Vehicle Localization using Road Markings

Tao Wu and Ananth Ranganathan

Abstract— Reliable lane-level localization is a requirement for many driver-assistance methods as well as for autonomous driving. Localization using cameras is desirable due to ubiquity and cheapness of sensors but is hard to achieve reliably. We propose a method towards reliable visual localization using traffic signs painted on the road such as arrows, pedestrian crossings, and speed limits. These road markings are relatively easily detected since they are designed to be highly conspicuous. Our method automatically recognizes road markings and uses features detected within them to compute the location of the vehicle. This provides an absolute global localization if the road markings have been surveyed before hand, and relative positioning information otherwise. We demonstrate using experiments and with groundtruth data that our method provides accurate lane-level visual localization under various lighting conditions and using various types of road markings.

I. INTRODUCTION

Reliable vehicle localization has been a topic of interest to the research community for a while [11]. This is because vehicle position is essential for autonomous driving, as well as for applications ranging from providing intuitive turn-by-turn directions to use of vehicle-to-vehicle communications. The most common positioning method is to use GPS, which while providing a number of advantages, has also a few well-known limitations. The two main limitations are overall accuracy and inability to function in urban areas and mountainous terrain where a large portion of the sky is blocked off. Overall accuracy of GPS used in current vehicles is of the order of 3 to 10 meters under best conditions [21], [25], which is insufficient for lane-level localization.

We are interested in performing localization using cameras. While the motivating factor is the low cost of the sensor as compared to others (for instance, lidars [11]), the corresponding challenges in making a robust system are also higher. Harsh lighting conditions and featureless environments such as broad highways pose special difficulties.

In this paper, we propose a novel method for using road markings painted on the road surface to estimate the precise location of the vehicle. Road markings are relatively easy to detect under various lighting conditions as they are designed to be conspicuous. Hence, use of these markings can be expected to improve the robustness of a general purpose vision-based localization system. An example of our system in action is shown in Figure 1.

We propose the use of a stereo camera pair for this work and use our road marking detector presented in [24] as the

Tao Wu is with the Department of Electrical & Computer Engineering, University of Maryland, College Park, MD, 20742, USA `taowu@umiacs.umd.edu`

Ananth Ranganathan is with Honda Research Institute, 425 National Ave, Mountain View, CA, 94043, USA `aranganathan@honda-ri.com`
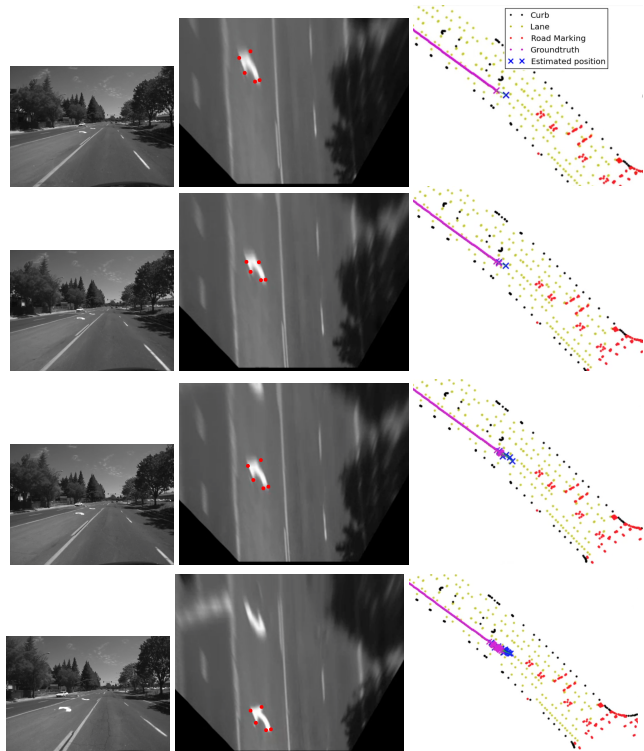
Fig. 1. A sequence of images showing pose estimation from a single road mark that is repeatedly detected as the vehicle drives past it. Each row shows (from left to right) the input image, the inverse perspective mapped image, and the vehicle position shown against the surveyed map data. Blue crosses mark the estimated position and purple crosses show the GPS groundtruth from the vehicle. Other points are part of the surveyed map data, which includes the GPS locations of the road markings(in red).

basis of our localization system. The detector uses a feature-based template approach. Once the road marking has been detected, we compute corners inside the detected road marking area and use these corners for estimating the pose of the vehicle. We present two methods for estimating the vehicle pose from the detected corners, one using triangulation and the second using a linear transform on a bird's eye view image. These two methods are compared using groundtruth.

Evaluation of our algorithm is performed using a groundtruth dataset of surveyed road markings which contains GPS coordinates of pre-specified corners within road markings on certain roads. The results demonstrate the accuracy and practicality of our approach.

II. RELATED WORK

Vehicle localization has a vast and varied literature based on the use of GPS devices. A number of methods try to

augment GPS with inertial sensors [4] and map matching methods [17], [19] to overcome GPS dropouts due to occlusion of satellites. Our interest is in low-cost accurate solution and hence, we do not focus on differential GPS and other augmentation services [8].

Apart from GPS, the use of any other sensors for localization results in drift over time due to accumulation of errors. In this work, we avoid drift by localizing against road markings that, when surveyed previously, provide absolute location. To our knowledge, such an approach has not been tried before. Our method, while it can be used by itself, is more suited for use after integration with a general purpose localization system consisting of visual odometry and location recognition, such as [12].

The common method for driftless localization in the literature is the use of a map registered to global coordinates. Such maps are built using simultaneous localization and mapping (SLAM) methods [20] using stereo vision [12], monocular vision [2], and lidars [11].

Road marking detection is a subject that has not been explored much previously. Veit et al. [23] provide an overview of existing road marking detection methods. However, they focus only on detection and not on classification of the road markings. Vacek et al. [22] use detection of lane markers to guide road marking detection and thus, cannot function in the absence of clear lane markings. Noda et al. [16] compute an eigenspace on a template by perturbing an ideal template under different illumination, scale, blur, and motion blur conditions. Classification is performed using eigenspace analysis. A method based on shape matching [10] is closest to our system in terms of applicability as it operates in real-time and can detect and classify a large number of complex road markings. However, we claim that our system is much simpler than the one presented therein, which uses a two-stage neural network classifier for which the number of hidden nodes and their connectivity has been carefully chosen by hand for the experiments.

III. ROAD MARKING DETECTION

We begin by providing an overview of our road marking detector which forms the basis of our pose estimation system. Our road marking detection system is based on learning feature-based templates of the markings using training images. In particular, our choice of features for defining the templates are FAST corner features [18]. During runtime, these templates are matched using a two step process of first selecting promising feature matches and subsequently performing a structural matching to account for the shape of the markings.

Templates are learned from training images which contain groundtruth bounding boxes for the road markings. The training images are first rectified to compensate for lens and perspective distortions, the latter in particular being done with an inverse perspective transform [1]. Subsequently, FAST corners are detected within the regions of interest, and the bounding box and the corners along with their HOG descriptors [3] are stored as the template for that particular
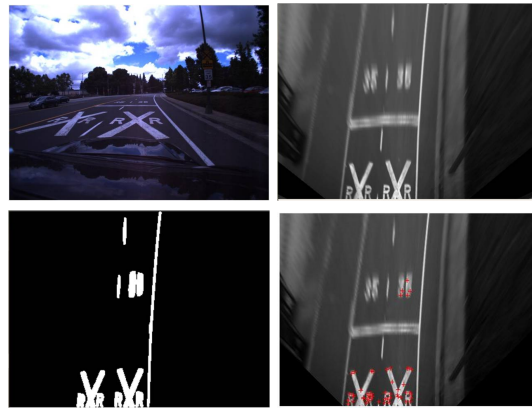

Fig. 2. The flow of our algorithm for road marking detection (clockwise from top left) the original image, rectified inverse perspective map image (birds-eye view), detected regions of interest (MSER features), corner features detected by the FAST corner detector (in red).

road marking type. Each road marking type (Stop, Bike lane, turn arrows etc.) may have many templates corresponding to various views, and road and lighting conditions.

At runtime, inverse perspective rectification, MSER detection, FAST corner detection, and HOG descriptor computation are performed on each test image. The signs in the testing images are then detected and identified based on the corner features. First, we find putative matching pairs of corners based on their HOG descriptors. Subsequently, we refine the result through a structural matching algorithm that matches the 2D geometry of the corners within the road marking. The geometry matching takes into account the possibility of multiple road markings in the image as well as the possibility of some features in a template not being detected due to changes in lighting and perspective. More details of the algorithm can be found in [24].

This detection method seamlessly handles complex road markings with multiple disconnected components such as the railroad marking shown in Figure 2. Further it is efficient enough to run in real time without the use of special hardware such as GPUs, and scales well with the number of types of road markings to be detected.

IV. SHADOW BOOSTING FOR IMPROVING ROAD MARKING DETECTION

We now describe our first contribution in this paper which extends the above explained road marking detection method to handle markings in shadow. In outdoor environments, especially in urban areas, shadows from trees and buildings on the road surface are not unusual. These can significantly change the appearance of road markings. An example of a road marking in shadows of trees is shown in Figure 3. We propose an algorithm which boosts the brightness of white road markings in non-uniform shadows in daytime outdoor environments and operates in real time.

The problem of color constancy is often handled in color images by converting to some color space such as HSI or HSV where the illumination and color occupy independent channels. Thus, colored road markings can be easily ex-

Fig. 3. Examples of detected road markings in shadow regions under various lighting conditions.

Fig. 4. An example image of the ratio between blue and red channel of an outdoor daytime scene. Note that the ratios are higher (image is brighter) in shadowed areas.

tracted from shadows. However, white road markings and gray road surfaces share the same hue and saturation so that the only difference between them is brightness. This makes it unreliable to extract these road markings from non-uniform shadows by simply converting the images to a different color space.

Although many shadow removal algorithms exist and have achieved good results, they either have a heavy computational burden and thus are not suitable for real-time applications, or are designed specifically for indoor environments. For instance, Finlayson, et. al. [5] present an entropy minimization-based algorithm to remove shadows from color images while Grest et al. [7] propose to remove shadows in indoor environments using color similarity metric. Liu et. al. [13] describe an approach that preserves texture consistency but, again, is quite computationally intensive. Miyazaki et. al. [15] describe a shadow removal method based on a hierarchical graph cut algorithm, which however is not suitable for our application as it requires user interaction.

A. Camera sensor model and modeling shadows

Sensors in color cameras convert visible lights to vectors in RGB color space. For color calibrated cameras, this process can be formulated using the following model [6]:

$$R_k = \sigma \int E(\lambda)S(\lambda)Q_k(\lambda)d\lambda, \ k \in \{R, G, B\} \quad (1)$$

where $E(\lambda)$, $S(\lambda)$ and $Q_k(\lambda)$ are the power spectrums of the illumination, the spectral reflectance and the camera sensitivity in channel k, with $k \in \{R, G, B\}$. σ is a normalizing constant.

In practice, the sensors are usually sensitive to a narrow frequency range in each color channel. So the sensitivity function $Q_k(\lambda)$ can be approximated by a Dirac delta function:

$$Q_k(\lambda) = \delta(\lambda - \lambda_k) \quad (2)$$

and thus, the camera sensor model becomes

$$R_k = \sigma E(\lambda_k)S(\lambda_k)Q_k, \ k \in \{R, G, B\} \quad (3)$$

In the camera sensor model, the sensitivity function is determined by the intrinsic property of the camera and is invariant to object surfaces. The spectral reflectance function, on the other hand, is constant for the same object surface. Hence, when a non-uniform shadow is present on the surface of a object, the difference between the shadow region and the non-shadow region is the illumination spectrum function $E(\lambda_k)$.

In a daytime outdoor environment, shadows are formed when the light from the sun is occluded and thus cannot reach the object surface. However, the objects in the dark shadow region are still visible. This implies that there is a ambient light source which provides illumination to all object surfaces. Let the illumination power spectrum of the sun light and the ambient light be denoted by $E_s(\lambda_k)$ and $E_a(\lambda_k)$. Let $alpha$ indicate whether the pixel is in shadow region. The pixel is in non-shadow region if $alpha = 1$, and is in shadow region if $alpha = 0$. Then the sensor model is:

$$R_k = \sigma(\alpha E_s(\lambda_k) + E_a(\lambda_k))S(\lambda_k)Q_k, \ k \in \{R, G, B\} \quad (4)$$

B. Shadow boosting

Our aim is to process the images so that the white road markings can be extracted more easily from shadows during the daytime. Assuming that the object surface is white, we have:

$$R_R = R_G = R_B \quad (5)$$

And the ratio between channel i and j is

$$r_{i,j} = \frac{R_i}{R_j} = \frac{\alpha E_s(\lambda_i) + E_a(\lambda_i)}{\alpha E_s(\lambda_j) + E_a(\lambda_j)} \quad (6)$$

We observe that during daytime, the dominated ambient light source is from the blue sky, which contains a large portion of $E_a(\lambda_B)$. This implies that for white objects, the shadow information is encoded in the ratio between the blue and the red channel $r_{B,R}$, which is one in non-shadow area, and is greater than one in shadow area. An example image of $r_{B,R}$ is shown in Figure 4. It shows that on the road surface, the shadow region is highlighted.

To study the relationship between the ratio $r_{B,R}$ and the illumination of the pixel, we manually cropped regions from the same object from outdoor videos and plotted the ratio $r_{B,R}$ against the logarithm of the illumination. The result suggests that there is roughly a linear relationship between them, i.e.,

$$Illumination = \exp(ar_{B,R} - b) \quad (7)$$

Fig. 5. (left) The inverse perspective mapped (IPM) image of a road marking in shadow region; (right) The boosted IPM image.

where a and b are parameters and can be learned from training images using a simple line fitting algorithm. Equation (7) can thus be applied on color images so that the shadow regions are boosted. An example of the road sign in shadow region and after shadow boosting is shown in Figure 5.

As can be observed from (7), our shadow boosting method is extremely efficient and suitable for real-time use as a pre-processing step in our road marking detection algorithm.

After shadow boosting method, the illumination in the shadow region is at the same level as in the non-shadow region, so that region extraction algorithms, such as MSER [14], can be applied to extract the regions of interest for our road marking detection and recognition, resulting in greater accuracy.

V. POSE ESTIMATION FROM ROAD MARKINGS

We now come to the main thrust of this paper, which involves using road markings to compute vehicle position. We tried two techniques for this :

- **From Perspective-n-Point:** The first technique involves solving the well-known Perspective-n-Point (PnP) problem in computer vision [], wherein the pose of the camera is calculated from known 3D to 2D point correspondences.
- **From the IPM image:** The second technique involves calculating the position of the camera within the inverse perspective mapped (IPM) image. If the extrinsic calibration of the camera is known, each pixel in the IPM image corresponds to some physical distance on the ground, and hence, given the pixel locations of the corners within the road markings in an IPM image, the relative pose of the camera to these corner features can be estimated.

The perspective-n-point method requires knowledge of 3D location so of the features on the road markings. We employed the visual odometry method of [12] for this purpose, which tracks points and computes their 3D locations using stereo-based triangulation and local bundle adjustment. However, in our experience, the results from the perspective-n-point method were worse than that using the IPM images. Hence, we focus on the IPM image method here on.

A. Relative pose estimation using IPM images

The position of the car relative to the detected road signs can be estimated based on the position of the corners of the road signs in the IPM images. Let the positions of the corners in the IPM images and their true locations in a reference coordinate system be denoted as $[U_d, V_d]$ and $[X_t, Y_t]$, respectively, where U_d, V_d, X_t, Y_t are $n \times 1$ vectors, and n is the number of corners.

There exists a scale s, a 2×2 rotation transform matrix R_s and a 2×1 translation vector t, such that the following mean square error is minimized.

$$s, R_s, t = \arg\min_{s, R_s, t} \| [X_t, Y_t]^T - (sR_s [U_d, V_d]^T + t) \|_2^2 \quad (8)$$

Let the position of the car on the IPM image be denoted by $[u_{car}, v_{car}]$ which can be obtained using extrinsic calibration, assuming the ground is flat. Then the position of the car $[x_{car}, y_{car}]$ in the same coordinate system as the road sign can be estimated by:

$$[x_{car}, y_{car}]^T = sR_s [u_{car}, v_{car}]^T + t \quad (9)$$

A good estimation of s and t can be obtained by:

$$s = s_d/s_t \quad (10)$$
$$t = t_d - t_t \quad (11)$$

where t_d and t_t are the mean coordinates of the corners in the IPM image $[U_d, V_d]$ and the coordinates in the reference coordinate system $[X_t, Y_t]$, respectively; s_d and s_t are the mean scale of the detected and template shapes. Let

$$[U_d', V_d']^T = [U_d, V_d]^T - t_d \quad (12)$$
$$[X_t', Y_t']^T = [X_t, Y_t]^T - t_t \quad (13)$$

Then, s_d and s_t are obtained by:

$$s_d = \frac{1}{n} \| [U_d, V_d]^T - t_d \|_2 \quad (14)$$
$$s_t = \frac{1}{n} \| [X_t, Y_t]^T - t_t \|_2 \quad (15)$$

And the rotation matrix R_s is:

$$R_s = V^T U, \text{ where } U\Sigma V = SVD([U_d', V_d']^T [X_t', Y_t']) \quad (16)$$

Thus the pose of the car relative to the road signs can be estimated. The advantage of this method is that it does not require any extra knowledge such as location of the 3D points in space. The drawback is the assumption of a flat ground or, at least, known pitch of the road, and also the requirement of the extrinsic calibration of the camera. However, as we demonstrate, the accuracy of the method compensates for these drawbacks.

B. Absolute Position estimation

As mentioned above, if the road markings have been surveyed previously, they can be used to obtain the absolute pose of the camera using the method provided in the previous section. However, some modifications to the template representation and road marking detection are required. We now describe these.

To obtain absolute camera pose from road markings, we built a database of corner features within various road markings containing high-accuracy GPS positions of all these corners. For each type of road marking, a specific and fixed
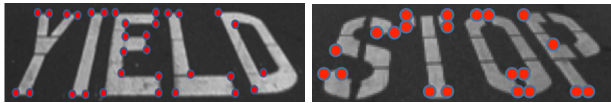
Fig. 6. Surveyed points for some road markings used for computing absolute camera pose

	w/o shadow boosting	w/ shadow boosting
TPR in non-shadow	90.1%	92.0%
TPR in shadow	12.5%	62.5%
FPR	0.9%	1.0%

set of corners were surveyed, a few examples of which are shown in Figure 6. The database is queried using a rough GPS location and returns a list of corner points that are close to the query location, along with their GPS coordinates, and corresponding road marking type. Then the absolute location of the car can be estimated with high-accuracy by comparing the position of the detected corners from the IPM images with the GPS positions of the corresponding corners in the list of the adjacent road signs.

The template used for road marking detection as described in Section III relies on detecting arbitrary corner points within the road markings. However, we now have specific corners (Figure 6) that we need to detect since only these have been surveyed. We achieve this by augmenting the templates used for the road marking detection with the set of required corners so that each template now contains a manually labeled set of corners that correspond to the surveyed GPS points. During runtime, the positions of these corners are then detected using the snake algorithm [9] on the IPM images.

The snake algorithm is typically used to find the contour of a given shape represented by a set of points, in our case the corner features. The algorithm iteratively finds locations for the points that minimize an energy function and lock on to a suitable contour. In our case, the contour is the road marking and the contrast in illumination between the markings and the road is sufficient for the algorithm to work reliably. More details of the snake algorithm are provided in [9].

Once the required corner features and the road marking have been detected, the GPS coordinates of the corners are queried in the database using the current rough location of the vehicle obtained from an inexpensive, commercial GPS device. Then the procedure of Section V-A can be applied with the reference coordinate system of the points used being the GPS coordinates. This yields the pose of the camera in the same coordinates.

VI. EXPERIMENTS

We present experiments for validating the shadow boosting algorithm for detecting shadowed roadmarks and for validating the road marking-based pose estimation algorithm. All the experiments were performed on a test car setup with a PointGrey Grasshopper camera running at 20 fps. The car was equipped with differential GPS, which was used only for groundtruth location.

A. Road marking detection with shadow boosting

We tested the proposed shadow boosting method on day-time videos using the road marking dataset presented in [24]. The videos were captured by a color camera mounted on

the roof of the car. Some examples of the detected white road markings in shadow regions are shown in Figure 3. The results demonstrate that the proposed method is able to detect white road markings in shadows, even in patterned shadow regions.

We compared the performance of detecting road markings with and without our shadow boosting method. The true positive detection rates (TPR) in both shadow region and non-shadow region and the false positive rates of both methods are summarized in Table I. It can be seen that shadow boosting effectively increases the detection rate in shadow regions while maintaining a high detection rate in non-shadow regions as well. The false positive rate is almost the same for both methods and the slight increase in the false positive rate with shadow boosting is not significant. Summarizing these results, shadow boosting improves overall road marking recognition significantly without any observable adverse effects in any mode of operation.

The shadow boosting preprocessing algorithm is fast since it has linear complexity on the size of the images. It takes less than 1 ms to process an 800x600 image on an Intel-i5 processor.

B. Pose estimation with groundtruth

We now present results for the pose estimation from road markings. We used road markings that were surveyed apriori using GPS so as to enable comparison to groundtruth. The difference of the surveyed GPS points on the road marks and the location of the car as reported by the on-board DGPS, was used as groundtruth for the pose estimation algorithm. Since this only provides groundtruth on location and not on the orientation of the vehicle, we only report results on position error.

The algorithm was tested, as before, on videos taken by cameras mounted on a car. Example of the detected corners on various road marking using the Snake method on the IPM images are shown in Figure 7 (left column).

Examples of the estimated car locations and the groundtruth car trajectory based on various detected road markings are shown in Figure 7. The map of the lanes, curbs and road markings are from our surveyed data. We tested the proposed method on six types of surveyed landmarks including turn arrows, stop signs, and railroad crossing signs. The mean absolute error between the estimated car location and the groundtruth was 0.99 meter on average based on more than 60 pose estimation instances. Since, in north america lane widths on most roads are on the order of 3 meters, this provides us a level of positioning accuracy that is much below the desired lane-level positioning.
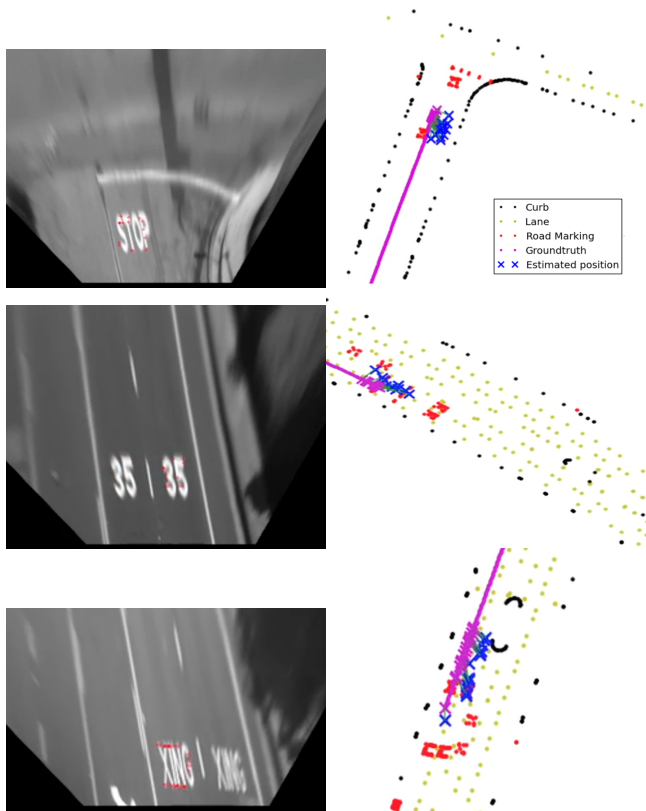
SLAM framework.

<div style="text-align:center">REFERENCES</div>

[1] M. Bertozzi, A. Broggi, and A. Fascioli. Stereo inverse perspective mapping: theory and applications. *Image and Vision Computing*, 8(16):585–590, 1998.

[2] L. Clemente, A. Davison, I. Reid, J. Neira, and J. Tardós. Mapping large loops with a single hand-held camera. Jun 2007.

[3] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 886–893, 2005.

[4] J. Farrell. *Aided navigation: GPS with high rate sensors*. McGraw-Hill New York, NY, USA:, 2008.

[5] G. D. Finlayson, M. S. Drew, and C. Lu. Entropy minimization for shadow removal. *Int. J. Computer Vision*, pages 35–57, 2009.

[6] G. D. Finlayson, S. D. Hordley, C. Lu, and M. S. Drew. Removing shadows from images. In *In ECCV 2002: European Conference on Computer Vision*, pages 823–836, 2002.

[7] D. Grest, J. michael Frahm, and R. Koch. A color similarity measure for robust shadow removal in real time. In *In Vision, Modeling and Visualization*, pages 253–260, 2003.

[8] E. Kaplan and C. Hegarty. *Understanding GPS: principles and applications*. Artech House Publishers, 2006.

[9] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *INTERNATIONAL JOURNAL OF COMPUTER VISION*, 1(4):321–331, 1988.

[10] A. Kheyrollahi and T. P. Breckon. Automatic real-time road marking recognition using a feature driven approach. *Machine Vision and Applications*, pages 1–11, 2010.

[11] J. Levinson, M. Montemerlo, and S. Thrun. Map-based precision vehicle localization in urban environments. In *Proceedings of Robotics: Science and Systems*, 2007.

[12] J. Lim, J.-M. Frahm, and M. Pollefeys. Online environment mapping. 2011.

[13] F. Liu and M. Gleicher. Texture-consistent shadow removal. In D. Forsyth, P. Torr, and A. Zisserman, editors, *Computer Vision ECCV 2008*, volume 5305 of *Lecture Notes in Computer Science*, pages 437–450. Springer Berlin Heidelberg, 2008.

[14] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide baseline stereo from maximally stable extremal regions. In *Proceedings of the British Machine Vision Conference*, pages 414–431, 2002.

[15] D. Miyazaki, Y. Matsushita, and K. Ikeuchi. Interactive shadow removal from a single image using hierarchical graph cut. In *Proceedings of the 9th Asian conference on Computer Vision - Volume Part I*, ACCV'09, pages 234–245, Berlin, Heidelberg, 2010. Springer-Verlag.

[16] M. Noda, T. Takahashi, D. Deguchi, I. Ide, H. Murase, Y. Kojima, and T. Naito. Recognition of road markings from in-vehicle camera images by a generative learning method. In *IAPR Conference on Machine Vision Applications*, 2009.

[17] M. Quddus, W. Ochieng, and R. Noland. Current map-matching algorithms for transport applications: State-of-the art and future research directions. *Transportation Research Part C: Emerging Technologies*, 15(5):312–328, 2007.

[18] E. Rosten, R. Porter, and T. Drummond. Faster and better: A machine learning approach to corner detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 32:105–119, 2010.

[19] T. Senlet and A. Elgammal. A framework for global vehicle localization using stereo images and satellite and road maps. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 2034–2041.

[20] S. Thrun. Robotic mapping: a survey. In *Exploring artificial intelligence in the new millennium*, pages 1–35. Morgan Kaufmann, Inc., 2003.

[21] US Dept. of Defense and GPS Navstar. Global positioning system standard positioning service performance standard, 2008.

[22] S. Vacek, C. Schimmel, and R. Dillmann. Road-marking analysis for autonomous vehicle guidance. In *Proceedings of the 3rd European Conference on Mobile Robots*, 2007.

[23] T. Veit, J.-P. Tarel, P. Nicolle, and P. Charbonnier. Evaluation of road marking feature extraction. In *Proceedings of the International IEEE Conference on Intelligent Transportation Systems*, 2008.

[24] T. Wu and A. Ranganathan. A practical system for road marking detection and recognition. In *Intelligent Vehicles Symposium (IV), 2012 IEEE*, pages 25 –30, june 2012.

[25] P. A. Zandbergen. Accuracy of iphone locations: A comparison of assisted gps, wifi, and cellular positioning. *Transactions in GIS*, 13(s1):5–25, July 2009.

Fig. 7. Corner points detected within road marks in the IPM images (in red, left column) are used to calculate position of a camera mounted on a vehicle (blue crosses, right column). The blue crosses in each image correspond to multiple road marking detections, only one of which is shown in the left column. Groundtruth from DGPS is shown as purple crosses while black dots represent road boundaries. The difference between the purple and blue crosses is the pose estimation error, which is, in general, less than 1 meter.

VII. CONCLUSIONS AND DISCUSSION

We presented a novel method for estimating camera pose from corner features detected on road markings. The method builds on our previous work on road marking detection [24] and uses extrinsic camera caibration and feature locations in a road marking on an inverse perspective mapped image to estimate pose of the camera with respect to the corner features.

As our pose estimation method relies on road marking detection, we also improved the road marking detector to work under difficult shadow conditions using a novel shadow boosting method proposed by us. This method improves overall road marking detection accuracy without affecting the false positive rate adversely.

The current method for pose estimation assumes a flat ground or knowledge of the pitch angle of the vehicle. This assumption can be removed by either detecting the ground plane or through the perspective-n-point method, which is more general. Our current results using this method demonstrated insufficient accuracy. However, we intend to improve this method in the future. It is also future work to integrate our pose estimation method into a visual SLAM algorithm as another type of measurement. This would enable continuous pose estimation and automatic drift correction within the