

Backward Segmentation and Region Fitting for Geometrical Visibility Range Estimation

Erwan Bigorgne and Jean-Philippe Tarel *

LCPC (ESE), 58 Boulevard Lefèbvre,
F-75732 Paris Cedex 15, France
bigorgne@lcpc.fr tarel@lcpc.fr

Abstract. We present a new application of computer vision: continuous measurement of the *geometrical visibility range* on inter-urban roads, solely based on a monocular image acquisition system. To tackle this problem, we propose first a road segmentation scheme based on a Parzen-windowing of a color feature space with an original update that allows us to cope with heterogeneously paved-roads, shadows and reflections, observed under various and changing lighting conditions. Second, we address the under-constrained problem of retrieving the depth information along the road based on the flat world assumption. This is performed by a new region-fitting iterative least squares algorithm, derived from half-quadratic theory, able to cope with vanishing-point estimation, and allowing us to estimate the geometrical visibility range.

1 Introduction

Coming with the development of outdoor mobile robot systems, the detection and the recovering of the geometry of paved and /or marked roads has been an active research-field in the late 80's. Since these pioneering works, the problem is still of great importance for different fields of Intelligent Transportation Systems. A precise, robust segmentation and fitting of the road thus remains a crucial requisite for many applications such as driver assistance or infrastructure management systems. We propose a new infrastructure management system: automatic estimation of the *geometrical visibility range* along a route, which is strictly related to the shape of the road and the presence of occluding objects in its close surroundings. Circumstantial perturbations such as weather conditions (vehicles, fog, snow, rain ...) are not considered. The challenge is to use a single camera to estimate the geometrical visibility range along the road path, i.e the maximum distance the road is visible.

When only one camera is used, the process of recovering the projected depth information is an under-constrained problem which requires the introduction of generic constraints in order to infer a unique solution. The hypothesis which is usually considered is the flat world assumption [1], by which the road is assumed

* Thanks to the French Department of Transportation for funding, within the SARI-PREDIT project (<http://www.sari.prd.fr/HomeEN.html>).

included in a plane. With the flat world assumption, a precise detection of the vanishing line is crucial. Most of the past and recent single camera algorithms are based on this assumption but differ by the retained model for the road itself [2–7]. One group of algorithms moves aside from the flat world assumption and provides an estimation of the vertical curvature of the road. In [8–10] the constraint that the road generally keeps an approximately constant width and does not tilt sideways is used.

In a general way, it should be noted that the quoted systems, which often relate to applications of lane-tracking /-following, work primarily on relatively 'not so far' parts of the road. In our case, the *geometrical visibility range* should be monitored along an interurban route to check for instance its compatibility with speed limits. We are thus released from the requirements of a strictly real-time application; however, both parts of the system, the detection and the fitting of the road, should manage the far extremity of the perceptible road, a requisite for which a road detection-based approach appears to be more adequate than the detection of markings.

This article is composed of two sections. The first section deals with the segmentation of the image. We restrict ourselves to structured road contrary to [11]. The proposed algorithm operates an adaptative supervised classification of each pixel in two classes: *Road* (\mathcal{R}) and *Other* (\mathcal{O}). The proposed algorithm is robust and benefits from the fact that the process is off-line. The second section deals with the region-fitting algorithm of the road working on the probability map provided by the segmentation step. The proposed algorithm follows an alternated iterative scheme which allows both to estimate the position of the vanishing line and to fit the borders of the road. The camera calibration being known, the positions of the vanishing line and of the far extremity of the perceptible road are enough to estimate the geometrical visibility range.

2 An adaptative probabilistic classification



Fig. 1. Examples of road scenes to segment, with shadows and changes in pavement material.

A dense detection of the road has been the object of many works which consider it as a two class pixel classification problem either in a supervised way [1, 12–14], or not [15, 16]. All these works face the same difficulty: the detection should be performed all along a road, when the appearance of the road is likely to strongly vary because of changes in the pavement material or because of

local color heterogeneity; the lighting conditions can also drastically modify the appearance of the road, see Fig. 1:

- The shadows in outdoors environment modify intensity and chromatic components (blue-wards shifting).
- The sun at grazing angles and/or the presence of water on the road causes specular reflections.

Several previously proposed systems try to tackle these difficulties. The originality of our approach is that we take advantage of the fact that the segmentation is off-line by performing backward processing which leads to robustness. We use a classification scheme able to cope with classes with possibly complex distributions of the color signal, rather than searching for features that would be invariant to well-identified transformations of the signal. In our tests, and contrary to [11], no feature with spatial or textured content (Gabor energy, local entropy, moments of co-occurrence matrix, etc.) appeared to be sufficiently discriminant in the case of paved roads, whatever is the environment. In practice, we have chosen to work in the La^*b^* color-space which is quasi-uncorrelated.

2.1 Parzen-windowing

In order to avoid taking hasty and wrong decisions, the very purpose of the segmentation stage is restricted to provide a probability map to be within the Road class, which will be used for fitting the road. The classification of each pixel is performed using a Bayesian decision: the posterior probability for a pixel with feature vector $\mathbf{x} = [L, a^*, b^*]$ to be part of the road class is:

$$P(\mathcal{R}/\mathbf{x}) = \frac{p(\mathbf{x}/\mathcal{R})P(\mathcal{R})}{p(\mathbf{x}/\mathcal{R})P(\mathcal{R}) + p(\mathbf{x}/\mathcal{O})P(\mathcal{O})} \quad (1)$$

where \mathcal{R} and \mathcal{O} denotes the two classes. We use Parzen windows to model $p(\mathbf{x}/\mathcal{R})$ and $p(\mathbf{x}/\mathcal{O})$, the class-conditional probability density functions (*pdf*). We choose the anisotropic Gaussian function with mean zero and diagonal covariance matrix Σ_d as the Parzen window.

Parzen windows are accumulated during the learning phases in two 3-D matrices, called $P^{\mathcal{R}}$ and $P^{\mathcal{O}}$. The matrix dimensions depend on the signal dynamic and an adequacy is performed with respect to the bandwidth of Σ_d . For a 24-bit color signal, we typically use two 64^3 matrices and a diagonal covariance matrix Σ_d with $[2, 1, 1]$ for bandwidth. This particular choice indeed allows larger variations along the intensity axis making it possible to cope with color variations caused by sun reflexions far ahead on the road. A fast estimation of $p(\mathbf{x}/\mathcal{R})$ and $p(\mathbf{x}/\mathcal{O})$ is thus obtained by using $P^{\mathcal{R}}$ and $P^{\mathcal{O}}$ as simple Look-Up-Tables, the entries of which are the digitized color coordinates of feature vectors \mathbf{x} .

2.2 Comparison

We compared our approach with [16] which is based on the use of color saturation only. We found although saturation usually provides good segmentation

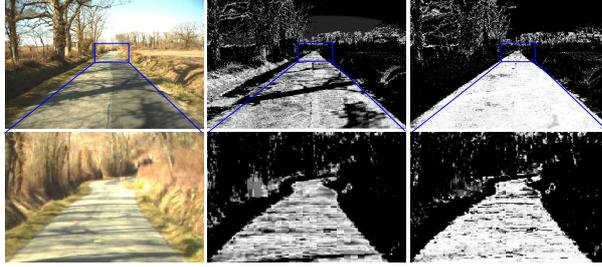


Fig. 2. Posterior probability maps based on saturation (middle) and $[L, a^*, b^*]$ vectors (right).

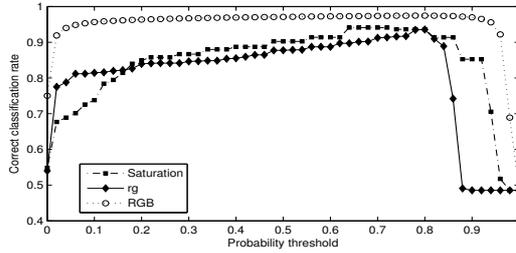


Fig. 3. Correct classification rate comparison for different types of feature.

results, this heuristic fails in cases too complex, where separability is no longer verified, see Fig. 2. Fig. 3 shows the correct pixel classification rate for a variable threshold applied on the class-conditional pdf $p(\mathbf{x}/\mathcal{R})$. Three types of feature have been compared on twenty images of different road scenes with a ground-truth segmented by hand: 1) the color saturation $\mathbf{x} = S = 1 - \frac{\text{Min}(R,G,B)}{\text{Mean}(R,G,B)}$, 2) the chromatic coefficients $\mathbf{x} = [r = \frac{R}{R+G+B}, b = \frac{B}{R+G+B}]$ and 3) the full color signal $\mathbf{x} = [L, a^*, b^*]$. The obtained results show the benefit of a characterization based on this last vector, which is made possible by the use of Parzen-windowing.

2.3 Robust Update

The difficulty is to correctly update the class-conditional pdfs along a route despite drastic changes of the road appearance. In case of online processing, thanks to the temporal continuity, new pixel samples are typically selected in areas where either road or non-road pixels are predicted to take place [12, 1]. In practice, this approach is not very robust because segmentation prediction is subject to errors and these errors imply damaged class-conditional pdfs that will produce a poor segmentation on the next image.

Due to our particular application which is off-line, we greatly benefit from a backward processing of the entire sequence: being given N images taken at regular intervals, the $(N - k)$ -th one is processed at the k -th iteration. For this image, new pixel samples are picked up in the bottom center part of the image



Fig. 4. 20 of the detected road connected-components in a image sequence. This particular sequence is difficult due to shadows and pavement material changes.

to update the '*Road*' pdf. The advantage is that we **know for sure** that these pixels are from the '*Road*' class since the on-board imaging system grabbing the sequence is on the road. Moreover, these new samples belongs to the **newly** observable portion of the road, and thus no prediction is needed. The update of the '*Other*' pdf is only made on pixels that have been labeled '*Other*' at the previous iteration. In order to lower as much as possible the risk of incorrect learning of the '*Road*' class, and then to prevent any divergence of the learning, the proper labeling of pixels as '*Road*' is performed by carrying out a logical-AND operation between the fitted model explained in the next section and the connected-component of the threshold probability map which is overlapping the bottom center-part of the image. The '*Other*' class is then naturally defined as the complementary. This process drastically improves the robustness of the update compared to online approaches.

Fig. 4 shows the detected connected-component superimposed on the corresponding original images with a probability threshold set at 0.5. These quite difficult frames show at the same time shadowed and overexposed bi-component pavement materials. The over-detections in the three first frames of the fourth row are due to a partially occluded private gravel road. This quality of results cannot be obtained with online update.

3 Road fitting

As explain in the introduction, the estimation of the shape of the road is usually achieved by means of edge-fitting algorithms, which are applied after the detection of some lane or road boundaries. Hereafter, we propose an original approach based on region-fitting which is more robust to missing data and which is also able to cope with vanishing line estimation.

3.1 Road models

Following [4], we use two possible curve families to model the borders of the road. First we use polynomial curves. $u_r(v)$ (*resp.* $u_l(v)$) models the right (*resp.* left) border of the road and is given as:

$$u_r = b_0 + b_1v + b_2v^2 + \dots + b_dv^d = \sum_{i=0}^d b_i v^i \quad (2)$$

and similarly for the left border. Close to the vehicle, the four first parameters b_0, b_1, b_2, b_3 are proportional respectively to lateral offset, to the bearing of the vehicle, to the curvature and to the curvature gradient of the lane. Second we use hyperbolic polynomial curves which better fit road edges on long range distances:

$$u_r = a_0(v - v_h) + a_1 + \dots + a_d \frac{1}{(v - v_h)^{d-1}} = \sum_{i=0}^d a_i (v - v_h)^{1-d} \quad (3)$$

and similarly for the left border. v_h is the position in the image of the vanishing line. The previous equations are rewritten in short in vector notations as $u_r = A_r^t X_{v_h}(v)$ (*resp.* $u_l = A_l^t X_{v_h}(v)$).

3.2 Half quadratic theory

We propose to set the region fitting algorithm as the minimization of the following classical least-square error:

$$e(A_l, A_r) = \int \int_{Image} [P(\mathcal{R}/\mathbf{x}(u, v)) - \Omega_{A_l, A_r}(u, v)]^2 dudv \quad (4)$$

between the image $P(\mathcal{R}/\mathbf{x}(u, v))$ of the probability to be within the road class and the function $\Omega_{A_l, A_r}(u, v)$ modeling the road. This region is parametrized by A_l and A_r , the left and right border parameters. Ω_{A_l, A_r} must be one inside the region and zero outside. Notice that function $A^t X_{v_h}(v) - u$ is defined for all pixel coordinates (u, v) . Its zero set is the explicit curve parametrized by A and the function is negative on the left of the curve and positive on its right. We thus can use the previous function to build Ω_{A_l, A_r} . For instance, function $g\left(\frac{A^t X_{v_h}(v) - u}{\sigma}\right) + \frac{1}{2}$ is a smooth model of the region on the right of the curve for any increasing odd function g with $g(+\infty) = \frac{1}{2}$. The σ parameter is useful to tune the smoothing strength. For a two-border region, we multiply the models for the left and right borders accordingly:

$$\Omega_{A_l, A_r} = \left(g\left(\frac{A_l^t X_{v_h}(v) - u}{\sigma}\right) + \frac{1}{2} \right) \left(\frac{1}{2} - g\left(\frac{A_r^t X_{v_h}(v) - u}{\sigma}\right) \right) \quad (5)$$

By substitution of the previous model in (4), we rewrite it in its discrete form:

$$e_{A_l, A_r} = \sum_{ij \in Image} \left[P_{ij} - \left(g\left(\frac{A_l^t X_i - j}{\sigma}\right) + \frac{1}{2} \right) \left(\frac{1}{2} - g\left(\frac{A_r^t X_i - j}{\sigma}\right) \right) \right]^2 \quad (6)$$

The previous minimization is non-linear due to g . However, we now show that this minimization can be handled with the half-quadratic theory and allows us to derive the associated iterative algorithm. Indeed, after expansion of the square in (6), the function g^2 of the left and right residuals appears: $g^2\left(\frac{A_l^t X_i - j}{\sigma}\right)$ and $g^2\left(\frac{A_r^t X_i - j}{\sigma}\right)$. Function $g^2(t)$ being even, it can be rewritten as $g^2(t) = h(t^2)$. Once the problem is set in these terms, the half-quadratic theory can be applied in a similar way as for instance in [6] by defining the auxiliary variables $\omega_{ij}^l = \left(\frac{A_l^t X_i - j}{\sigma}\right)$, $\omega_{ij}^r = \left(\frac{A_r^t X_i - j}{\sigma}\right)$, $\nu_{ij}^l = \left(\frac{A_l^t X_i - j}{\sigma}\right)^2$ and $\nu_{ij}^r = \left(\frac{A_r^t X_i - j}{\sigma}\right)^2$. The Lagrangian of the minimization is then obtained as:

$$\begin{aligned} \mathcal{L} = & \sum_{ij} \left[h(\nu_{ij}^l)h(\nu_{ij}^r) + \frac{1}{4}(h(\nu_{ij}^l) + h(\nu_{ij}^r)) + (2P_{ij} - 1)g(\omega_{ij}^l)g(\omega_{ij}^r) \right. \\ & \left. + (P_{ij} - 1/4) [-g(\omega_{ij}^l) + g(\omega_{ij}^r)] - h(\nu_{ij}^l)g(\omega_{ij}^r) + h(\nu_{ij}^r)g(\omega_{ij}^l) \right] \\ & + \sum_{ij} \lambda_{ij}^l \left(\omega_{ij}^l - \frac{A_l^t X_i - j}{\sigma} \right) + \lambda_{ij}^r \left(\omega_{ij}^r - \frac{A_r^t X_i - j}{\sigma} \right) \\ & + \sum_{ij} \mu_{ij}^l \left(\nu_{ij}^l - \left(\frac{A_l^t X_i - j}{\sigma} \right)^2 \right) + \mu_{ij}^r \left(\nu_{ij}^r - \left(\frac{A_r^t X_i - j}{\sigma} \right)^2 \right) \end{aligned} \quad (7)$$

The derivatives of (7) with respect to : the auxiliary variables, the unknown variables A_l and A_r , and the Lagrange coefficients λ_{ij}^l , λ_{ij}^r , μ_{ij}^l , μ_{ij}^r are set to zero. The algorithm is derived as an alternate and iterative minimization using the resulting equations.

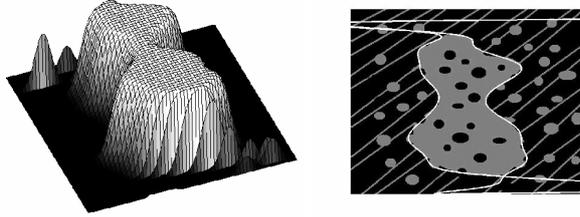


Fig. 5. 15th degree polynomial region fitting on a difficult synthetic image. Left: Ω_{A_l, A_r} 3D-rendering. Right: Obtained borders in white.

The proposed algorithm can handle a region defined either with polynomial curves (2) or with hyperbolic curves (3). It is only the design matrix (X_i) that changes. Fig. 5 presents a region fit on a difficult synthetic image with numerous outliers and missing parts. The fit is a 15th order polynomial. On the left side, the 3-D rendering of the obtained region model Ω_{A_l, A_r} is shown. Notice how the proposed region model is able to fit a closed shape even if the region borders are two explicit curves. We want to insist on the fact that contour-based fitting cannot handle correctly such images, with so many edge outliers and closings.

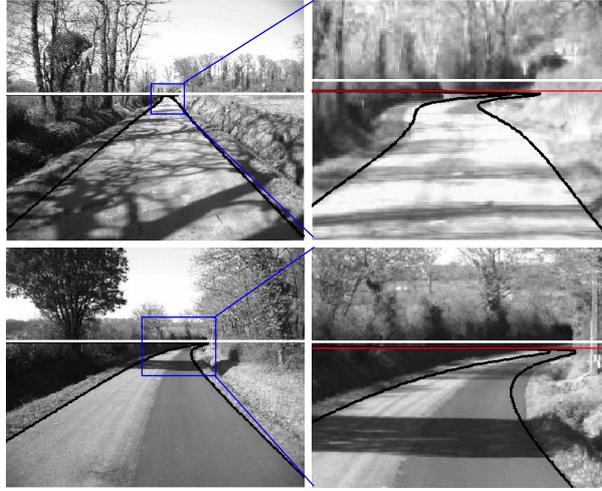


Fig. 6. Road region fitting results with 6th order hyperbolic polynomial borders. The images on the right provide a zoom on the far extremity of the road. The white line figures the estimated vanishing line; the red line figures the maximum distance at which the road is visible.

3.3 Geometrical Visibility Range

As explained in the introduction, for road fitting, it is of main importance to be able to estimate the position v_h of the vanishing line which parametrizes the design matrix. We solve this problem by adding an extra step in the previous iterative minimization scheme, where v_h is updated as the ordinate of the point where the asymptotes of the two curves intersect each other. In practice, we observed that the modified algorithm converges towards a local minimum. The minimization is performed with decreasing scales to better converge towards a local minimum not too far from the global one.



Fig. 7. Flat and non-flat road used for distance accuracy experiments.

Moreover, as underlined in [4], the left and right borders of the road are related being approximately parallel. This constraint can be easily enforced in the region fitting algorithm and leads to a minimization problem with a reduced number of parameters. Indeed, parallelism of the road borders implies $a_{ri} = a_{li}$, $\forall i \leq 1$, in (3). This constraint brings an improved robustness to the road region fitting algorithm as regards missing parts and outliers. Fig. 6 shows two images

taken from one of the sequences we experimented with. It illustrates the accuracy and the robustness of the obtained results when the local-flatness assumption is valid, first row. Notice the limited effect of the violation of this assumption on the second row at long distance. The white line shows the estimated vanishing line while the red line shows the maximum image height where the road is visible. The geometric visibility range of the road is directly related to the difference in height between the white and red lines.

Table 1. Comparison of the true distance in meters (true) with the distances estimated by camera calibration (calib.), and estimated using the proposed segmentation and fitting algorithms (estim.) for four targets and on two images. On the left is the flat road, on the right the non-flat road of Fig. 7.

target	true	calib.	estim.	true	calib.	estim.
1	26.56	26.95	27.34	33.59	23.49	34.72
2	52.04	56.67	68.41	59.49	60.93	61.67
3	103.52	98.08	103.4	111.66	111.58	114.08
4	200.68	202.02	225.96	208.64	1176.75	1530.34

Finally, we ran experiments to evaluate the accuracy of the estimated distances using one camera. On two images, one where the road is really flat and one where it is not the case, see Fig. 7, we compared the estimated and measured distances of white calibration targets set at different distances. The true distances were measured using a theodolite, and two kinds of estimation are provided. The first estimation is obtained using the camera calibration with respect to the road at close range and the second estimation is obtained using road segmentation and fitting. Results are shown on Tab. 1. It appears that errors on distance estimation can be important for large distances when the flat world assumption is not valid; but when it is valid the error is no more than 11%, which is satisfactory. A video format image is processed in a few seconds, but can be optimized further.

4 Conclusion

We tackle the original question of how to estimate the geometrical visibility range of the road from a vehicle with only one camera along inter-urban roads. This application is new and of importance in the field of transportation. It is a difficult inverse problem since 3D distances must be estimated using only one 2D view. However, we propose a solution based first on a fast and robust segmentation of the road region using local color features, and second on parametrized fitting of the segmented region using a priori knowledge we have on road regions. The segmentation is robust to lighting and road color variations thanks to a backward processing. The proposed original fitting algorithm is another new illustration of the power of half-quadratic theory. An extension of this algorithm is also proposed to estimate the position of the vanishing line in each image. We validated

the good accuracy of the proposed approach for flat roads. In the future, we will focus on the combination of the proposed approach with stereovision, to handle the case of non-flat roads.

References

1. Turk, M., Morgenthaler, D., Gremban, K., Marra, M.: VITS - a vision system for autonomous land vehicle navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **10**(3) (1988) 342–361
2. Liou, S., Jain, R.: Road following using vanishing points. *Comput. Vision Graph. Image Process.* **39**(1) (1987) 116–130
3. Crisman, J., Thorpe, C.: Color vision for road following. In: *Proc. of SPIE Conference on Mobile Robots*, Cambridge, Massachusetts (1988)
4. Guichard, F., Tarel, J.P.: Curve finder combining perceptual grouping and a kalman like fitting. In: *IEEE International Conference on Computer Vision (ICCV'99)*, Kerkyra, Greece (1999)
5. Southall, C., Taylor, C.: Stochastic road shape estimation. In: *Proceedings. Eighth IEEE International Conference on Computer Vision. Volume 1.* (2001) 205–212
6. Tarel, J.P., Ieng, S.S., Charbonnier, P.: Using robust estimation algorithms for tracking explicit curves. In: *Proceedings of European Conference on Computer Vision (ECCV'02). Volume I.*, Copenhagen, Denmark (2002) 492–507
7. Wang, Y., Shen, D., Teoh, E.: Lane detection using spline model. In: *Pattern Recognition Letters. Volume 21.* (2000) 677–689
8. Dementhon, D.: Reconstruction of the road by matching edge points in the road image. Technical Report Tech. Rep. CAT-TR-368, Center for Automation Research, Univ Maryland (1988)
9. Dickmanns, E., Mysliwetz, B.: Recursive 3D road and relative ego-state recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(2) (1992) 199–213
10. Chapuis, R., Aufrere, R., Chausse, F.: Recovering a 3D shape of road by vision. In: *Proc. of the 7th Int. Conf. on Image Processing and its applications*, Manchester (1999)
11. Rasmussen, C.: Texture-based vanishing point voting for road shape estimation. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition. Volume 1.* (2004) 470–477
12. Thorpe, C., Hebert, M., Kanade, T., Shafer, S.: Vision and navigation for the Carnegie-Mellon Navlab. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **10**(3) (1988) 362–373
13. Sandt, F., Aubert, D.: Comparaison of color image segmentations for lane following. In: *SPIE Mobile Robot VII*, Boston (1992)
14. Crisman, J., Thorpe, C.: Scarf: a color vision system that tracks roads and intersections. *IEEE Transactions on Robotics and Automation* **9**(1) (1993) 49–58
15. Crisman, J., Thorpe, C.: Unscarf, a color vision system for the detection of unstructured roads. In: *Proc. Of IEEE International Conference on Robotics And Automation.* (1991) 2496–2501
16. Charbonnier, P., Nicolle, P., Guillard, Y., Charrier, J.: Road boundaries detection using color saturation. In: *Proc. European Conference (EUSIPCO)*, Ile de Rhodes, Grce (1998) 2553–2556