

**Innovations Deserving
Exploratory Analysis Programs**

The logo features the word "IDEA" in a large, bold, serif font. A light gray rectangular box is positioned behind the letters "I" and "D". Two thin, dotted lines extend from the bottom corners of this box, angling downwards and outwards to the right.

IDEA

Highway IDEA Program

***COMPUTER VISION TRAFFIC SENSOR
FOR FIXED AND PAN-TILT-ZOOM CAMERAS***

Final Report for Highway IDEA Project 140

Prepared by:

Stan Birchfield, Wayne Sarasua, and Neeraj Kanhere, Clemson University, Clemson, SC

August 2010

TRANSPORTATION RESEARCH BOARD
OF THE NATIONAL ACADEMIES

INNOVATIONS DESERVING EXPLORATORY ANALYSIS (IDEA) PROGRAMS MANAGED BY THE TRANSPORTATION RESEARCH BOARD (TRB)

This NCHRP-IDEA investigation was by Research & Technology Corp. completed as part of the National Cooperative Highway Research Program (NCHRP). The NCHRP-IDEA program is one of the three IDEA programs managed by the Transportation Research Board (TRB) to foster innovations in highway and intermodal surface transportation systems. The other two IDEA program areas are TRANSIT-IDEA, which focuses on products and results for transit practice, in support of the Transit Cooperative Research Program (TCRP), and ITS-IDEA, which focuses on products and results for the development and deployment of intelligent transportation systems (ITS), in support of the U.S. Department of Transportation's national ITS program plan. The three IDEA program areas are integrated to achieve the development and testing of nontraditional and innovative concepts, methods, and technologies, including conversion technologies from the defense, aerospace, computer, and communication sectors that are new to highway, transit, intelligent, and intermodal surface transportation systems.

For information on the IDEA Program contact IDEA Program, Transportation Research Board, 500 5th Street, N.W., Washington, D.C. 20001 (phone: 202/334-1461, fax: 202/334-3471, <http://www.nationalacademies.org/trb/idea>)

The project that is the subject of this contractor-authored report was a part of the Innovations Deserving Exploratory Analysis (IDEA) Programs, which are managed by the Transportation Research Board (TRB) with the approval of the Governing Board of the National Research Council. The members of the oversight committee that monitored the project and reviewed the report were chosen for their special competencies and with regard for appropriate balance. The views expressed in this report are those of the contractor who conducted the investigation documented in this report and do not necessarily reflect those of the Transportation Research Board, the National Research Council, or the sponsors of the IDEA Programs. This document has not been edited by TRB.

The Transportation Research Board of the National Academies, the National Research Council, and the organizations that sponsor the IDEA Programs do not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to the object of the investigation.

COMPUTER VISION TRAFFIC SENSOR FOR FIXED AND PAN-TILT- ZOOM CAMERAS

IDEA Program Final Report for Award NCHRP-140

Prepared for the IDEA Program

Transportation Research Board

The National Academies

Stan Birchfield, Wayne Sarasua, and Neeraj Kanhere

Clemson University, Clemson, SC

August 2010

ACKNOWLEDGMENTS

The investigators would like to thank the officials at the New York State Department of Transportation, Maryland Department of Transportation, and South Carolina Department of Transportation, whose support and effort made this project possible. Gratitude is also expressed to the graduate students who spent many hours of tedious labor annotating the video data to enable ground truth comparison of the sensors.

CONTENTS

EXECUTIVE SUMMARY	1
IDEA PRODUCT	3
CONCEPT AND INNOVATION.....	5
COMPUTER VISION-BASED APPROACH.....	6
AUTOMATIC CAMERA CALIBRATION.....	7
INVESTIGATION	8
TESTBED DESIGN	8
MEETINGS.....	8
SENSOR INSTALLATION.....	9
<i>Maryland installation</i>	<i>9</i>
<i>New York installation</i>	<i>10</i>
<i>South Carolina roadside setup.....</i>	<i>10</i>
<i>Other installations.....</i>	<i>10</i>
SENSOR REFINEMENT.....	11
<i>Camera issues</i>	<i>11</i>
<i>Algorithm improvements.....</i>	<i>12</i>
<i>Solar power issues.....</i>	<i>13</i>
SENSOR OPERATION.....	15
DATA COLLECTION AND GROUND TRUTH.....	17
ANALYSIS AND EVALUATION OF VISION SENSOR.....	18
<i>Volume accuracy</i>	<i>18</i>
<i>Speed and classification accuracy</i>	<i>27</i>
<i>Additional results in a variety of challenging scenarios</i>	<i>29</i>
PLANS FOR IMPLEMENTATION	30
CONCLUSION	31
REFERENCES	33

EXECUTIVE SUMMARY



This project, entitled *Computer Vision Traffic Sensor for Fixed and Pan-Tilt-Zoom Cameras*, involved the development, installation, and evaluation of a next generation computer vision traffic sensor capable of collecting traffic parameters such as count, speed, and classification data from a single camera sensor. In contrast to existing video-based commercial systems relying upon *virtual detectors* (i.e. localized presence detection), the proposed sensor relies upon recent developments in computer vision to actually track vehicles through the camera's field of view. This tracking technology enables the system to overcome many of the limitations of previous approaches, such as distractions due to reflections, shadows, perspective effects (occlusions and spillover), and congestion. Some highlights of the approach include its ability to

- operate in a variety of lighting and environmental conditions;
- operate with cameras that are high above the ground (50-70 feet), as well as low to the ground (25-30 feet);
- operate with fixed or pan-tilt-zoom cameras;
- dynamically recalibrate itself when the camera undergoes pan and tilt movements, after an initial simple six-click calibration procedure (without the delicate heuristics required for virtual detectors);
- work in congested scenarios; and
- produce accurate count, speed, and classification data.

The figure above shows the versatility of the system operating in a variety of camera configurations, road characteristics, lighting conditions, and environmental scenarios.

The system was installed at two permanent automatic traffic recorder (ATR) sites in Maryland and New York, as well as one Traffic Management Center (TMC) in South Carolina. From these installations, as well as from a temporary roadside setup in South Carolina, hundreds of hours of video data have been collected for analysis and testing. The output of the vision-based system has been extensively compared with those of the two existing ATR sensors in Maryland and New York in order to estimate its accuracy. In addition, to facilitate a more objective evaluation, more than 72 hours of video have been manually annotated to generate ground truth data with which to compare both the vision sensor and the existing ATR sensors.

Overall, the system was found to be able to produce hourly per-lane counts with 5-10% error, with the error reducing to less than 0.5% when aggregated over a 24-hour period. The system estimates instantaneous per-vehicle speeds with less than 10 mph error, which when averaged over several minutes reduces to less than 1 mph error. Using three length-based bins, the system classifies vehicles with less than 3% error. The system's performance was measured under different environmental conditions, at different times of day/night, with and without shadows, and with cameras both low and high off the ground. The accuracy of the vision-based system was found to be roughly equivalent to the existing ATR sensors at both locations in terms of accuracy in estimating volume, speed, and three-bin classification.

The system was tested with a camera height ranging from 26 feet to 70 feet. At higher camera locations, the perspective effects are minimized, occlusion is reduced, and more lanes of traffic fall within the field of view, thereby enabling the system to more accurately produce counts and speeds. At 70 feet, the system has been demonstrated on six lanes of traffic (both directions) simultaneously. As the height of the camera is reduced to a minimum of 26 feet, the number of lanes that can be covered decreases to three lanes, but surprisingly the error in count, speed, and classification results are not measurably affected. In fact, classification is more accurate at lower heights due to the increased pixel resolution that facilitates more accurate estimation of vehicle lengths. This ability of the system to operate at such low heights is unprecedented and represents a significant departure from previous video-based systems.

The system was tested with the camera placed so that the distance to the nearest road edge ranged from 20 to 46 feet. Not surprisingly, results improved as this distance is decreased, due to the reduced occlusion when the camera is closer to the road. However, it was found that there is actually an advantage to placing the camera on the side of the road rather than directly over the middle lane, because the side view facilitates measuring vehicle lengths.

Other factors affect performance of the system. At night, care must be taken to turn off the automatic gain control and to increase the shutter speed in order to stabilize the image brightness and to reduce blooming, respectively. Even so, classification at night is not currently possible due to the lack of image data when no ambient light is present. In addition, when the camera exhibits significant lens distortion, the length measurements can be affected, thus reducing accuracy in classification and speeds. A final factor to consider is congestion.

Compared with existing sensors such as loop detectors, the system exhibited remarkably improved performance in highly congested situations, which is another significant advantage of the system. However, congestion does present a challenge for the system, particularly when the occlusion resulting from a low camera viewing multiple lanes prevents the system from seeing vehicles in the farthest lanes. In such situations, accuracy drops noticeably, though the system was still found to outperform the installed loop detector by a significant amount.

Considerable progress has been made toward producing a system that is capable of fully automatic calibration. At installation or setup time, the camera is calibrated manually using a simple six-click calibration procedure that does not involve any of the delicate heuristics so common with virtual detection-based approaches. Then, when the camera's pan or tilt changes, the system is able to automatically recalibrate itself. However, recalibration under changing zoom has proved to be more difficult than originally anticipated – this problem will be addressed by future research.

IDEA PRODUCT

The purpose of this project was to develop a next generation vision-based traffic sensor to detect, track, and classify vehicles using both fixed and pan-tilt-zoom (PTZ) cameras. In contrast with existing machine vision-based sensors, the proposed sensor uses recent developments in computer vision to actually track vehicles through the video sequence, thus providing more robust and versatile behavior, able to work with low-angle cameras or existing pan-tilt-zoom cameras. The system is expected to impact two application scenarios important for transportation officials: roadside data collection and real-time traffic management. The first scenario is illustrated in Figure 1. In this scenario, a camera placed beside the road, and its video feed is sent to a local processor to gather traffic parameters (volume, speed, and classification). The user can then remotely access this information via a wireless modem. The camera can be either digital (network IP) or analog, it can be fixed or pan-tilt-zoom, and it can be supplied with either solar DC power or AC power. One important and novel feature of the system is its ability to work with cameras that are installed low to the ground (25-30 feet), thereby reducing installation costs substantially and making it feasible to support transient studies via a portable system.

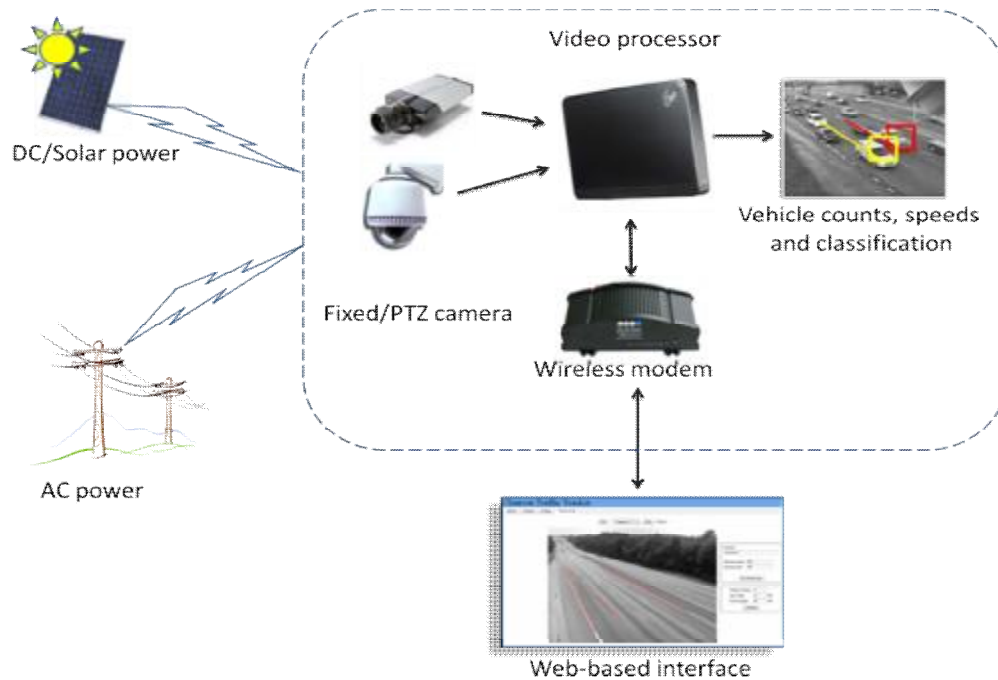


FIGURE 1: Roadside camera scenario in which the video feed from a camera is analyzed in real time by a local processor to estimate traffic parameters (volume, speed, and classification).

The second scenario involves processing video data obtained by existing cameras. There are literally thousands of such cameras installed throughout the nation, primarily in urban areas, which are used for the sole purpose of manual verification of incident reports from other sensors. For the organizations that are currently relying on traditional machine vision-based sensors to obtain this information, they must install a separate set of cameras for this purpose, which substantially increases both installation and maintenance costs. The proposed sensor is designed to process video data from existing cameras, thereby obtaining real-time information about per-lane traffic counts and speeds. This information can then be used to alert Traffic Management Center (TMC) operators to incidents with extremely low latency compared with other sensors. An important aspect of this approach is the ability of the sensor to automatically recalibrate when the user has changed the pan or tilt of the camera. This scenario is illustrated in Figure 2, which shows a computer server processing multiple video feeds (either analog or digital) from existing cameras whose data is being fed already to a central TMC. The user can access this information via a web browser either locally or remotely.

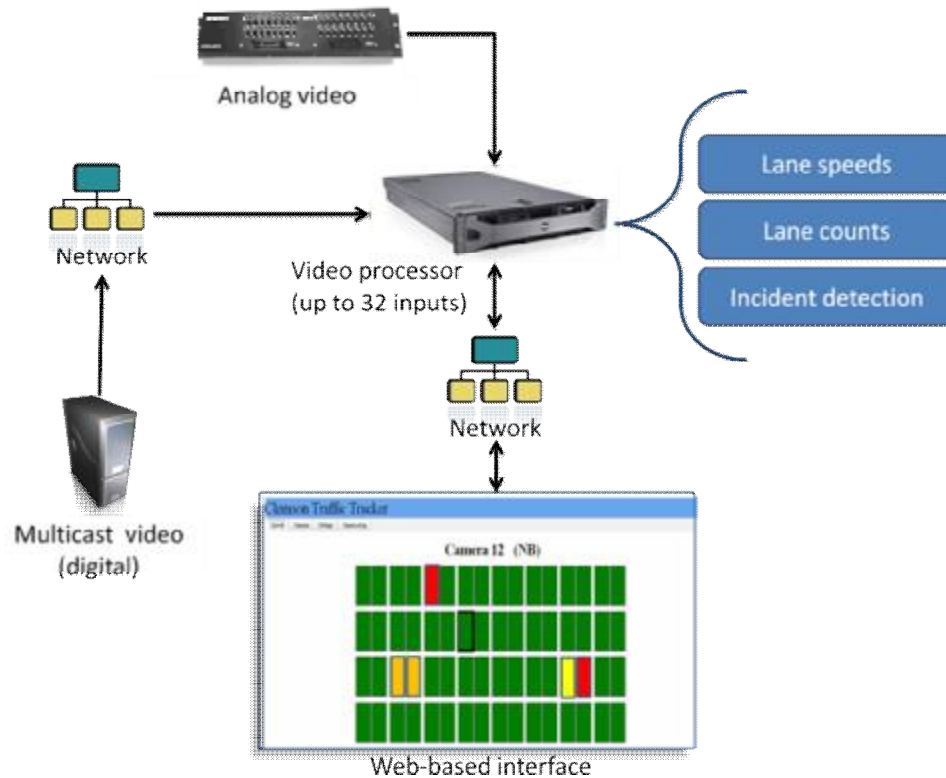


FIGURE 2: TMC scenario. Multiple video feeds from existing cameras are fed to a video processor on a central server which analyzes the video to determine lane counts and speeds for real-time incident detection.

CONCEPT AND INNOVATION

Traffic counts, speeds, and vehicle classification data are fundamental for a variety of transportation applications ranging from transportation planning to modern intelligent transportation systems (ITS), as well as for assessing the condition of our highway system (1,2,3). Traditional sensors such as inductive loop detectors, microwave radar, infrared devices, piezos and road tube sensors typically perform well in low to moderate congestion but deteriorate as highways reach capacity and require costly periodic maintenance(4,5,6,7,8). Many commercially available machine vision-based systems rely on *virtual detectors* that require fixed cameras mounted high above the road, have limited vehicle classification capabilities, and are prone to errors caused by vehicle occlusion or spillover into adjacent lanes when cameras are mounted relatively low to the ground. In fact, poor accuracy was one of the reasons why video image detection received low ratings in a survey of user satisfaction(9).

COMPUTER VISION-BASED APPROACH

In contrast with the traditional approach of virtual detectors, the proposed sensor builds upon recent developments in computer vision. By using a combination of edge detection, feature tracking, and vehicle base-fronts to detect and track vehicles, the robust tracking approach overcomes several of the limitations of machine vision sensors. In particular, the sensor is versatile, able to work even with cameras mounted as low as 25 feet above the ground because it properly handles the perspective effects that cause occlusion and spillover, as demonstrated in Figure 3. The system was originally introduced several years ago in (10), and since then various improvements have been demonstrated and incorporated, such as real-time performance (11,12,13), the ability to classify motorcycles (14), and automatic calibration (15,16).



FIGURE 3: Because of its ability to properly handle perspective effects, the proposed sensor is not prone to over-counting caused by spillovers. (a) With a traditional machine-vision sensor, a large vehicle triggers multiple virtual detectors when the camera is low to the ground. (b) With the proposed computer-vision sensor, the vehicle is correctly tracked as a single entity.

One of the noteworthy characteristics of the proposed system is its versatility. The reliance upon low-level features such as intensity edges makes the system robust to changing lighting and environmental conditions, and it also allows the system to operate under congested conditions which have been problematic for traditional machine vision sensors. Its ability to work with cameras that are low to the ground facilitates the use of much shorter poles than those required by existing commercial machine vision systems, opening up the possibility of a portable video-based system. The use of shorter poles will significantly reduce cost, simplify installation and maintenance, and improve safety because the reduced mass will be less dangerous to an unfortunate vehicle that strikes the pole. Moreover, the unique approach of tracking vehicles allows the system to collect parameters that are unavailable with current technology, such as space-mean speed. Further, because of the longer detection zones, acceleration and deceleration characteristics can be collected, as well as specific lane-change maneuvers. Beyond the specific project objectives, it is anticipated that the system presented here will serve as a foundation for additional capabilities in the future, such as incident detection, collecting acceleration and deceleration characteristics, measuring lane-change maneuvers, counting turn movements, monitoring pedestrian activity, and detecting dangerous or erratic driver behavior for the purpose

of alerting workers to improve work zone safety. These capabilities will be made possible by the sensor's novel approach of detecting and tracking individual vehicles over time.

The system is currently capable of classifying vehicles into three length-based classes. Although the Federal Highway Administration (FHWA) requires 13 classes for reporting Highway Performance Monitoring System (HPMS) data, the *Highway Capacity Manual* (HCM) includes only three classes for doing capacity analysis(17). The *FHWA Traffic Monitoring Guide* (TMG) suggests that a broad 3 to 4 class scheme can meet most of the data demands in traffic analysis by using factors from axle studies to convert classification data to 13 classes (2). With such conversion procedures, the current system is expected to be useful for a variety of applications even with the limited number of classes.

AUTOMATIC CAMERA CALIBRATION

Commercial machine vision sensors require the user to manually draw virtual detection zones at setup time. For such systems the cameras must be fixed in place, because any movement of the camera would require the virtual detection zones to be manually redefined – a tedious process. The proposed sensor takes a first step toward overcoming these limitations by including an automatic calibration module that uses image processing to estimate the two vanishing points, from which the camera height, focal length, and pan and tilt angles are calculated. This automatic calibration is envisioned not only to ease the installation of fixed or temporary cameras, but it also promises to open the possibility of *automatically* analyzing video from pan-tilt-zoom (PTZ) cameras. With automatic calibration, the system would be able to gather traffic parameters from the thousands of existing cameras that are currently being used only for manual surveillance of traffic operations, leading to a cost-effective way to greatly increase data collection by utilizing existing hardware.

Significant progress has been made toward developing an automatic calibration module. The investigators have conducted a thorough analysis of the problem of road-side camera calibration by developing a taxonomy and an experimental comparison of various techniques (16). Moreover, an automatic calibration algorithm has been developed that will recalibrate the camera after its pan or tilt angles have changed, assuming that the camera has been manually calibrated at setup time. See Figure 4. This manual calibration procedure requires just six mouse clicks to define the edges of the road and a single known length, thus greatly simplifying the process over the heuristic-based detection zone identification of existing machine vision-based sensors. The problem of fully automatic calibration even when the camera zoom has changed has proved more difficult than originally envisioned and will be investigated in future work.



FIGURE 4: When the PTZ camera is moved from its current view (left) to a new view (middle), the vision-based sensor automatically recalibrates after a few seconds (right).

INVESTIGATION

TESTBED DESIGN

At project initiation, the members of the team met to discuss the various software and hardware components of the system to be installed at locations in Maryland and New York. It was decided to use two different cameras for the two locations to provide some experience with the different hardware to guide future purchasing decisions with regard to the importance of pan-tilt-zoom capability and digital versus analog interface. Thus, an analog pan-tilt-zoom camera was selected for the Maryland installation, while a fixed digital network camera was selected for New York. For both locations, a low-power computer capable of running a standard operating system was chosen to minimize issues involving porting the software.

MEETINGS

The advisory committee met on May 19, 2009 on the Clemson University campus. A total of 13 members participated in the meeting, some attending in person while others relied on teleconferencing and video-conferencing. The meeting covered a summary of the project, a presentation demonstrating the idea behind the proposed sensor, issues regarding the solar-powered installation, possible locations for a test-bed installation in South Carolina, and accuracy requirements of the sensor. Feedback from committee members was greatly helpful in steering the future decisions of the project.

On Jan 11, 2010 the team met in Baltimore with officials from the Maryland State Highway Administration and presented qualitative and quantitative results of the year-long sensor tests. As shown later in this report, the sensor at that time had been able to capture volume information 24 hours a day with accuracy slightly better than the existing loop detector at the ATR site. During this meeting the SHA officials showed interest in testing the sensor

further at a different location, perhaps along the Capital Beltway due to heavy traffic and congestion.

Another advisory committee meeting was held at the Traffic Management Center (TMC) in Columbia, SC. on Jan 28, 2010. Several members of the advisory committee were present, including SCDOT officials from both the TMC and data collection divisions. The meeting involved a presentation summarizing the technology behind the proposed sensor and the results of the sensor tests, similar to what was shared in Maryland. Discussion during the meeting revealed several suggestions for future improvements and also included the possibility of installing a roadside video-based ATR testbed in South Carolina, similar to those in Maryland and New York.

In addition to these stated meetings, the members of the team were in frequent contact with personnel in Maryland, New York, and South Carolina throughout the project via email and phone calls. Additional correspondence with personnel in other cities and states provided additional information for guiding decisions regarding project direction.

SENSOR INSTALLATION

The proposed sensor was installed at the two test locations in Maryland and New York, with cooperation of personnel at the respective DOTs. This involved mounting the cameras, setting up power and communications, and configuring the software. Both installations were designed to allow for remote access from the investigators' offices in Clemson.

Maryland installation

The first sensor was installed in Columbia, MD, on December 16, 2008 along Hwy-29, Northbound at Gates Lane (+39° 12' 37.60", -76° 51' 14.86"), as shown in Figure 5. The team installed a pan-tilt-zoom (PTZ) dome camera, a frame grabber to convert the analog video signal from the camera into a digital data stream, a processing board to run the algorithm on the real-time video, and a wireless modem for communications and for monitoring the sensor remotely.




FIGURE 5: Installation at Maryland along Hwy 29. The red oval indicates the video processor.

New York installation

The second sensor was installed in Long Island, NY on December 17, 2008 along I-495 East, near exit 23 ($+40^{\circ} 44'25.96''$, $-73^{\circ} 49' 13.13''$), as shown in Figure 6. The only source of power available at this location is a pair of solar panels mounted on the pole. The installation in New York was more time-consuming than the one in Maryland, due to the need to manually aim and focus the fixed camera. To conserve power, a USB-controlled relay was used to programmatically control the heater/fan in the camera housing.



FIGURE 6: Installation in New York along I-495. The red oval indicates the video processor.

South Carolina roadside setup

Instead of mounting a third permanent installation, a temporary system was set up in South Carolina along I-85 next to a rest area ($+34.573101, -82.741184$), by installing a camera to a retractable mast on a transportation van, shown in Figure 7. Data was collected continuously for 24 hours for later processing, from a location significantly farther from the road than the systems in either Maryland or New York.



FIGURE 7: Temporary setup in South Carolina along I-85.

Other installations

On July 27, 2010, the software was installed on a server located in the TMC in Columbia, SC. This server has access to hundreds of existing TMC cameras and will enable the vision-based system to be refined and tested on a wide scale. Prior to this installation, video data was

collected from 54 existing cameras that feed into the TMC in Greenville, SC. In addition, the investigators have been in contact with personnel in several other cities and states both for the purpose of collecting data for further testing, as well as to discuss possible future installations.

SENSOR REFINEMENT

After the installations in Maryland and New York, both the hardware and software components of the two systems were refined over the next several months to address various issues that arose.

Camera issues

At installation time the Maryland camera was oriented to span a large area in its field of view. However, it was noticed later that such an orientation causes severe degradation of the image quality when the sun is low on the horizon. Thankfully, the PTZ capability of the camera made it easy to correct this problem by tilting the camera down to reduce glare, as shown in the top of Figure 8. Similarly, the initial accuracy of the sensor at night was found to be far below the expected value. Closer inspection revealed that the problem was due to the default settings on both cameras, which were geared toward providing as much detail as possible in low-light situations. The algorithm, however, relies on detecting and tracking the headlights in its nighttime mode of operation. As a result, the adverse effect of headlight-blooming and image saturation prevented the algorithm from distinguishing headlights from the surrounding areas. The camera settings were changed to provide images more apt for the sensor, as shown in the bottom of Figure 8.



FIGURE 8: Tilting the camera down to reduce glare (top), and decreasing shutter speed to reduce blooming (bottom).

Another issue that arose is the automatic switching of the camera to a low-light mode of operation when the light intensity in the scene falls below a certain internal threshold, which is a characteristic of the camera installed in Maryland. This switching results in a sudden change in illumination, as shown in Figure 9. The algorithm was modified to handle this sudden change in illumination without degrading the performance of the tracker.



FIGURE 9: An abrupt change in the overall brightness of the image occurs when the camera switches between regular and low-light mode. In less than one second the camera switched from regular (left) to low-light mode (middle). Several minutes later it switched back to regular mode (right).

A final cause of concern involved the default settings of the thermostat in the camera housing in New York, which caused the batteries to drain quickly. To alleviate this problem, a software-controlled relay was installed to override the thermostat, resulting in considerable power savings. However, this change occasionally caused condensation inside the camera housing, resulting in poor quality images as shown in Figure 10. Turning the heater back on alleviated the problem.



FIGURE 10: The blurred image resulting from leaving the heater off for more than 24 hours (left) became sharper (middle and right) as a result of turning on the heater.

Algorithm improvements

The installations in Maryland and New York provided a realistic suite of data over a variety of lighting and environmental conditions in which to test and refine the software. Over the course of the project the approach can be modified considerably, introducing a number of improvements to increase the robustness of the system. Particular attention was paid to ensure proper handling of the long shadows and changing illumination conditions that occur in the early morning and late afternoon hours near sunrise and sunset. Throughout the software development process, a change management system was used to keep track of new versions of the software,

and the remote access was used to upload new versions to the two installation locations when they became available. Manually annotated ground truth data, described later, were used to ensure that new versions achieved higher accuracy than previous ones.

Figure 11 shows the improvements in the sensor accuracy (as a result of updates to the algorithm) over the several months in 2009 based on the data collected at Maryland. Each point along the plot is the mean of the absolute difference between the loop-detector counts and vision-sensor counts over an entire month calculated for each hour. After several months of algorithm refinement, the system noticeably improved. Note, however, that these plots *do not* show the error in the vision sensor but only the difference between the two sensors. The actual error will be shown in the quantitative analysis later in the report.




FIGURE 11: Mean absolute difference between loop-detector counts and vision-sensor counts for each lane at the location in Maryland (Top: slow lane, Bottom: fast lane), shown for three months (March, June, and September 2009) with successive improvements in the algorithm. Each data point shows the mean difference for that particular hour over all the days in the respective month.

Solar power issues

Since its installation, the solar-powered sensor in NY had trouble operating consecutively for more than a few days at a time. The solar controller was configured to disconnect the load when the voltage across the batteries drops below 11.5 V, and to reconnect the load when the batteries charge to 12.5 V. A voltage-logging device was installed to monitor the voltage across the batteries. The plots in Figure 12 demonstrate that the batteries were drained faster than the solar panels were able to recharge them. As a result the voltage continued to drop each day until the controller disconnected the load. Once this occurred, it required approximately four days of charging the batteries before the controller reconnected the load.




FIGURE 12: The battery voltage in NY over the period of a week in 2009. The voltage continued to drop until the solar controller disconnected the load on March 26th. After the batteries were recharged, the controller reconnected the load on March 31st.

To analyze the observed behavior of the sensor in NY and to study the power consumption issues, a local setup was implemented similar to the one in NY. The required solar panels were loaned by a fellow faculty member in the ECE department at Clemson. Instead of using the same processing board used in NY, an even lower-power processing board was used. The new board reduced the power required for processing the video data by a four-fold factor. The local sensor shown in Figure 13 functioned continuously without power loss for months at a time with half the solar panels as the system in NY. This success motivated a complete redesign of the sensor in NY to reduce the power consumption, including removing the power losses resulting from DC-AC conversions.



FIGURE 13: A local test sensor with a solar panel and a revised version of the sensor with an even lower-power processor.

On July 16, 2009, the processing board for the NY sensor was replaced by this more power-efficient board. However, even with the new board and simpler design, the sensor still had difficulty operating continuously without losing power. The NYSDOT was extremely helpful in analyzing the problem and performed a short-circuit tests on the solar panels. From these tests, it was discovered that the effective output from the panels was only a fraction of the rated expected value (1.25A short-circuit current as opposed to 8A mentioned in the technical specifications of the panels). The primary cause of this loss appears to be the fact that the solar panels were placed in the shade of some nearby trees, which prevents the panels from delivering sufficient current to keep even the new lower-power sensor operating continuously.

Despite the difficulty of achieving continuous operation, the NY installation should be considered a success. The primary purpose of the installation was to provide a testbed for collecting data and demonstrating the capability of the vision-based sensor to produce accurate results in the midst of severe congestion. Both of these objectives were met, as seen in the results later in this report. Moreover, the new, simplified design was successfully tested on a local setup, thus ensuring that a solar-power-based installation will be possible at any location where sufficient sunlight is available, when the opportunity arises again in the future. However, additional resources were not put into solving the particular problem in NY due to the identification of a commercialization strategy that will not involve solar-powered systems in the near term.

SENSOR OPERATION

This section briefly describes the sensor operation, as well as some qualitative results. As shown in Figure 14, the user calibrates the system using the following procedure:

- 1) The user specifies the number of lanes, the lane width, and a known length measurement.
- 2) The user clicks four times along the edges of the road (two clicks for each edge – red lines).
- 3) The user clicks two points in the road corresponding to the known length measurement (green line).

With the above information the sensor calculates the mapping between the world and image coordinate systems which enables it to measure distances in the road plane for tracking vehicles and measuring speeds. The sensor conveys this mapping to the user by drawing a cuboid as shown in the rightmost image of the figure. This quick visual feedback helps to inform the user if a mistake was made in any of the steps above.



FIGURE 14: Simple six-click calibration procedure.

The output from the sensor is stored as a text file in comma separated value (CSV) format. A new file is created for each day beginning at midnight. The file has an entry for each vehicle detected and tracked by the sensor. The entry consists of the vehicle's id, timestamp, lane, class, and speed. In addition the sensor also stores a smaller file for each day containing user-defined interval counts.

Some sample images demonstrating the results of the system in Maryland are shown in Figure 15. The middle image in the top row shows the robustness of the sensor to the glare resulting from smudges on the camera's dome. The bottom-left picture shows vehicles being tracked on snow-covered road. Notice in the bottom-center image that vehicles are tracked in spite of an icicle formed on the dome of the camera. The color of each rectangle indicates the class of vehicle, with green representing a passenger vehicle and red indicating a truck.



FIGURE 15: Snapshots of the results of the sensor showing the tracked vehicles in variety of weather and lighting conditions in Maryland. Green rectangles indicate passenger vehicles, while red indicates trucks.

Figure 16 shows some sample output frames from the sensor installed in New York. Because poor weather is correlated with reduced solar input, the sensor tends to lose power when

the weather is adverse. As a result, all of the data collected from New York was obtained in good weather conditions.



FIGURE 16: Sample output frames from the sensor in New York.

To show the versatility of the sensor, Figure 17 displays some sample frames obtained by other cameras. Two of these images (left and right) were captured by a pan-tilt-zoom TMC camera in Greenville, SC that is approximately 70 feet above the road. The middle image was captured by a portable tripod camera setup that was temporarily installed at Myrtle Beach, SC during bike week in order to test the sensor’s ability to classify motorcycles (shown in blue). This analysis, published in (14), showed that the sensor was able to achieve motorcycle volume within 6% of the actual motorcycle count on a large dataset involving thousands of motorcycles.



FIGURE 17: Results of testing the sensor at other locations. Notice the sensor’s ability to detect motorcycles (shown in blue) as well as tracking vehicles on both sides of a multi-lane highway with a single camera.

DATA COLLECTION AND GROUND TRUTH

For this report, several days’ worth of video data was collected at both the Maryland and New York sites, as well as along I-85 in South Carolina, and using one of the Greenville TMC cameras (selected somewhat arbitrarily as a representative), as mentioned earlier. As shown in Table 1, together these four cameras span a range of scenarios of both camera placement and road size. Due to its higher placement above the ground, the Greenville TMC camera was able to track vehicles on both the north- and south-bound lanes, whereas the other locations were only able to achieve results for the near side of the road.

	height of camera (feet)	distance to road edge (feet)	number of lanes
MD	28.4	20.5	2
NY	26.3	30.0	3
I-85	28.0	45.9	2
Greenville TMC	70.4	31.8	6

TABLE 1: Parameters of the four camera setups used in the quantitative analysis of the system. Shown are the height of the camera above the road plane, the perpendicular distance on the road plane from the camera to the edge of the nearest lane, and the number of lanes processed by the system.

Co-located with the vision system in Maryland is an ATR with a loop detector providing counts, while the system in New York is co-located with an ATR containing both a loop detector and a piezo sensor, thus enabling counts, speeds, and 13-category classification results to be obtained at that location. While data provided by these existing sensors allowed the performance of the vision-based sensor to be compared with their performance, discrepancies between them are difficult to interpret since it is not known whether the difference is due to error in one sensor or the other (or both). To overcome this problem, graduate students at Clemson supported by the project were assigned the tedious task of manually counting and classifying the vehicles in the various videos. At present, more than 72 hours of data have been manually annotated. This ground truth is indispensable for objectively calculating the accuracy of any sensor, as it provides a way of comparing sensors to the actual quantity intending to be estimated.

ANALYSIS AND EVALUATION OF VISION SENSOR

In this section, quantitative results are presented by comparing the output of the vision-based sensor with ground truth and with the existing ATR sensors (where available). An extensive comparison is made for counts, followed by additional experiments comparing speeds and classification.

Volume accuracy

Figures 18 and 19 show the output of the loop detector and the vision-based sensor in Maryland for two typical days – one weekday and one weekend day – using hourly bins due to the granularity of the loop detector setup. The plots on the left display the hourly counts reported by the loop detector and our vision sensor, while the plots on the right show the absolute value of the percentage difference between the two outputs, along with a 10% dashed line for reference. For most hours, the difference in per-lane volumes is well below 10%.

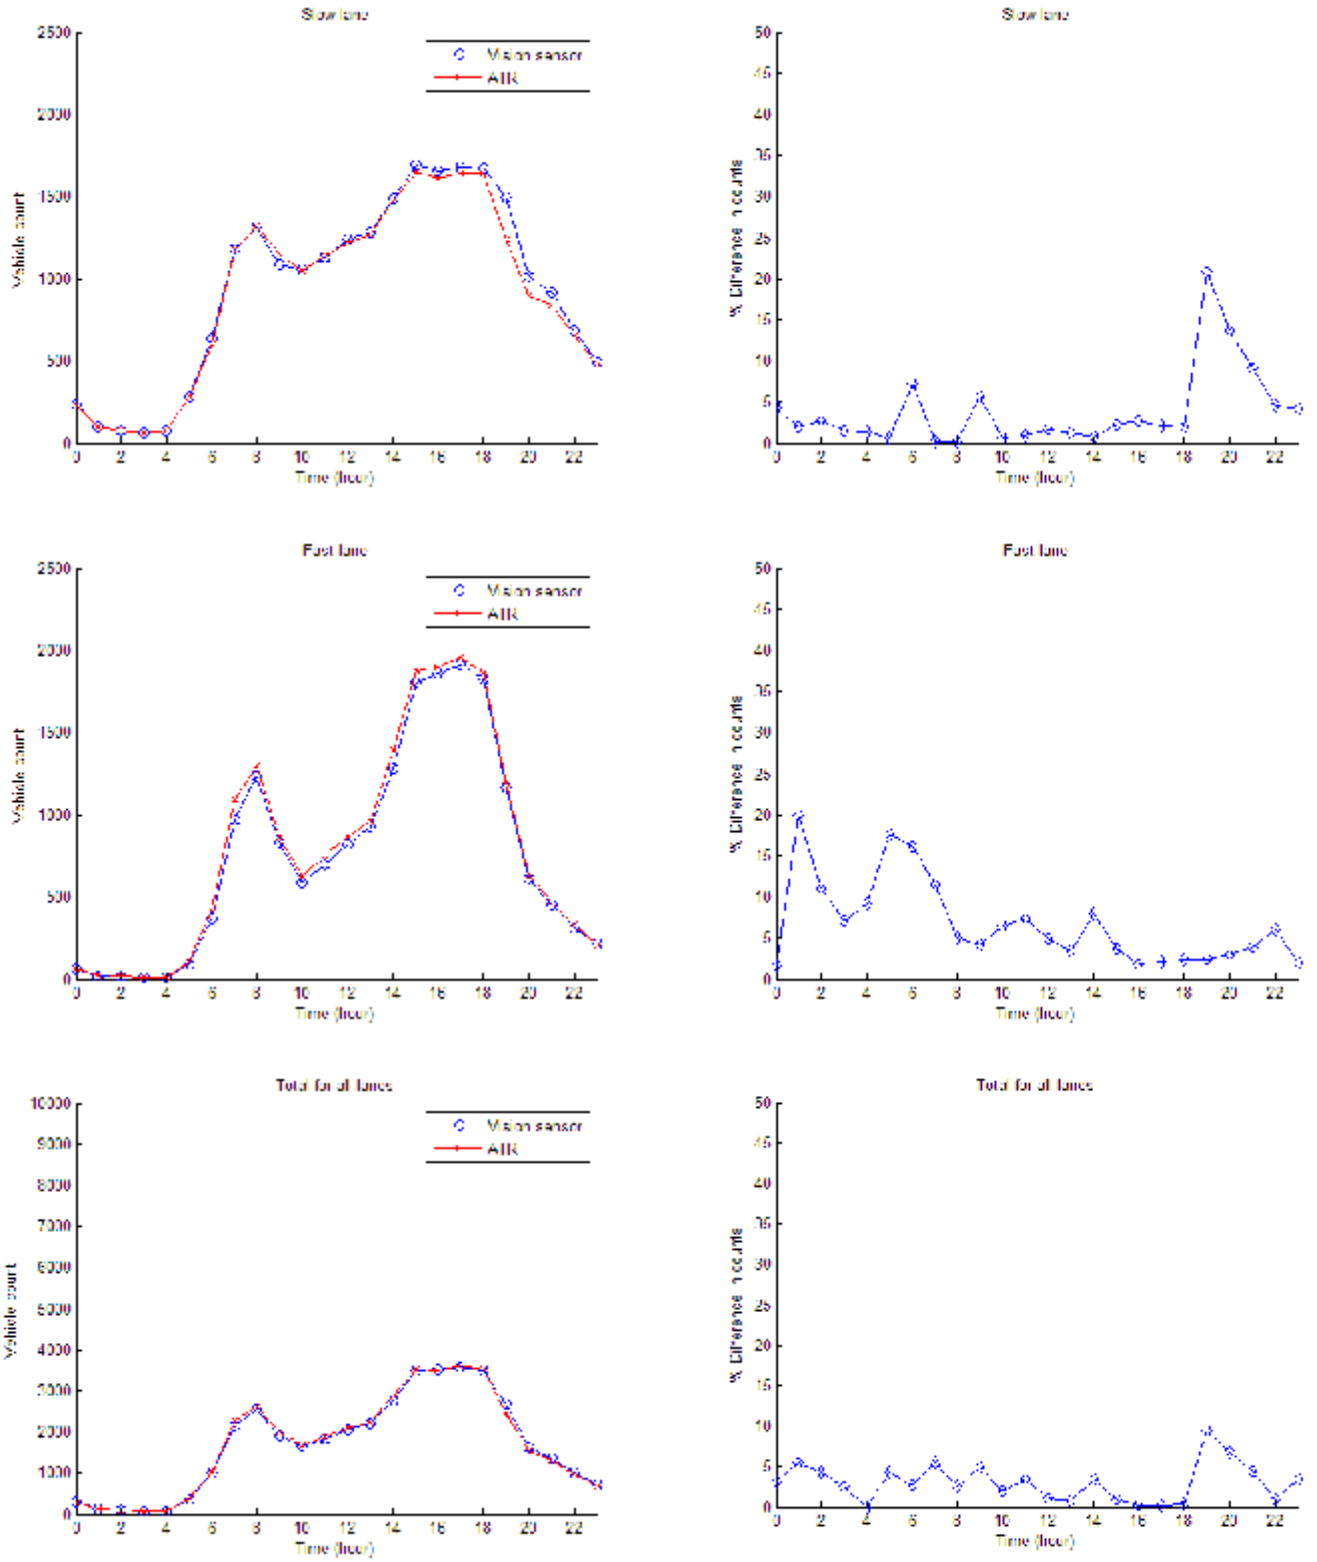


FIGURE 18: Comparison of vision-sensor counts and loop detector counts in Maryland shown for Friday, Sept 11, 2009. The horizontal axis shows the hour of day, from midnight to midnight.

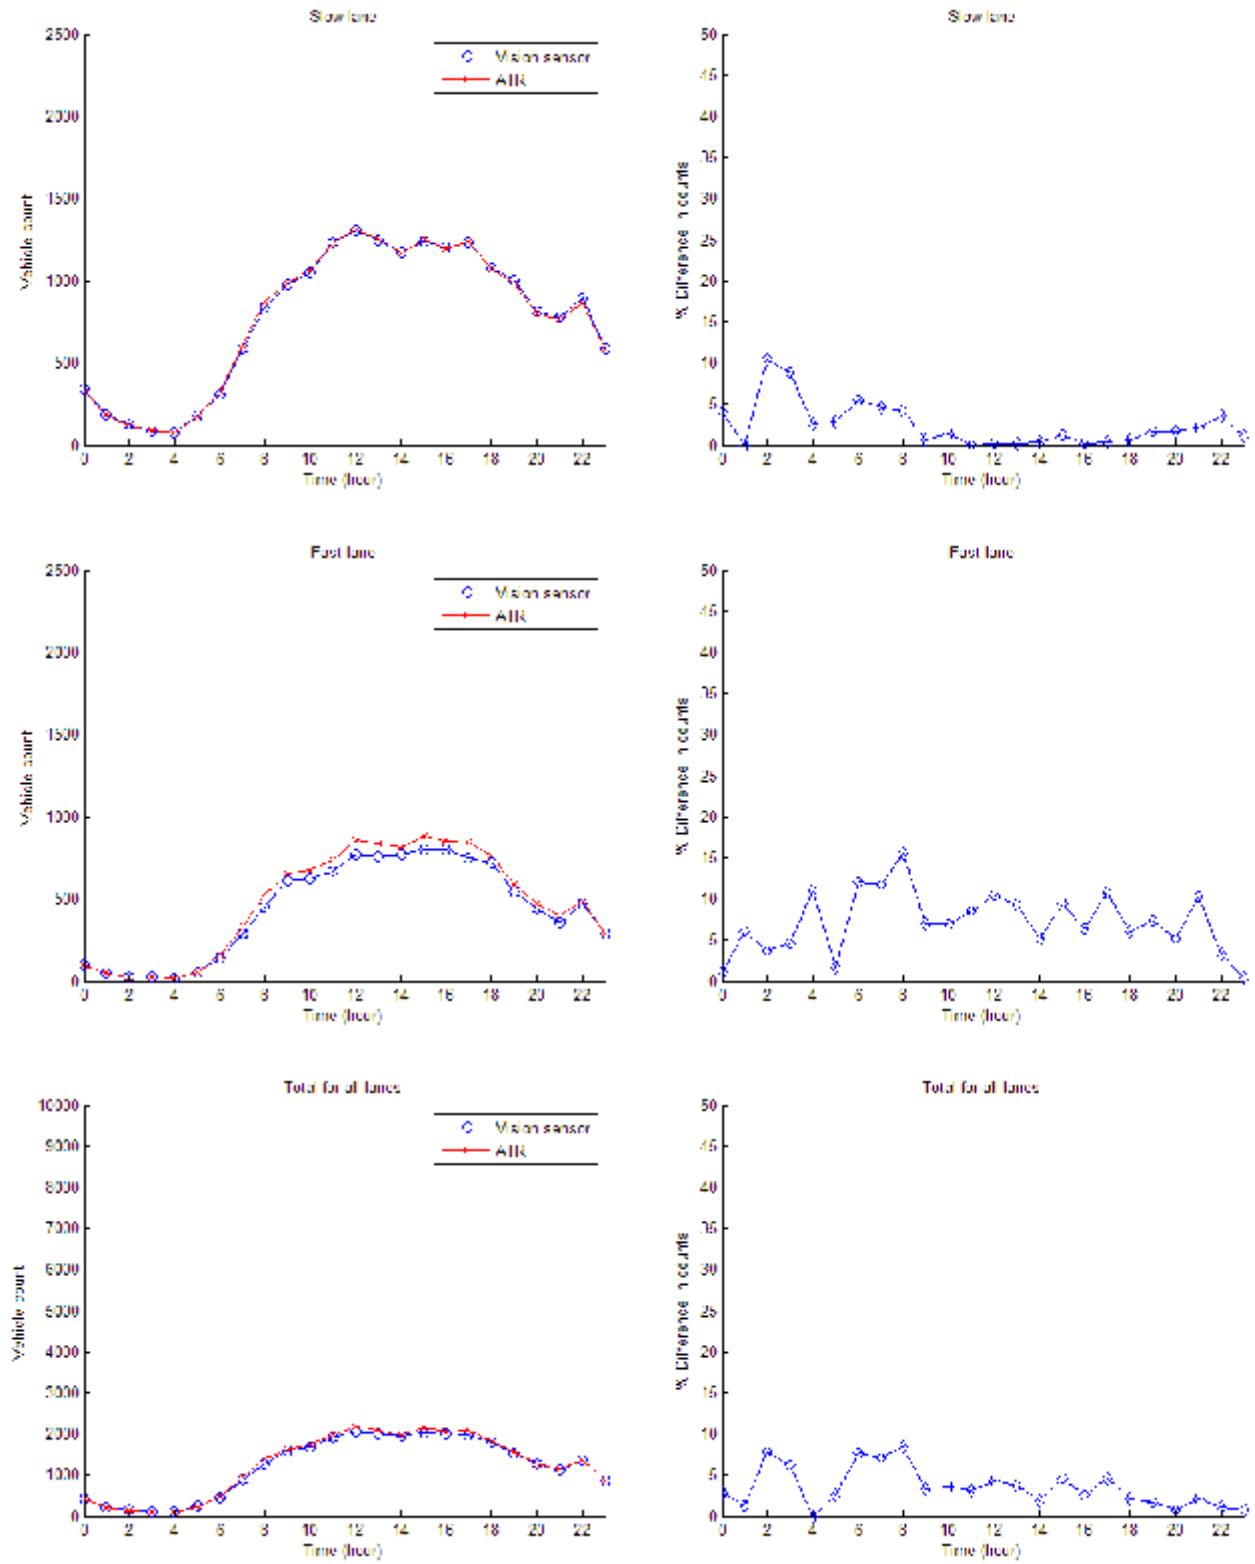


FIGURE 19: Comparison of vision-sensor counts and loop detector counts in Maryland shown for Saturday, Sept 12, 2009. The horizontal axis shows the hour of day, from midnight to midnight.

Despite the general agreement between the sensors, low volumes during the early morning hours result in large percentage differences. To analyze this phenomenon in detail, we manually annotated several hours of video data. Table 2 shows the result of this analysis. From the plots in Figure 18 we see that at 1:00 AM on Sept 11, the difference between the two sensor counts is almost 20% in the fast lane. However when compared to the ground truth counts it is revealed that the vision sensor actually performs better than the loop detector for this lane. Similarly at 8:00 AM on Sept 12 the difference in the fast lane is almost 15% whereas in reality the error is only 10%, just 4% worse than the loop detector.

Date				Manual counts (Ground truth)			Loop detector			Vision-sensor			% Error (Loop detector)			% Error (Vision-sensor)		
				Slow lane	Fast lane	Total	Slow lane	Fast lane	Total	Slow lane	Fast lane	Total	Slow lane	Fast lane	Total	Slow lane	Fast lane	Total
2009	9	11	1	104	22	126	102	25	127	100	20	120	-1.92	13.64	0.79	-3.85	-9.09	-4.76
2009	9	11	4	74	10	84	74	11	85	75	10	85	0.00	10.00	1.19	1.35	0.00	1.19
2009	9	11	6	618	391	1009	600	438	1038	643	367	1010	-2.91	12.02	2.87	4.05	-6.14	0.10
2009	9	11	7	1139	1027	2166	1176	1090	2266	1178	965	2143	3.25	6.13	4.62	3.42	-6.04	-1.06
2009	9	12	2	127	24	151	114	27	141	126	26	152	-10.24	12.50	-6.62	-0.79	8.33	0.66
2009	9	12	8	835	498	1333	870	528	1398	834	446	1280	4.19	6.02	4.88	-0.12	-10.44	-3.98

TABLE 2: Comparison of loop-detector and vision-sensor counts with ground-truth counts. Numbers in red indicate hours during which the loop detector error exceeds that of the vision-sensor. Negative signs indicate undercounting.

Figure 20 shows the accuracy of the two sensors compared with the ground truth counts for another 24-hour period, Thursday, February 19, 2009. From the per-lane count results (left plots in the figure), both the vision-sensor and the loop detector produce values that are almost indistinguishable from the ground truth. Interestingly, though, when the actual errors of the sensors are examined (right plots), it is discovered that the accuracy of the vision-based sensor is even slightly better than the loop detector for a majority of the hours. In fact, the vision system exceeded 5% error only one time in each lane during the 24 hour period, and the total 24-hour 2-lane volume is within 0.23% of ground truth. However, the per-lane errors of the loop detector cancel each other, yielding slightly better results for the total count, as seen in the 24-hour results summarized in Table 3.

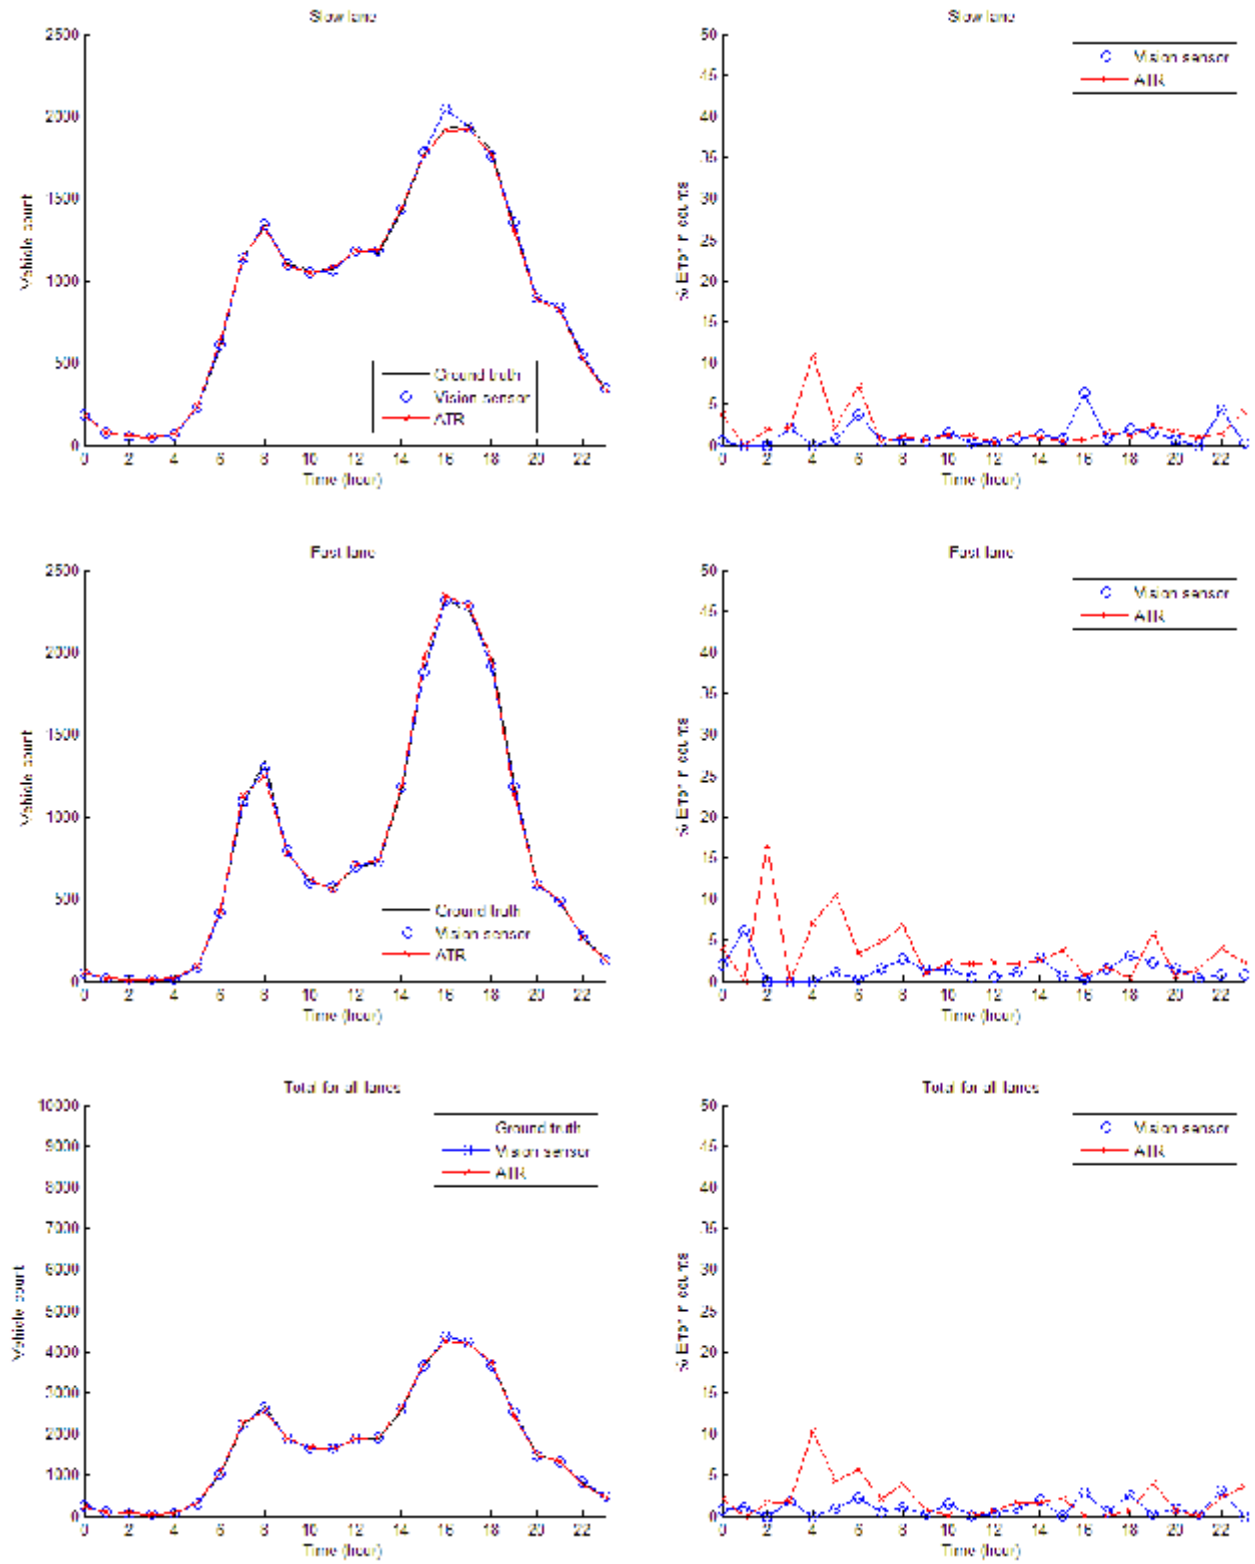








FIGURE 20: Accuracy of the vision-based sensor and the loop detector in Maryland by comparing with manually annotated ground truth, shown for Thursday, February 19, 2009.

	Manual counts (Ground truth)			Loop detector			Vision-system		
	Slow lane	Fast lane	Total	Slow lane	Fast lane	Total	Slow lane	Fast lane	Total
Counts	22109	18662	40771	22024	18750	40774	22256	18609	40865
% Error				-0.38	0.47	0.01	0.66	-0.28	0.23

TABLE 3: 24-hour comparison of loop-detector and vision-system counts with ground-truth counts in Maryland, Thursday, February 19, 2009.

Due to the dense traffic, frequent occlusions, and additional lane at the site in New York, coupled with the frequent loss of power, manual annotation of the video data at that location was much more time consuming. Nevertheless, shown below in Figure 21 is the lane-by-lane comparison of the vision-sensor and loop-detector counts, compared with ground truth, for a 24-hour video collected on Thursday, November 5, 2009. Overall, the vision sensor generates significantly less error than the loop detector, especially in the slow lane where the latter generated errors as high as nearly 40% in one hour of counting while the error from the former sensor never exceeded 10%. In the other lanes the vision sensor performed worse than the loop detector for a few hours just before midnight, but otherwise performed better. Notice that in all three lanes the maximum error of the loop detector is significantly higher than the maximum error of the vision sensor. These results were especially promising because no special refinements were made to the algorithm to account for the different camera or hardware setup. Table 4 shows the summary of the results for the entire 24-hour period for both sensors. Despite the fact that the loop detector errors are quite pronounced in the per-hour plots, the errors tend to cancel each other over time, leading to 24-hour counts with error just twice that of the vision-based sensor.



(FIGURE 21 continued)




FIGURE 21: Results for the loop detector and vision-based sensor in New York for a 24-hour period (midnight to midnight) on Thursday, Nov. 5, 2009, compared with ground truth.

	Manual counts (Ground truth)				Loop detector				Vision-system			
	Slow lane	Middle lane	Fast lane	Total	Slow lane	Middle lane	Fast lane	Total	Slow lane	Middle lane	Fast lane	Total
Counts	31883	31491	31377	94751	29418	29983	30617	90018	31432	31592	29400	92424
% Error					-7.7	-4.8	-2.4	-5.0	-1.4	0.3	-6.3	-2.5

TABLE 4: 24-hour comparison of loop-detector and vision-system counts with ground-truth counts in New York, Thursday, Nov. 5, 2009.

Figure 22 shows the results of the vision-based sensor at the location along I-85 in South Carolina for a 24-hour period on Wednesday, February 24, 2010, compared with ground truth. No ATR exists at this location with which to compare the vision sensor. Per-hour counts are slightly better than those in New York but slightly worse than those in Maryland, due to the large distance from the road edge coupled with the low height of the camera. Nevertheless, total per-hour count errors are less than 5% for all hours, while per-lane per-hour counts are generally less than 10%. Table 5 shows that the per-lane errors cancel over the entire day, leading to a total error in count over the 24-hour period of just two vehicles out of more than 15,000, or an error rate of a mere -0.01%.




FIGURE 22: Results for vision-based sensor along I-85 in South Carolina for a 24-hour period (midnight to midnight) on Wednesday, February 24, 2010, compared with ground truth.

	Manual counts (Ground truth)			Vision-system		
	Slow lane	Fast lane	Total	Slow lane	Fast lane	Total
Counts	10138	5312	15450	10493	4955	15448
% Error				3.5	-6.7	-0.0

TABLE 5: 24-hour comparison of vision-system counts with ground-truth along I-85 in South Carolina, Wednesday, February 24, 2010.

Speed and classification accuracy

While the Maryland site did not include a sensor to collect either speed or classification data, the piezoelectric sensor in New York was able to capture both. Figure 23 shows the difference in speeds – aggregated hourly – between the piezo sensor and the vision-based system for the 24-hour period of November 5, 2009. The hourly average speed difference between the two sensors never exceeded 4.5 miles per hour (mph), and the average difference weighted by traffic volume over the 24-hour period is 2.1 mph. Thus, the two sensors are in fairly close agreement with each other.




FIGURE 23: Hourly difference in speeds between the piezoelectric sensor and the Clemson system in New York on Thursday, Nov. 5, 2009.

As mentioned earlier in the report, the vision-based system is currently unable to classify vehicles when there is not sufficient ambient light to estimate vehicle length. Therefore, the classification comparison of the two sensors was performed only on the daylight data between the hours of 8:00 AM to 4:00 PM on November 5, 2009. Because it is an axle counter, the

piezoelectric sensor is capable of classifying vehicles into 13 categories. For this comparison, the 13 FHWA classes were aggregated into a three-class scheme as follows:

Class A: FHWA classes 1,2,3

Class B: FHWA classes 4-7

Class C: FHWA classes 8-13.

Note that the vision-based system was shown to accurately detect motorcycles in (14). However, the video data included in this report contain almost no motorcycles, and it has not yet been determined how to set the correct sensitivity of the motorcycle classifier so as not to introduce false positives. Therefore, motorcycles were included in Class A along with the other small vehicles.

Table 6 shows the results of the vehicle classification comparison between the piezo sensor and the vision-based system. The results for the two devices are almost identical in terms of the proportions in each class. The differences aggregated over the 8-hour period are usually less than 1% for each travel lane. Table 7 provides a summary of the number of unclassified vehicles for each sensor during the 8-hour period, showing that the piezo sensor produces far fewer unclassified vehicles than the vision-based system.

Class	Piezoelectric Sensor				Clemson System			
	Slow	Middle	Fast	Total	Slow	Middle	Fast	Total
A	92.07	84.68	99.68	92.44	92.11	83.31	99.29	91.51
B	5.63	8.76	0.17	4.71	5.22	10.37	0.24	5.33
C	2.30	6.55	0.02	2.85	2.67	6.32	0.47	3.16

TABLE 6: Summary of classified vehicles in NY from 8:00 AM to 4:00 PM on November 5, 2009.

Lane	Piezoelectric Sensor			Clemson System		
	Unclassified	Total	% Unclassified	Unclassified	Total	% Unclassified
Slow	353	14030	2.5%	1163	14928	7.8%
Medium	403	13522	3.0%	1695	13755	12.3%
Fast	150	14459	1.0%	905	12562	7.2%
Total	906	42011	2.2%	3763	41245	9.1%

TABLE 7: Summary of unclassified vehicles in NY from 8:00 AM to 4:00 PM on November 5, 2009.

Additional results in a variety of challenging scenarios

Manually annotating video data for speed and classification results is much more time-consuming than doing so for counts. As a result, it is prohibitive to obtain manual ground truth for such results over many consecutive hours as was done for counts. Instead, nine video clips – each several minutes long – were selected from among the four camera locations mentioned previously and manually annotated for counts, speeds, and classification. These clips were chosen to represent a variety of camera locations, lighting conditions, and weather conditions. Sample images from the video clips are displayed in Figure 24. The number appended to the name of a video clip shows the hour in which the video was collected, so that “-09” means the video was extracted in the morning, between the hours of 9:00am and 10:00am; “-21” means between 9:00pm and 10:00pm, and so forth. The Greenville TMC clip was collected in the afternoon, while the MD-icicle and MD-snow clips were captured during the middle of the day. Note that it was actually snowing during the latter video clip.

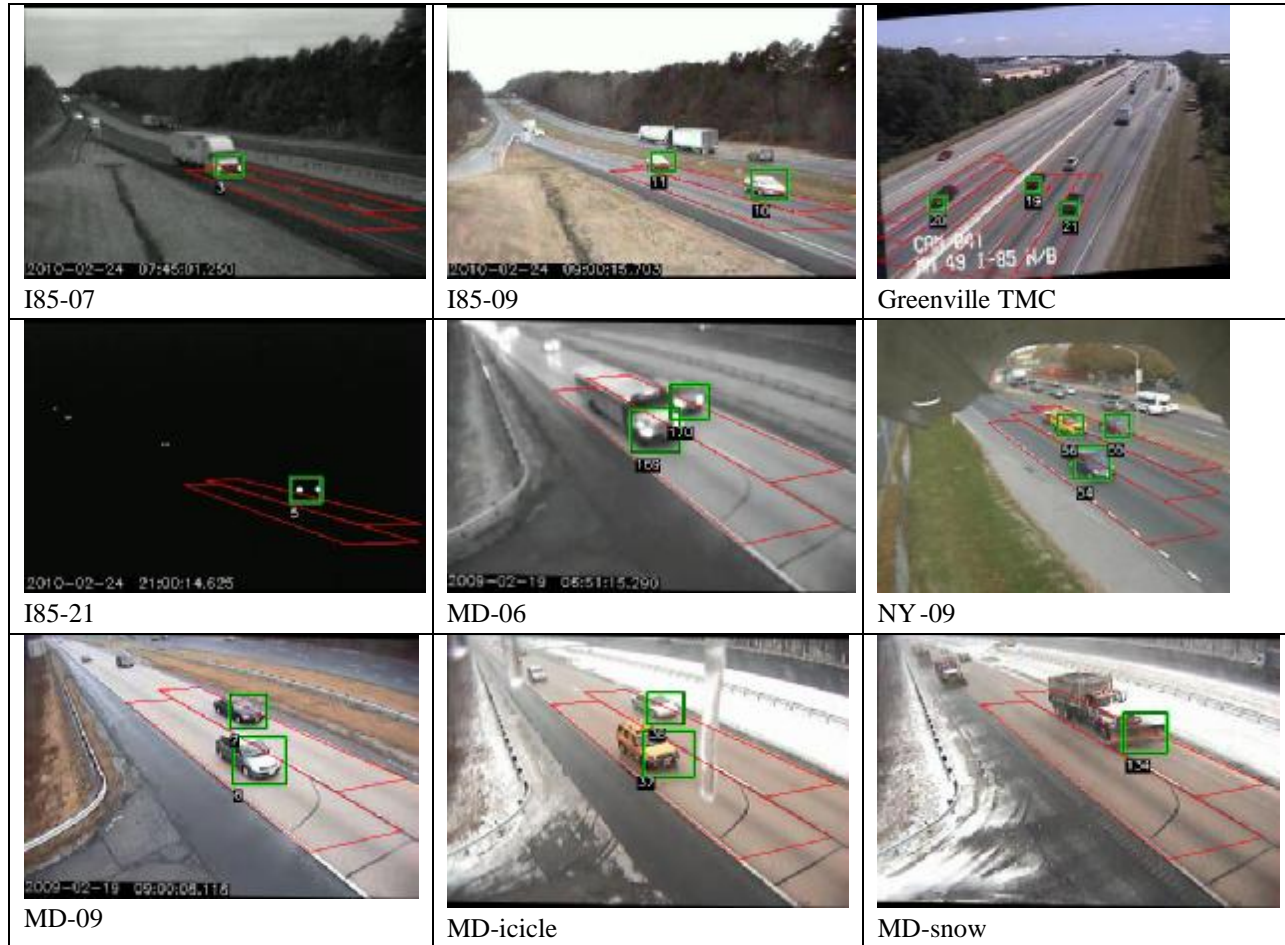


FIGURE 24: Images of the nine video sequences used for evaluating speeds and classification results.

Table 8 shows a summary of the results of the video-based system for these nine different video clips, compared with ground truth. For measuring speeds, the travel lengths were as

follows: 80 feet for MD, 100 feet for NY, 160 feet for I-85, and 120 feet for TMC. The detection error ranged from 0.5% to 6.6%, while the classification error (using the three classes described previously) ranged from 0.8% to 2.6%. The only videos in which vehicles were unclassified were those captured along I-85, due to the vantage point of the camera, exacerbated by the low contrast during the 7:00 AM hour. The speed error was computed on a per-vehicle basis, so that the maximum speed error among all vehicles in all videos remained under 10 mph. The average speed, which provides a more realistic measure of the accuracy of the system by averaging over all vehicles in the video clip, never differed from the ground truth by more than 1.4 mph.

	Video duration (min:sec)	# vehicles	# detected	Detect. error	Class. error	# unclass.	Max speed error (mph)	Avg. speed error (mph)
TMC	02:53	228	237	3.9%	2.2%	0	3.6	0.6
I85-07	15:04	235	231	1.7%	1.7%	10	7.7	0.6
I85-09	09:59	155	148	4.5%	2.6%	4	3.6	0.2
I85-21	15:03	150	140	6.6%	N/A	N/A	3.0	0.1
NY-09	07:29	378	393	4.0%	1.6%	0	6.6	0.3
MD-06	22:20	370	381	3.0%	1.9%	0	9.0	0.1
MD-09	14:59	356	350	1.7%	0.8%	0	6.0	1.4
MD-snow	15:03	199	198	0.5%	1.5%	0	3.0	0.9
MD-icicle	15:00	169	172	1.8%	1.8%	0	5.0	0.1

TABLE 8: Results of the video-based sensor computing counts, speeds, and classification on a variety of video clips.

By way of comparison, it is worth noting that a simple study evaluating the popular Autoscope image detection system concluded that it produced a classification average error of 4.4% versus manual ground truth for two classes (18). Another evaluation found that accuracy using 5 vehicle classes varied from 65% to 90% depending on the type of facility and the number of lanes, with the intensity of shadows or the presence of rain degrading the average accuracy by an additional 10%. Thus, these results above seem to indicate that the proposed vision-based system may be able to perform competitively.

PLANS FOR IMPLEMENTATION

The primary customer base of the technology is state and local transportation agencies. In this project, the research team partnered with three state departments of transportation (Maryland, New York, and South Carolina). Working prototypes of the sensor were installed in Maryland and New York. This allowed the technology to be further refined to enhance its robustness and applicability and to demonstrate its capabilities to potential early adopting customers. Moreover,

the team has spent considerable effort establishing contacts with various potential customers through presentations at regional and national conferences such as TRB Annual Meeting, SDITE, and NATMEC.

The team has partnered with a local software company to bring this research idea to market. Significant progress has been made in developing a business strategy, producing the necessary ancillary software such as the user interface, and identifying potential customers. A booth at the upcoming TRB conference in January 2011 has been secured, at which the product will be revealed to the public under the name of TrafficVision.

Moreover, the team has an ongoing relationship with the Clemson University Research Foundation (CURF), whose purpose is to commercialize the intellectual property developed at the university through technology transfer. To protect the intellectual property developed in this project, a patent application covering the technology was filed in December 2009.

CONCLUSION

The project involved development and testing of a next generation vision-based traffic sensor to collect traffic parameters such as volume, speed, and classification using both fixed and pan-tilt-zoom cameras. The prototype sensor was deployed at two locations in Maryland and New York, and additional data was collected at both a temporary setup in South Carolina as well as at a TMC in South Carolina using existing cameras. The accuracy of the sensor in terms of vehicle counts, speeds, and classification (three length-based categories) was found to be comparable to those of the loop detectors and piezos present at the corresponding sites. Manually annotated ground truth was able to further validate the system, showing that under a variety of camera locations, lighting changes, and environmental conditions, the system was able to produce per-lane per-hour errors less than approximately 5% in general, and nearly always less than 10%.

Following are some notable conclusions from this project:

1. The proposed sensor is versatile in nature. It processes video gathered from analog or digital (network IP) cameras which can be either fixed or pan-tilt-zoom. The sensor produces vehicle counts, speeds and classification as its output. Due to the ability of the sensor to track vehicles it is possible to estimate more microscopic parameters such as space-mean-speed, lane changes, acceleration profile, etc., with additional software development.
2. More than 72 hours of data was manually annotated to provide more accurate validation of the sensor. These data will also enable regression tests on future versions of the software.

3. The ability of the sensor to handle perspective effects at low camera heights opens up the possibility of using the sensor in portable short-term data collection as well as work-zone safety applications.
4. The sensor has demonstrated its ability to work reliably in varying lighting and weather conditions.
5. Considerable progress has been made toward the goal of complete automatic calibration of the sensor. In the present implementation a sensor calibrated for one view of the camera is able to automatically recalibrate itself when the camera undergoes pan-tilt motion. Extension of this functionality to handle variable zoom is expected in the near future.
6. Thousands of cameras across the country are currently used only for manual surveillance and incident verification. The ability of the proposed vision-based sensor to automatically recalibrate itself will facilitate turning such cameras into intelligent sensors.
7. Significant progress has been made toward developing a solar-powered system. By reducing the power consumption of the device considerably, it is expected that the present version of the system will operate in any location with sufficient solar panels that are not covered by shade.

REFERENCES

1. J. Mergel. An Overview of Traffic Monitoring Programs in Large Urban Areas, Center for Transportation Information of Volpe National Transportation Systems Center, Cambridge, MA, July 1997.
2. Office of Highway Policy Information. *Traffic Monitoring Guide*, FHWA, USDOT, May 2001.
3. Office of Highway Policy Information. *Traffic Monitoring Guide Supplement*, Section 4S FHWA, USDOT, April 2008.
4. Dan Middleton, Ryan Longmire, and Shawn Turner. *State of the Art Evaluation of Traffic Detection and Monitoring Systems Volume I – Phases A & B: Design*. Publication FHWA-AZ-07-627(1). FHWA, USDOT, February 2007.
5. P. T. Martin, Y. Feng, and X. Wang. *Detector Technology Evaluation*. (November, 2003. <http://www.mountain-plains.org/pubs/html/mpc-03-154/index.php> Accessed July 27, 2009.
6. Jerry Kotzenmacher, Erik Minge, and Bingwen Hao. *Evaluation of Portable Non-Intrusive Traffic Detection System*. Final Report: MN-RC-2005-37, Minnesota Department of Transportation, September 2005.
7. French Engineering, LLC. Traffic Data Collection Methodologies, Final report for Pennsylvania Department of Transportation, Bureau of Planning and Research, April 2006.
8. James H. Banks. Evaluation of Portable Automated Data Collection Technologies: Final Report, California PATH Research Report, UCB-ITS-PRR-2008-15, August 2008.
9. S. L. Skaszek. State-of-the-Art Report on Non-Traditional Traffic Counting Methods, FHWA-AZ-01-503, October 2001.
10. Neeraj K. Kanhere, Shrinivas J. Pundlik, and Stanley T. Birchfield. Vehicle segmentation and tracking from a low-angle off-axis camera. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1152–1157, June 2005.
11. N. K. Kanhere, S. T. Birchfield, and W. A. Sarasua. Vehicle Segmentation and Tracking in the Presence of Occlusions, *Transportation Research Record: Journal of the Transportation Research Board*, No. 1944, pp. 89-97, 2006.
12. N. Kanhere, S. Birchfield, W. Sarasua, and T. Whitney. Real-Time Detection and Tracking of Vehicle Base Fronts for Measuring Traffic Counts and Speeds on Highways, *Transportation Research Record: Journal of the Transportation Research Board* No. 1993, pp. 155-164, 2007.
13. N. Kanhere and S. Birchfield. Real-Time Incremental Segmentation and Tracking of Vehicles at Low Camera Angles Using Stable Features. *IEEE Transactions on Intelligent Transportation Systems*, 9(1):148-160, March 2008.
14. N. Kanhere, S. Birchfield, W. Sarasua, and S. Khoeini. Traffic Monitoring of Motorcycles During Special Events Using Video Detection, *Transportation Research Record: Journal of the Transportation Research Board*, 2010 (in press).

15. N. Kanhere, S. Birchfield, and W. Sarasua. Automatic Camera Calibration Using Pattern Detection for Vision-Based Speed Sensing, *Transportation Research Record: Journal of the Transportation Research Board*, No. 2086, pp. 30-39, 2008.
16. N. Kanhere and S. T. Birchfield. A Taxonomy and Analysis of Camera Calibration Methods for Traffic Monitoring Applications. *IEEE Transactions on Intelligent Transportation Systems*, 11(2): 441-452, June 2010.
17. *Highway Capacity Manual*. Transportation Research Board, Washington, D.C.. 2000.
18. B. Auffray, K. A. Tufte, Z. Horowitz, S. Matthews, and R. Bertini. Evaluation of Single-Loop Detector Vehicle-Classification Algorithms using an Archived Data User Service System, presented at the ITE District 6 Annual Meeting, Honolulu, June 2006.
19. X. Yu, P. Prevedouros, and G. Sulijoadikusumo. Evaluation of Autoscope, SmartSensor HD, and TIRTL for Vehicle Classification Detectors. Presented at 89th Annual Meeting of the Transportation Board, paper #10-1846, Transportation Research Board, Washington, D.C., 2004.