would like to thank all our sponsors as well as our collaborators and CVRR colleagues for continued assistance.

References

[1] (2009) Highway statistics 2006 - FHWA. US Department of Transportation, Federal Highway Administration. [Online]. Available: http://www.fhwa.dot.gov/policy/ohim/hs06/index.htm

[2] (2001) 2001 national household travel survey. U.S. Department of Transportation (DOT). [Online]. Available: http://nhts.ornl.gov/2001

[3] (2009) FARS encyclopedia. U.S. Department of Transportation - National Highway Traffic Safety Administration. [Online]. Available: http://www-fars.nhtsa.dot.gov/Main/index.aspx

[4] B. Morris and M. Trivedi, "Learning and classification of trajectories in dynamic scenes: A general framework for live video analysis," in *Proc. IEEE International Conference on Advanced Video and Signal based Surveillance*, Santa Fe, New Mexico, Sep. 2008.

[5] ——, "An adaptive scene description for activity analysis in surveillance video," in *Proc. IEEE Inter. Conf. on Pattern Recog.*, Tampa, Florida, Dec. 2008.

[6] W. Hu, X. Xiao, D. Xie, T. Tan, and S. Maybank, "Traffic accident prediction using 3-D model-based vehicle tracking," *IEEE Trans. Veh. Technol.*, vol. 53, no. 3, pp. 677–694, May 2004.

[7] N. Saunier, T. Sayed, and C. Lim, "Probabilistic collision prediction for vision-based automated road safety analysis," in *Proc. IEEE Conf. Intell. Transport. Syst.*, Seattle, Washington, Sep. 2007, pp. 872–878.

[8] (2006) PeMS performance measurement system 7.0. Caltrans, The University of California, Berkeley, and PATH. [Online]. Available: http://pems.eecs.berkeley.edu

[9] B. T. Morris and M. M. Trivedi, "Real-time video based highway traffic measurement and performance monitoring," in *Proc. IEEE Conf. Intell. Transport. Syst.*, Seattle, Washington, Sep. 2007, pp. 59–64.

[10] Y.-K. Jung and Y.-S. Ho, "Traffic parameter extraction using video-based vehicle tracking," in *Proc. IEEE Int'l. Conf. Intell. Transport. Syst.*, Oct. 1999, pp. 764–769.

[11] S. Kamijo, H. Koo, X. Liu, K. Fujihira, and M. Sakauchi, "Development and evaluation of real-time video surveillance system on highway based on semantic hierarchy and decision surface," in *IEEE Int'l. Conf. Systems, Man, and Cybernetics*, Oct. 2005, pp. 840–846.

[12] B. T. Morris and M. M. Trivedi, "Robust classification and tracking of vehicles in traffic video streams," in *Proc. IEEE Conf. Intell. Transport. Syst.*, Toronto, Canada, Sep. 2006, pp. 1078–1083.

[13] C. Cardelino, "Daily variability of motor vehicle emissions derived from traffic counter data," *Journal of the Air and Waste Management Association*, vol. 48, no. 7, Jul. 1998.

[14] (2008) Traffic monitoring guide. U.S. Department of Transportation - Office of Highway Policy Information. [Online]. Available: http://www.fhwa.dot.gov/ohim/tmguide/

[15] B. T. Morris and M. M. Trivedi, "Improved vehicle classification in long traffic video by cooperating tracker and classifier modules," in *Proc. IEEE International Conference on Advanced Video and Signal based Surveillance*, Sydney, Australia, Nov. 2006, pp. 9–14.

[16] ——, "Learning, modeling, and classification of vehicle track patterns from live video," *IEEE Trans. Intell. Transp. Syst.*, vol. 9, no. 3, pp. 425–437, Sep. 2008.

[17] C. Chen, Z. Jia, and P. Varaiya, "Causes and cures of highway congestion," *IEEE Control Syst. Mag.*, vol. 21, no. 6, pp. 26–32, Dec. 2001.

[18] C.-Y. Chan, "Defining safety performance measures of driver-assistance systems for intersection left-turn conflicts," in *Proc. IEEE Intell. Vehicles Symposium*, Jun. 2006, pp. 25–30.