

Toward a Comparative Study of Lane Tracking using Omni-directional and Rectilinear Images for Driver Assistance Systems

Shinko Y. Cheng and Mohan M. Trivedi

Abstract—With a panoramic view of the scene, a single omni-directional camera can monitor the interior and exterior of the vehicle simultaneously. Omni-directional cameras have been shown to produce useful data for determining head gaze orientation from within a vehicle. We examine the issues involved in integrating lane-tracking functions using the same omni-directional camera. We present analysis on the impact of the omni-directional camera’s reduced image resolution on lane tracking accuracy, as a consequence of gaining the expansive view. This was done by first presenting Omni-VioLET, a modified implementation of the vision-based lane estimation and tracking system (VioLET), and evaluating the performance of both lane-trackers by comparing their performance with ground-truth generated with images synchronously captured with the other images along the same freeway course. The results are surprising: With $1/10^{th}$ the number of pixels representing the same space and about $1/3^{rd}$ the horizontal image resolution as a rectilinear image of the same road, the omni-directional camera implementation results in only twice the amount of mean absolute error in tracking the left-lane boundary position. With 2.3 times the resolution of the rectilinear camera, an omni-directional camera is shown to be just as accurate (in terms of average spatial error) as the rectilinear version.

I. INTRODUCTION: OMNIDIRECTIONAL CAMERA FOR LOOKING IN AND LOOKING OUT

The main feature of omni-directional cameras is its ability to capture an image of the scene 360 degrees around the camera, and its potential use in monitoring many things in the scene at one time as illustrated in Fig. 1. In an intelligent driver assistance system, the panoramic view can be used to monitor the front, rear, and side views of the vehicle’s surroundings with a single sensor. The same camera can be positioned to capture the scene outside and inside the vehicle, eliminating the need for multiple cameras and any complex calibration maintenance algorithms between multiple cameras. Reducing redundancy being one of the goals in embedded systems, combining multiple functionalities into a single, simpler sensor reduces the cost associated with using individual sensors for each function. There is also increasing amounts of evidence that driver behavior should be an integral part of any effective driver assistance system [13], motivating the need for a suite of sensors that extracts cues from both outside as well as inside the vehicle. With these motivations, we investigate problems associated with integrating driver assistance functionalities, originally designed for multiple rectilinear cameras, on a single omni-directional camera. More specifically, we examine the issues involved in integrating lane-tracking functions using the omni-directional cameras in this multi-function context.

S. Y. Cheng and M. M. Trivedi are with the Laboratory for Intelligent and Safe Automobiles at the University of California, San Diego, 9500 Gilman Drive MC0434, La Jolla CA 92037-0434 sycheng@ucsd.edu, mtrivedi@ucsd.edu

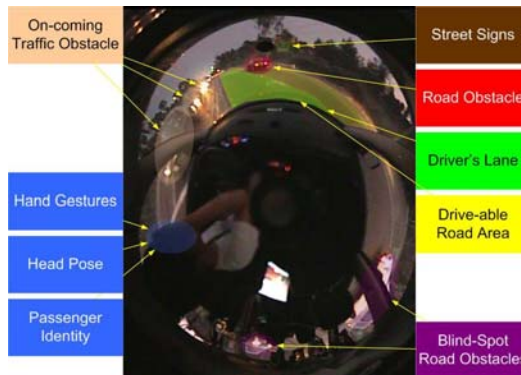


Fig. 1. Panoramic image captured by omni-directional cameras with a potential for holistic visual context analysis.

Huang, *et al* [8] demonstrated that an omni-directional camera can be used to estimate driver head pose to generate the driver’s view of the road. Knowledge of the driver’s gaze direction has many uses beyond driver-view synthesis. The driver head motion is the critical component that adds one second of warning time to a lane departure warning system in [11]. This human-centered driver support system use vehicle lane position from cameras looking out of the vehicle, vehicle speed, steering, yaw rate from the vehicle itself, as well as head motion from a camera looking in the vehicle to make predictions of when drivers make lane-change maneuvers. Estimates of driver head movement also improved intersection turn maneuver predictions [4]. Subsequently, each of these predictions can potentially describe the driver’s awareness of the driving situation. For example, given an obstacle in the vehicle’s path and continued driver preparatory movements to perform the maneuver, the assistance system can then conclude that the driver is unaware of the danger and take appropriate corrective action. Observing the driver also has applications in driver identity verification, vigilance monitoring, along with intention and awareness monitoring. It is clear that driver bodily movements are very significant cues in determining several driver behaviors, and visually extracting driver information using omni-directional cameras have been shown to be an effective approach.

Lane tracking is an important component in many intelligent driver assistance systems. Lane tracking has utility in generating warnings for unintentional lane departure, over-speed warnings for oncoming sharp turns, estimating ahead vehicle proximity for adaptive cruise control, avoiding obstacles, and many others [5]. Just as observing drivers will enhance driving safety, lane tracking

is an integral part in the same task.

This naturally leads to the following question: Can efficiency be improved and accuracy maintained in utilizing a single omnidirectional camera rather than two rectilinear cameras to perform these same functions of observing the driver and the road? Since the two systems operate on essentially a transformed image of the same scene, the answer should be yes, but to what extent? Omnidirectional cameras typically use the same kind of sensing electronics as rectilinear cameras, providing the same image resolution. This implies that omnidirectional cameras necessarily provide less detail of the scene at the cost of the 360 degree coverage. How accurate is lane-tracking with multi-purpose omnidirectional cameras? We attempt to answer these questions by describing results attained with a modified implementation of and conduct a comparative performance evaluation of VioLET, a vision-based lane estimation and tracking system [10] modified to operate from an omnidirectional camera as well as a monocular rectilinear camera. We then present analysis on the relationship between image resolution and spatial error in lane-tracking.

Our contributions can be listed in the following:

- 1) We introduce Omni-VioLET, a lane-tracking system using an omnidirectional camera based on VioLET, a well-tested, robust lane tracking algorithm. The same omnidirectional camera also captures a view of the driver for driver monitoring applications.
- 2) We present a system-level performance comparison of the lane-tracking system using a rectilinear camera and omnidirectional camera on the same road course with ground-truth.
- 3) We present a comparative analysis of lane tracking accuracy from using these two types of imaging sensors by describing and utilizing a simulation model characterizing the relationship between lane detection accuracy and image resolution.

II. RELATED RESEARCH IN VISION-BASED LANE TRACKING

Most previously proposed vision-based lane tracking systems follow a processing structure consisting of these three steps: 1) extract road features from sensors, 2) suppress outliers from the extracted road features, and 3) estimate and track lane model parameters.

There are several notable lane tracking approaches using rectilinear cameras. One by Bertozzi *et al.* [1] proposes to use stereo rectilinear camera for lane detection combined with obstacle detection. Their approach employs a flat-plane transformation of the image onto a birds-eye view of the road, followed by a series of morphological filters to locate the lane markings. A recent contribution by Nedevschi *et al.* [12] augments the usual flat-plane assumption of the road to a 3D model based on clothoids and vehicle roll angles. That system relies on depth maps calculated from a stereo camera and edges in images to infer the orientation of the 3D road model. A detailed survey of lane position and tracking techniques using monocular rectilinear cameras is presented in [10].

Ishikawa, *et al.* [9] proposes an approach using an omnidirectional camera to track lane position in an autonomous vehicle application. The approach first transforms the omnidirectional image to flat-plane, followed by hough transform to search for the left and right

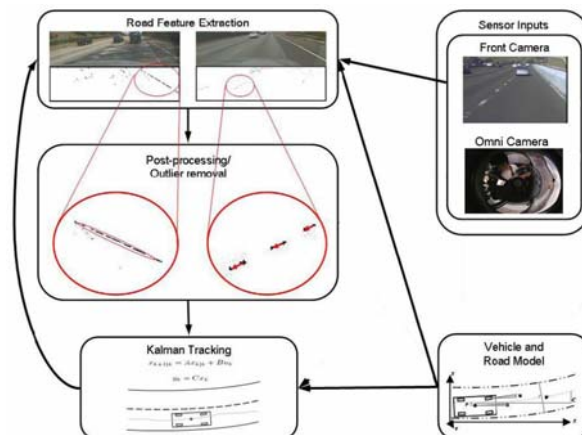


Fig. 2. System flow diagram for VioLET, a driver assistance focused lane position estimation and tracking system.

lane marking with a lane separation prior probability model. With an autonomous vehicle application in mind, the scene captured by this omnidirectional camera saw lane markers ahead and behind the vehicle, both aiding in determining the vehicle's position in the lane and lane width. Furthermore, lines perpendicular to the vehicle could also be detected with this system since the sides are also monitored as well. This work demonstrates that effective lane tracking can be achieved to an extent with omnidirectional images.

Because the Ishikawa approach was designed in the context of autonomous vehicles, the operating environment, although the setting was outdoors, was idealized with solid white lines for lane markings with constant lane widths. The central component of the VioLET system is the use of steerable filters, which has been shown to be highly effective in extracting circular reflectors as well as line segments, the prevalent kinds of lane markings in actual California freeways. Furthermore, the lane tracker in Ishikawa:omnilane lacked a mechanism to incorporate temporal history of last observed lane locations with the current estimate. We will use the Kalman filtering framework to update interesting statistics of lane position, lane width, and so forth using lane marking position measurements from the current as well as previous moments in time. Lastly, we are interested in examining the extent to which lane tracking can be accurate by using only the view of the road ahead of the vehicle, deferring the other areas of the omnidirectional image for other applications that monitor the vehicle interior.

III. OMNI VIOLET LANE TRACKER

In this section, we describe the modifications of the VioLET system for operation on omnidirectional images. An omnidirectional camera is positioned just below and behind the rear view mirror. From this vantage point, both the road ahead of the vehicle as well as the front passengers can be clearly seen. With this image, left lane marker position, right lane marker position, vehicle position within the lane, and lane width are estimated from the image of the road ahead. Fig. 2 shows a block diagram of the Omni-VioLET system for use in this camera comparison.




Fig. 3. This illustration shows the original images, the flat-plane transformed images, filtered response image from the flat-plane transformed images for circular reflectors and lane line markers.

The VioLET system operates on a flat-plane transformed image, which is a birds-eye view of the road. This is generated from the original image captured from the camera, with knowledge of the camera’s intrinsic and extrinsic parameters. The intrinsic and extrinsic parameters of the camera describe the many-to-one relationship between the 3D points in the scene in real-world length units and its projected 2D location on the image in image pixel coordinates. The planar road assumption allows us to construct a one-to-one relationship between 3D points on the road surface and their projected 2D locations on the image. This is one of the critical assumptions that allow lane-tracking algorithms to provide usable estimates of lane location and vehicle position, and is also the assumption utilized in the VioLET system.

The model and calibration of rectilinear cameras are very well studied, and many of the results translate to omni-directional cameras. We can draw direct analogs between the omni-directional camera model and the rectilinear camera model, namely its intrinsic and extrinsic parameters. Tools for estimating the model parameters have been also recently made available [14]. Table I summarizes the transformations from a 3D point in the scene $\mathbf{P} = (x, y, z)$ to the projected image point on the image $\mathbf{u} = (u, v)$.

Utilizing the camera parameters for either rectilinear and omni-directional cameras, a flat-plane image can be generated given knowledge of the world coordinate origin and the region on the road surface we wish to project the image onto. The world origin is set at the center of the vehicle on the road surface with the y-axis pointing forward and z-axis pointing upward. Examples of the flat plane transformation are shown in Fig. 3. Pixel locations of points in the flat-plane image and the actual locations on the road are related by a scale factor and offset.

The next step is extracting road features by applying steerable filters based on the second derivatives of a two-dimensional Gaussian density function on the transformed images. Two types of road features are extracted: circular reflectors (“Bots dot”) and lines. The circular reflectors are not directional so the filter responses are equally high in both the horizontal and vertical directions. The lines are highly directional and yield high responses for filters oriented along its length. The filtered images such as the one shown in Fig. 3 are then thresholded and undergo connected component analysis to isolate the most likely candidate road features.

The locations of the road features are averaged to find the new measurement of the lane boundary location. The average is weighted on its proximity to the last tracked location of the lane

boundary. The measurement for the other lane boundary is made the same way. The last estimated lane boundary location is estimated using a Kalman filter using lane boundary locations as observations, and vehicle position in the lane, left lane boundary location, right lane boundary location and lane width as hidden states. For more details on these steps of extracting road features and tracking these features using the Kalman filter, we refer the reader to the original paper [10]. The original implementation also takes advantage of vehicle speed, yaw-rate and road texture to estimate road curvature and refine the estimates of the lane model. We chose to omit those measurements in our implementation, and to focus on estimating lane boundary position and vehicle position in the lane, for which ground-truth can be collected, to illustrate the point that omni-directional images of a certain resolution can achieve a certain level of lane tracking accuracy.

Altogether, the VioLET system assumes a planar road surface model, assumes having knowledge of camera parameters, uses steerable filters to extract road features, and uses a Kalman filter to estimate lane model parameters from road feature location measurements. The outputs are lane boundary positions, lane width and vehicle position in the lane.

IV. EXPERIMENTAL PERFORMANCE EVALUATION AND COMPARISON

Omni-VioLET lane tracking system is evaluated with video data collected from three cameras. All were installed in a specially equipped test vehicle named the Laboratory for Intelligent and Safe Automobiles - Passat (LISA-P). Both the rectilinear camera, placed on the roof of the car, and the omni-directional camera, hung over the rear-view mirror, capture the road ahead of the vehicle. A third camera, mounted over the left side, is used to collect ground-truth. The vehicle was driven along actual freeways and video was collected synchronously along with various vehicle parameters. Details of the test-bed can be found in [4]

For evaluation, we collected groundtruth lane position data using video from the third calibrated camera. A flat-plane transformation of images from this camera is also generated such that the horizontal position of the transformed image represents the distance from the vehicle. A grid of points corresponding to known physical locations of the ground is used to adjust the orientation and position of the side camera. Fig. 4 shows the result of manually correcting the pose of the camera, and thus the grid of points in the image from the side camera. With this grid of points and its associated location in the image, a flat-plane image is generated as shown in the same figure. From the flat-plane image, lane positions are manually annotated to generate groundtruth. This groundtruth is compared against lane tracking results of both the rectilinear and omnidirectional camera based VioLET systems.

The test vehicle was driven at dusk on a multi-lane freeway at freeway speeds for several minutes. The image resolution of the flat-plane transformed image derived from the omni-camera was set at 100×100 , while the one derived from the rectilinear-image was set at 400×400 . These numbers were chosen because the lateral resolution of the road is approximately 100 pixels for the omnidirectional image and 400 pixels for the rectilinear image. The

TABLE I
 PROJECTIVE TRANSFORMATION FOR RECTILINEAR AND OMNI-DIRECTIONAL CAMERAS

	Rectilinear Camera Model	Omni-directional Catadioptric Camera Model
World to Camera coordinates	$\mathbf{P}_c = \begin{pmatrix} X_c \\ Y_c \\ Z_c \end{pmatrix} = \mathbf{R}\mathbf{P} + \mathbf{t}$	$\mathbf{P}_c = \begin{pmatrix} X_c \\ Y_c \\ Z_c \end{pmatrix} = \mathbf{R}\mathbf{P} + \mathbf{t}$
Camera to Homogeneous Camera Plane (Normalized Camera) Coordinates	$\mathbf{p}_n = \begin{pmatrix} x_n \\ y_n \end{pmatrix} = \begin{pmatrix} X_c/Z_c \\ Y_c/Z_c \end{pmatrix}$	-
Undistorted to Distorted Camera Plane Coordinates	$\mathbf{p}_d = \begin{pmatrix} x_d \\ y_d \end{pmatrix} = \lambda \mathbf{p}_n + dx,$ $dx = \begin{pmatrix} \lambda = 1 + \kappa_1 r^2 + \kappa_2 r^4 + \kappa_3 r^6, \\ 2\rho_1 xy + \rho_2(r^2 + 2x^2) \\ \rho_1(r^2 + 2y^2) + 2\rho_2 xy \end{pmatrix}, r^2 = x_n^2 + y_n^2$	$\begin{pmatrix} x_d \\ y_d \\ f(x_d, y_d) \end{pmatrix} = \begin{pmatrix} X_c \\ Y_c \\ Z_c \end{pmatrix},$ $f(x_d, y_d) = a_0 + a_1\rho + a_2\rho^2 + a_3\rho^3 + a_4\rho^4,$ $\rho^2 = x_d^2 + y_d^2$
Distorted Camera Plane to Image Coordinates	$\begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = \begin{pmatrix} f_x & \alpha f_x & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_d \\ y_d \\ 1 \end{pmatrix}$	$\begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = \begin{pmatrix} f_x & \alpha f_x & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_d \\ y_d \\ 1 \end{pmatrix}$



Fig. 4. Illustration of ground-grid alignment in the ground-truth camera, and the resulting flat-plane transformed image. Radial distortion is taken into account as can be seen by the bowed lines.

area of the road covered by the flat-plane image is 4 lanes side-to-side and 100 feet forward of the camera.

In aligning lane tracking results from the two systems with groundtruth, the groundtruth is kept unchanged for reference. The lane tracking estimates were manually scaled and offset to compensate for errors in camera calibration, camera placement and error in lane-width estimation. The operation consists of three things: 1) Global offset, 2) global scale, and 3) unwrapping amount. The global offset puts all 3 cameras on the same lateral position of the car. The global scale changes the scale of the estimates which results from errors in camera calibration. Unwrapping amount is specified to compensate for errors in lane-width estimates, which impact left-lane position estimates when the left-lane location is more than half a lane-width away. These alignment parameters are set once for the entire experiment.

Fig. 6 show results of left lane boundary estimation from both systems against groundtruth during the two segments of the test run. These segments show lane-following and lane-changing maneuvers. In these two segments, we can see the strength of the rectilinear

camera with approximately 3.01cm mean absolute error from to ground-truth, and omni-directional camera based lane tracking with 8.42cm mean absolute error. During a lane change maneuver, the distinction is not as pronounced. Using the same measure, the two systems approximately equalled in performance at 6.47cm and 6.18cm. For the entire sequence, Recti-VioLET and Omni-VioLET tracking compared with groundtruth resulted in mean absolute error of 5.87cm and 4.93cm, and RMS error of 5.39cm and 5.38cm, respectively.

As expected, the better performance achieved from using rectilinear images for lane tracking can be attributed to the higher resolution *regions* of the original images used for the flat-plane transformation. Although resolution is reduced to 1/10th the number of pixels in the omnidirectional image, its impact on the left lane boundary tracking performance, using the described method for comparison, resulted in only approximately 3 times more error during lane-following situations, and apparently no improvement in lane changing maneuvers.

V. ACCURACY COMPARISON

To quantify the difference between the minimum obtainable estimation error from the two lane-tracking systems, we performed a simulation to measure the error of road points introduced by projection and quantization processes of the system. The simulation model is motivated by the main processes of the VioLET system, namely 1) acquiring the image, 2) finding the flat-plane transformed image based on camera calibration, and 3) flat-plane image filtering to locate lane marker positions in the flat-plane image. The fourth step although it is integral to the system as well, 4) Kalman filtering of the detected marker positions and statistics, is not considered in the simulation model. Corresponding to each of these steps, first, uniformly random points lying on the virtual ground-plane

are generated and projected onto the two image planes using the respective camera parameters from the actual camera set-up. These projected points lying in the image frame are quantized to simulate the imaging process (step 1), then back-projected onto the ground-plane and quantized to a specified image resolution. The selected resolution depends on the actual resolution of the image and the subpixel accuracy of the feature detector (steps 2 and 3). The ability to acquire subpixel estimates of the lane marking position in the original or flat-plane image is modeled with subpixel accuracy parameter, which divides the quantization level by that factor. This raises the effective image resolution.

The Euclidean distance between the points' original locations to the new location is used as the primary error measure. Fig. 7 shows the result of the quantized back-projected points onto the ground-plane in 3D. Shown in the same figure is the range of errors, mean and maximum error envelope as a function of distance. The flat-plane transformed image resolution for the simulation is and 400×400 for the rectilinear and 200×200 for the omni-directional case. Given that the road is perfectly flat and the image resolution reflects the *effective* image resolution, the mean expected error is described by this figure. Specifically, if the lane tracker only uses the image of the road 10m away to track lanes, the least amount of error one can expect from the lane tracker on the average is the mean back-projection error of 4cm and 10cm from the rectilinear and omni-directional images respectively. Of course, because of filtering in steps 3 and 4, the accuracy of road feature locations is better than what the selected flat-plane image resolutions can achieve, which is currently set to the actual flat-plane image resolution of the flat-plane image in the road experiment. Still, this model is useful in determining the needed additional resolution in order for the omni-directional image based system to produce errors equal to the rectilinear image based system in the mean squared sense.

Omni-directional images with 720x480 resolution generates larger back-projection errors at all distances than a rectilinear image of the same size by approximately 2.3 times. The rate at which the error increases is also higher for the omni-directional image. Fig. 8 shows reduced rates of degradation as omni-directional image resolution is increased by 1.0x, 1.5x, and so on up to 3.0x. At 2.3 times the original image resolution, the omni-directional camera achieves the same level of accuracy in terms of average back-projection error as the rectilinear camera at the 720x480 resolution.

VI. DISCUSSION AND CONCLUDING REMARKS

We made a case that omni-directional cameras can produce useful lane tracking at 720x480 resolution. Omni-directional camera-based lane tracking can produce the same accuracy as from a rectilinear camera when the omni-directional image is 2.3 times the resolution. The next step would be to construct a system utilizing the environmental, vehicular, and driver attributes extracted from the single omni-directional camera for recognizing and predicting driver maneuvers, such as lane changing and intersection turning, with which to infer driver awareness. Fig. 9 shows the block diagram of this system. The input data consists of vehicle position in the lane as gathered from Omni-VIOLET, a head movement detection algorithm as proposed by Huang, *et al.* [8], a possible hand position estimation algorithm modified to operate on visible spectrum

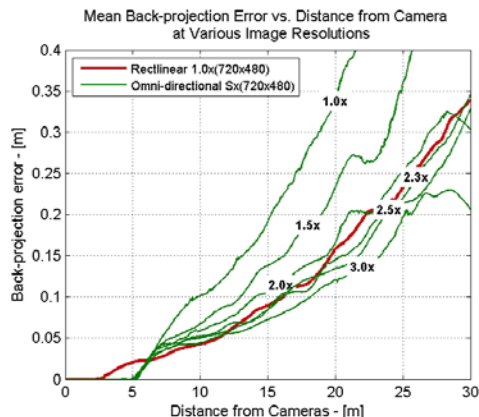


Fig. 8. The average back-projection error as a function of distance from camera are shown for various multiples of the original 720x480 image resolution for the omni-directional camera as compared to that for the 720x480 rectilinear image resolution




Fig. 9. Lane change intent classification using head pose and lane tracking from omni-directional cameras.

omni-directional images [3], and a number of vehicular dynamics information from the vehicle's CAN network. A sufficiently sized set of data would be collected, segmented, and annotated. Segments of the data would serve as a feature vector for a pattern classifier trained to classify segments corresponding to lane-changing [10], intersection-turning [4] or neither.

Through simulations and on-road experiments, we examined the amount of error that can be expected from detecting lane markers and estimating lane positions introduced in the imaging processing. We also introduced and evaluated a modified implementation of VioLET called Omni-VioLET. A systematic performance evaluation of the both systems were conducted by comparing the results against ground-truth, all from images captured along the same freeway course. We conclude that using an omni-directional camera results in about three times the amount of error at certain times and the same amount of error on the average, even though in producing a 360 degree field of view of the scene reduced the omni-directional image resolution, as compared to a rectilinear image, to $1/10^{th}$ of the number of pixels, and a $1/3^{rd}$ of the horizontal image resolution for imaging the road ahead.

REFERENCES

- [1] Massimo Bertozzi and Alberto Broggi. GOLD: a parallel real-time stereo vision system for generic obstacle and lane detection. *IEEE Transactions on Image Processing*, 7(1):62–81, January 1998.




Fig. 5. This illustrates the (a) Recti-VioLET and (b) Omni-VioLET lane tracking results from rectilinear images. The left column of images are the original images overlaid with the detected lane position. The right column of images are the flat-plane transformed images of the original. The two overlay lines in the transformed image depict the estimated left and right lane positions, and the boxes represent detected road features.




Fig. 6. The illustration shows a,d) the progression of the left and right lane-boundary position estimates over time as found by rectilinear and omnidirectional camera-based lane tracker and ground-truth. The distribution of error from ground-truth is shown for b,e) rectilinear camera based and c,f) omnidirectional camera based lane tracking.

[2] Massimo Bertozzi, Alberto Broggi, Massimo Cellario, Alessandra Fascioli, Paolo Lombardi, and Marco Porta. Artificial vision in road vehicles. *Proceedings of the IEEE*, 90(7), July 2002.

[3] Shinko Y. Cheng, Sangho Park, and Mohan M. Trivedi. Multi-spectral and multi-perspective video arrays for driver body tracking and activity analysis. *Computer Vision and Image Understanding* in press, 2007. doi:10.1016/j.cviu.2006.08.010.

[4] Shinko Y. Cheng and Mohan M. Trivedi. Turn intent analysis using body-pose for intelligent driver assistance. *IEEE Pervasive Computing, Special Issue on Intelligent Transportation Systems*, 5(4):28–37, October-December 2006.

[5] Wilfried Enkelmann. Video-based driver assistance—from basic functions to applications. *International Journal of Computer Vision*, 45(3):201–221, 2001.

[6] Tarak Gandhi and Mohan M. Trivedi. Vehicle surround capture: Survey of techniques and a novel omni video-based approach for dynamic panoramic surround maps. *IEEE Transactions on Intelligent Transportation Systems*, 7(3):293–308, September 2006.

[7] Itay Gat, Meny Benady, and Amnon Shashua. A monocular vision advance warning system for the automotive aftermarket. In *SAE 2005 World Congress and Exhibition*, April 2005.

[8] Kohsia Huang, Mohan M. Trivedi, and Tarak Gandhi. Driver’s view and vehicle surround estimation using omnidirectional video stream. In *IEEE International Symposium on Intelligent Vehicles*, pages 444–449, June 2003.

[9] Ken Ishikawa, Kazuyuki Kobayashi, and Kajiro Watanabe. A lane detection method for intelligent ground vehicle competition. In *Society of Instrument and Control Engineers (SICE) Annual Conference*, number 1, pages 1086–1089, August 2003.

[10] Joel C. McCall and Mohan M. Trivedi. Video based lane estimation and tracking for driver assistance: Survey, system and evaluation. *transits*, 7(1):20–37, March 2006.

[11] Joel C. McCall, David Wipf, Mohan M. Trivedi, and Baskar Rao. Lane change intent analysis using robust operators and sparse Bayesian learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 3, pages 59–67, 2005.

[12] Sergiu Nedevschi, Rolf Schmidt, Thorsten Graf, and Radu Danescu. 3d lane detection system based on stereovision. In *IEEE Intelligent Transportation Systems Conference*, pages 161–166, October 2004.

[13] Lars Petersson, Luke Fletcher, Alexander Zelinsky, Nick Barnes, and

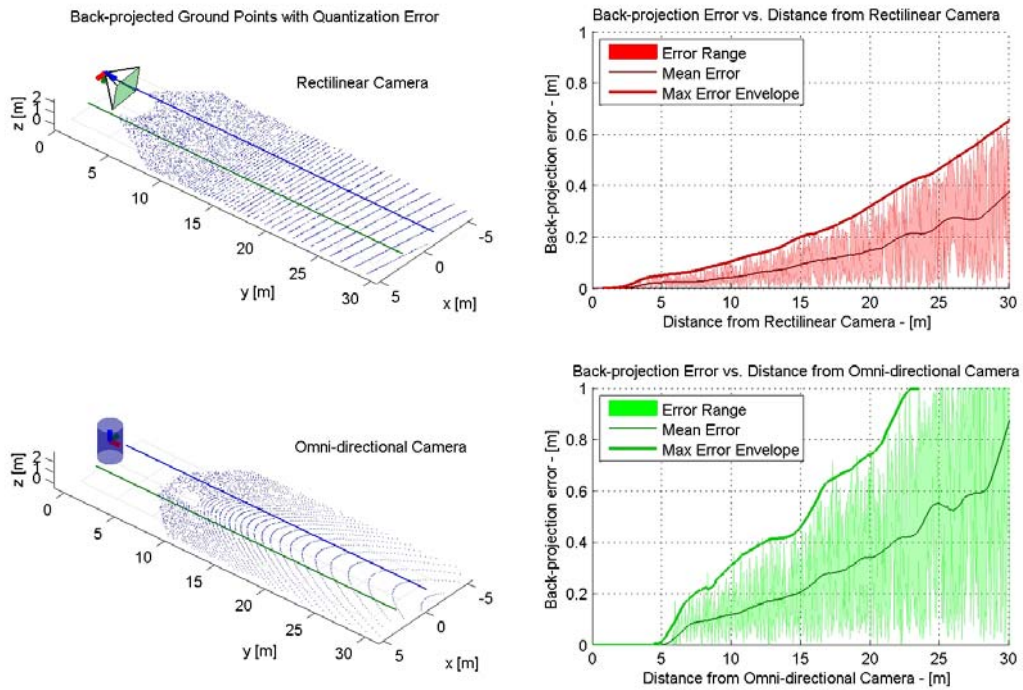


Fig. 7. Range, mean and envelope of maximum back-projection errors are plotted against distance of the point from the camera. The arrangement of the camera, ground plane and points are shown in the adjacent figure.

Fredrik Arnell. Towards safer roads by integration of road scene monitoring and vehicle control. *The International Journal of Robotics Research*, 25(1):53–72, January 2006.

[14] Davide Scaramuzza, Agostino Martinelli, and Roland Siegwart. A flexible technique for accurate omnidirectional camera calibration and structure from motion. In *IEEE International Conference on Computer Vision Systems*, pages 45–53, 2006.

[15] Zehang Sun, George Bebis, and Ronald Miller. On-road vehicle detection: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(5):694–771, May 2006.

[16] Mohan M. Trivedi, Tarak Gandhi, and Joel McCall. Looking in and looking out of a vehicle: Computer vision-based enhanced vehicle safety. *IEEE Transactions on Intelligent Transportation Systems*, January 2007.