HOG-like gradient-based descriptor for visual vehicle detection — **Source link**

Jon Arrospide, Luis Salgado, Javier Marinas

Institutions: Technical University of Madrid

Related papers:

- Histograms of oriented gradients for human detection

- On-road vehicle detection: a review

- A General Active-Learning Framework for On-Road Vehicle Recognition and Tracking

- Learning based symmetric features selection for vehicle detection

- Fast features for time constrained object detection

Share this paper:

HOG-like Gradient-based Descriptor for Visual Vehicle Detection

Jon Arróspide, Luis Salgado and Javier Marinas

Abstract— One of the main challenges for intelligent vehicles is the capability of detecting other vehicles in their environment, which constitute the main source of accidents. Specifically, many methods have been proposed in the literature for video-based vehicle detection. Most of them perform supervised classification using some appearance-related feature, in particular, symmetry has been extensively utilized. However, an in-depth analysis of the classification power of this feature is missing. As a first contribution of this paper, a thorough study of the classification performance of symmetry is presented within a Bayesian decision framework. This study reveals that the performance of symmetry-based classification is very limited. Therefore, as a second contribution, a new gradient-based descriptor is proposed for vehicle detection. This descriptor exploits the known rectangular structure of vehicle rears within a Histogram of Gradients (HOG)-based framework. Experiments show that the proposed descriptor outperforms largely symmetry as a feature for vehicle verification, achieving classification rates over 90%.

I. INTRODUCTION

The increasing awareness of road accident severity and impact has fostered the research on pre-crash sensing and, in general, on advanced driver assistance systems. In particular, most of the accidents are caused by other cars, therefore vehicle detection has become a central topic of investigation in the last years. Aside of active sensors such as radar and LIDAR, computer vision-based approaches for the detection of vehicles in traffic environments are gaining interest due to their flexibility, low cost and increased processing capabilities.

Most of video-based vehicle detection methods in the literature proceed in a two-stage fashion: hypothesis generation, and hypothesis verification. In the hypothesis generation stage, a quick search is performed throughout the image so that only a small subset of regions likely containing vehicles are retained. The search is typically based on some expected feature of vehicles, such as color [1], shadow [2], vertical edges [1], [3], or motion [4].

In the second stage, the presence of objects in the hypothesized regions is checked. Verification of hypothesis is usually addressed as a two-class supervised classification problem, in which a set of vehicle and non-vehicle samples are trained to derive representative features. Although complex features such as Wavelet decomposition and Histograms of Gradients are typically used for many applications, the need of real-time operation in traffic environment poses an important computational constraint. Therefore, simple features relating to the appearance of the vehicles are often preferred. In particular, the inherent symmetry of vehicle rear has been extensively used for vehicle verification [5], [6], [7], [8], [9].

However, an study of the real classification potential of symmetry is missing. Indeed, although this feature is very appealing due to its low computational load and its intuitively high perceptual power, we feel that its intrinsic simplicity may also pose an important limitation in the achievable separability between vehicle and non-vehicle classes. Therefore, in this paper an in-depth study of the vehicle classification performance of symmetry is carried out by adopting the classical symmetry feature definition in [10] within a Bayesian decision framework. Specifically, the distribution of the vehicle and the non-vehicle classes in the symmetry feature space is analyzed using a public database, and the best-fit distributions are found for each of them. The study reveals that the performance of symmetry-based classification is indeed limited to a rate of around 80%.

In order to overcome this limitation, in this paper a new vehicle descriptor based on gradient analysis within an HOG-like framework is proposed. In contrast to traditional HOG-based approaches, which involve heavy computation and training requirements, the descriptor is designed to be simple and fast by taking advantage of the knowledge of the structure of vehicles in order to comply with the real-time operation constraint. In particular, two different properties of a canonical vehicle in terms of gradient are exploited. First, vehicle rear features mainly horizontal and vertical edges, since both its contour and its inner elements (e.g., rear window, license plate) have a nearly rectangular shape. Second, vehicles typically have larger edge density than a random image in traffic environments. These properties are exploited by dividing the image into smaller cells (as done in HOG) and locally comparing the observed features with the expected patterns. Exhaustive experiments are performed in the same database, and the proposed descriptor is proven to largely outperform symmetry and to yield very good classification results.

II. SYMMETRY

Symmetry is one of the most outstanding shape features for recognition of many types of objects, including vehicles. In particular, mirror symmetry with respect to a vertical axis is characteristic of vehicle rears when projected onto the image plane. This feature has been widely used by many authors for applications relating to car-following situations. Many of them (e.g., [5],[9]) rely on the symmetry feature

definition introduced by Zielke in [10]. This method searches for symmetry around a vertical axis in the intensity of a grayscale image of size $R \times C$ by first scanning horizontal lines. The symmetry axis, x_s, is shifted along the definition interval, and all possible widths are considered up to size of the input image, $w \leq C$. Taking into account the range of these parameters, for each row of the image a matrix of symmetry values is computed as:

$$S_{1D}(x_s, w, y_0) = \frac{\int E_n(u, x_s, w, y_0)^2 du - \int O(u, x_s, w, y_0)^2 du}{\int E_n(u, x_s, w, y_0)^2 du + \int O(u, x_s, w, y_0)^2 du}$$

where E_n and O represent the even and odd parts of the image intensity function, I, and u scans the hypothesized width, $u = x - x_s$, $-w/2 \leq u \leq w/2$. These parts are given by [10]:

$$E_n(u, x_s, w, y_0) = E(u, x_s, w, y_0) - \frac{1}{w} \int_{-w/2}^{w/2} E(v, x_s, w, y_0) dv$$

$$E(u, x_s, w, y_0) = \begin{cases} \frac{1}{2}\left(I(x_s + u, y_0) + I(x_s - u, y_0)\right), & u \in [\frac{-w}{2}, \frac{w}{2}] \\ 0, & \text{otherwise} \end{cases}$$

$$O(u, x_s, w, y_0) = \begin{cases} \frac{1}{2}\left(I(x_s + u, y_0) - I(x_s - u, y_0)\right), & u \in [\frac{-w}{2}, \frac{w}{2}] \\ 0, & \text{otherwise} \end{cases}$$

The final symmetry measure is the average of the 1D symmetry values along the vertical direction. In addition, $-1 \leq S_{1D}(x_s, w, y_0) \leq 1$, so appropriate scaling is performed so that the final measure is within the normalized range $[0, 1]$. Hence, it is:

$$S_{2D}(x_s, w) = \frac{1}{2}\left(\frac{1}{R}\sum_{y=1}^{R} S_{1D}(x_s, w, y) + 1\right)$$

The parameters (x_s^o, w^o) maximizing the matrix $S_{2D}(x_s, w)$ determine the potential vehicle boundaries within the input image, and the value $S_{2D}(x_s^o, w^o)$ relates to the probability that it be a vehicle according to symmetry analysis.

III. SYMMETRY-BASED BAYESIAN CLASSIFIER

The described symmetry feature definition has been adopted in this study. The purpose is to analyze whether this feature conveys relevant information for vehicle classification, both individually or in a multi-cue scheme. Let us first visualize the distribution of the symmetry feature in our image database (presented in [11], available in the Internet [12]) in order to analyze the vehicle/non-vehicle separability. This database contains images extracted from highway video sequences, which are divided in four categories according to the relative pose of the vehicle with respect to the camera: close/middle range in the front, in the left, and in the right, and far range. For each of these regions, the database consists of 1000 positive (i.e. containing vehicles) and 1000 negative samples, selected to emulate the output of a hypothesis generator. As an example, Fig. 1 displays the normalized symmetry histogram for vehicle and non-vehicle classes in the front close/middle range. The lines joining the histogram
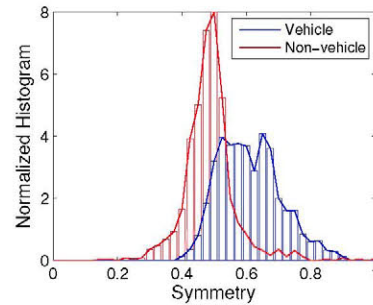

Fig. 1. Normalized histogram of symmetry feature for vehicle (blue) and non-vehicle (red) classes in the front close/middle range.
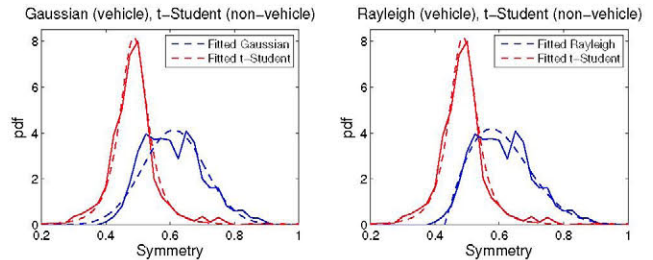

Fig. 2. Fitting of different pdfs to the observed data distribution for the front close/middle range. For the non-vehicle class (red), a t-Student distribution with $\nu = 3$ degrees of freedom, a shift of $\Delta = 0.4873$ and a scaling of $s = 0.045$ with respect to the standard t-Student is fitted, whereas for the vehicle class (blue) Gaussian ($[\mu, \sigma] = [0.6112, 0.0967]$) and Rayleigh ($[\sigma, \Delta] = [0.1452, 0.43]$) pdfs are proposed. The solid lines correspond to the data distribution in Fig. 1.

values have also been painted to get an approximation to the underlying probability density function. As expected, there is a high-degree of overlapping between the classes, which limits the performance of symmetry-based classification.

A. Likelihood Modeling

A Bayesian classification framework is proposed for the symmetry feature. As opposed to traditional approaches based on hard decision, this allows not only to perform classification but also to have the probability that a given new sample belongs to the vehicle class. This is valuable information, especially in multi-feature classification schemes, where it can be used to weigh the contribution of each feature to the final decision.

We first aim at finding the distributions that best fit the data. As regards the non-vehicle class observe that the curve is similar to Gaussian, i.e., symmetric and bell-shaped, but has larger tails. This is suitably modeled by a t-Student distribution, characterized by its degree of freedom ν. The parameter ν that best adapts the curve is selected for each image region, and the distribution is shifted and scaled according to the observed mean and variance. The vehicle class, in turn, seems more challenging to characterize. In search of simplicity, let us first attempt to fit a Gaussian distribution to the data. Parallelly, let us make a second hypothesis: the mass of the distribution is slightly concentrated on the left side, i.e., it is right-skewed, therefore we shall propose a Rayleigh distribution. Since by definition this distribution is zero at the origin, this must be appropriately shifted.

TABLE I

RESULTS OF KOLMOGOROV-SMIRNOV TEST

K-S-Test	Gaussian			Rayleigh		
	h	p-value	ks	h	p-value	ks
Front	1	0.0144	0.0495	0	0.1407	0.0363
Left	1	0.0029	0.0570	0	0.7930	0.0204
Right	1	0.0028	0.0571	0	0.1721	0.0348
Far	1	0.0091	0.0517	0	0.7436	0.0214

TABLE II

PERFORMANCE OF SYMMETRY FEATURE

Accuracy	Front	Left	Right	Far
	80.18	80.46	78.54	80.50

Fig. 2 shows the fitting of the aforementioned distributions to the data for the front close/middle range. Specifically, the left and right figures display Gaussian and Rayleigh fitting to the vehicle class, respectively, while the t-Student models non-vehicle data in both figures. Observe that the t-Student distribution fits almost perfectly the non-vehicle class data. In turn, visual inspection reveals that Rayleigh distribution adapts the vehicle data better than the Gaussian. Fig. 3 shows analogous results for the remaining regions. The better behavior of the Rayleigh distribution is confirmed by the Kolmogorov-Smirnov test. This test compares a sample with a reference probability distribution, and gives an idea of the probability that the sample is drawn from the distribution (null hypothesis). In particular, the test returns a statistic ks that measures the supremum of the difference between the theoretical and the empirical cumulative distribution function, and a p-value, which is the probability of obtaining a statistic at least as extreme as the one that was actually observed, assuming that the null hypothesis is true. If the p-value is above a given significance level, the null hypothesis cannot be rejected ($h = 0$), otherwise it is rejected ($h = 1$). Table I summarizes the results of the test assuming that the vehicle data are either Gaussian or Rayleigh-distributed. As expected, with the conventional significance level of 5%, the Rayleigh hypothesis cannot be rejected for any of the image regions and the p-value is well above the significance level, which supports its election to model the symmetry distribution associated to the vehicle class.

B. Experiments and Results

Once the probability densities have been modeled for both classes, classification experiments are performed in the database [12] using 50% holdout cross-validation 5-fold. As stated, a Bayesian classifier is employed: prior probabilities of the classes are assumed to be equal, thus classification is performed only with regard to conditional densities. The results of the experiments are enclosed in Table II in terms of accuracy (i.e., proportion of correctly classified test samples) As can be observed, there are no big differences in performance between the image regions and the accuracy is around 80%. This value is rather low, which leads us to the proposal of a new descriptor based on gradient within a HOG framework, which largely outperforms symmetry, as will be
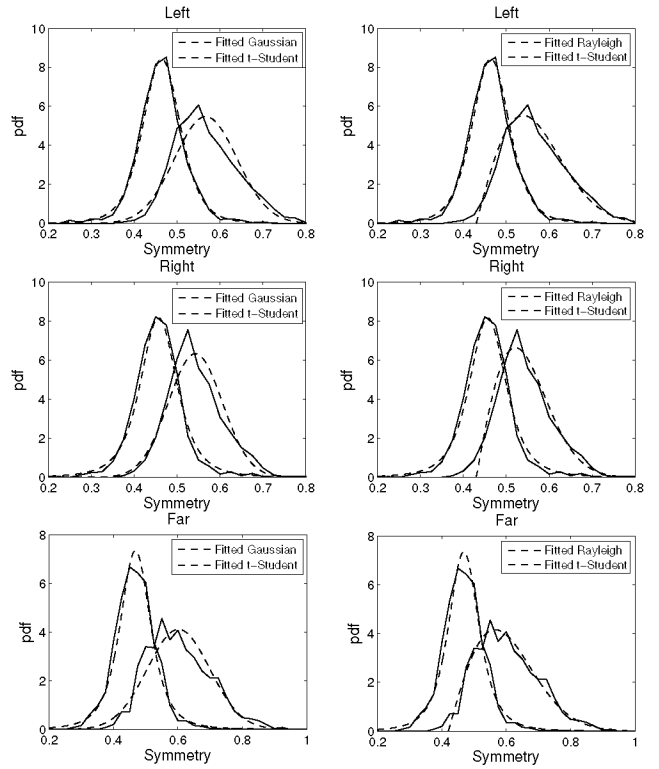

Fig. 3. Fitting of pdfs to the vehicle and non-vehicle classes for the left, right and far regions. The solid lines correspond to the observed data distribution in these regions. The parameters of the t-Student, Gaussian and Rayleigh functions are $(\nu = 5, \Delta = 0.4639, s = 0.045),(\nu = 5, \Delta = 0.4568, s = 0.045);(\nu = 3, \Delta = 0.4679, s = 0.05);(\mu = 0.5667, \sigma = 0.0729),(\mu = 0.5426, \sigma = 0.063),(\mu = 0.6019, \sigma = 0.097);(\sigma = 0.1096, \Delta = 0.43),(\sigma = 0.0912, \Delta = 0.43),(\sigma = 0.1458, \Delta = 0.42)$ respectively for the left, right and far regions.

shown in Section V.

IV. GRADIENT-BASED DESCRIPTOR

Once the symmetry feature has been proven to have limited performance for vehicle classification, the challenge remains to find another descriptor that, on the one hand, involves lower processing requirements than traditional complex features such as HOG or Wavelet features, thus enabling real-time operation, and on the other hand maintains good performance. In this study, we propose to use the previous knowledge about the structure of the vehicle to define a new feature based on HOG [13], which alleviates the severe computational load of traditional HOG approaches while achieving much better performance than symmetry. In particular, two different properties are considered that characterize a canonical vehicle instance in terms of gradient. On the one hand, vertical and horizontal gradients are expected to be dominant. On the other hand, due to the rich texture content of vehicles, they are supposed to typically have larger gradient density than a random image in traffic environments. These properties are exploited within a HOG-like framework: the image is divided in cells of size s, and for each cell a histogram of the gradient[1] orientations over the pixels

---

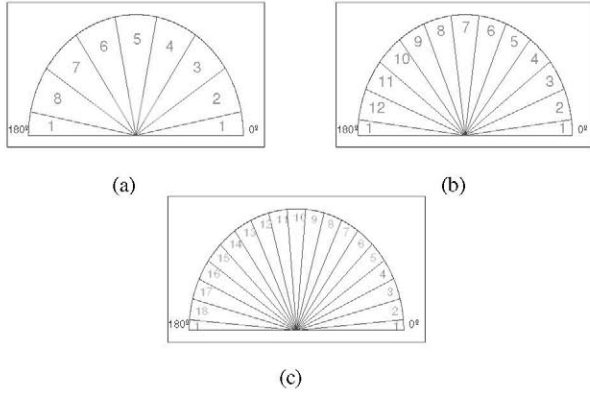[1]Gradient is computed using a 3 × 3 Sobel operator.

Fig. 4. Possible configurations of HOG regarding the number of orientation bins. Configurations with (a) 8, (b) 12, and (c) 18 bins are considered.
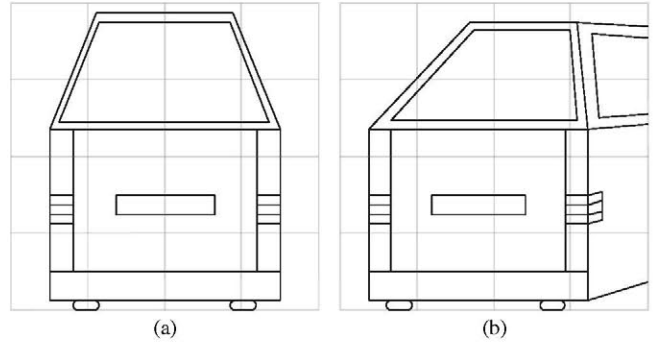

Fig. 5. Sketches of vehicle rear views from different perspectives: (a) the vehicle is in the front of the observer, (b) the vehicle is shifted to the left hand side of the observer.

is constructed according to a previously defined orientation binning. Specifically, the range of gradient orientation angles [0,180) is divided in uniformly spaced sectors. Pixels with gradient orientations inside each sector accumulate to the corresponding bin of the histogram proportionally to the magnitude of their gradient. Three possible configurations are considered, with $\beta = 8, 12$ and 18 bins respectively, as shown in Fig. 4.

A two-feature descriptor is defined to discriminate between vehicle and non-vehicle samples. The first feature measures the distance of each cell to a vertical/horizontal gradient cell. Specifically, the dominant orientation, o_d, of the HOG associated to the cell is retained, and the number of bins in between the dominant orientation and the vertical (o_v) and horizontal (o_h) orientations, is counted:

$$d_v = d(o_d - o_v)$$
$$d_h = d(o_d - o_h) \quad (1)$$

where $d(\cdot)$ denotes distance in number of bins, and d_v and d_h are the distances to the vertical and horizontal orientations, respectively. Then, the first feature is the distance to the vertical or the horizontal orientation, which is defined as $f_1 = \min(d_v, d_h)$. For example, if $\beta = 18$ it is $o_v = 1$ and $o_h = 10$ (see Fig. 4). Then, if $o_d = 4$ (which means that the gradient orientation histogram of the cell has a maximum between $25°$ and $35°$), the distances are $d_v = 3$ and $d_h = 6$, therefore the first feature computed in cell c is $f_1^c = 3$. As stated, this feature measures the distance to the vertical or horizontal orientations, which are expected in a typical vehicle rear pattern (see Fig. 5 (a)). Note that the used database contains instances including vehicles only partially, hence shifts of the expected vehicle structure in Fig. 5 (a) can occur both sidewards or up/downward. That is the reason why vertical and horizontal gradients are searched for in all the cells instead of particularizing the search according to the cell position.

Although the defined feature fits roughly the expected pattern of the vehicle, there is a part of the image in which vertical and horizontal gradients are not exclusive. Indeed, due to the perspective effect, the upper-left and right contours

of the vehicle are often tilted. Hence, in the right upper-half cells, gradients with angle below a certain threshold t_r (experimentally set to $25°$) are considered consistent with the expected vehicle pattern and included in the first orientation bin. Analogously, for the left upper-half cells, gradients with angle above $t_l = 180 - t_r = 155°$ are also included in the expected pattern.

The second feature counts the number of gradient-wise significant cells. Indeed, a typical vehicle instance is expected to have more gradient content than random background images in traffic environments, which usually contain homogeneous patches corresponding to the road and the sky. The HOG cell-structure allows to locally measure the gradient density and to define a global feature as the count of gradient-wise significant cells. In order not to lose sensitivity, a very low gradient magnitude threshold is defined to consider a pixel significant. Then, a cell is deemed to be significant if the proportion of pixels in the cell is above a certain threshold t_p (this is set experimentally, as shown in Section V).

The final gradient-based descriptor is composed of the number of significant cells, f_2, and the mean distance to the vertical and horizontal bins taking into account the significant cells, i.e., $f_1 = \frac{1}{f_2} \sum_{C_s} f_1^c$ (where C_s is the set of significant cells in the image). If $f_2 = 0$, it indicates that the image is nearly homogeneous, therefore it is considered to belong to the non-vehicle class.

A. Descriptor for Left, Right and Far Regions

The descriptor explained above is designed for vehicles rears. It must be taken into account that the projection of vehicles in the left, right and far regions does not exactly project the pattern in Fig. 5 (a) into the image, since the perspective is not perpendicular to the vehicle rear. In many cases, the hypothesized bounding boxes for vehicles (and accordingly the samples included in the database) contain a part of the vehicle side, which in turn typically leads to additional gradients in the image.

For instance, a vehicle samples of the left close/middle region could have a structure similar to that shown in Fig. 5 (b). This pattern displays several differences with respect to the pattern in Fig 5 (a). To begin with, a lower edge appears

in the right side of the vehicle pointing to the vanishing point (which is in the horizon line, i.e., the edge has a positive slope, and the gradient a negative slope). Therefore, in order to include this in the designed model, the descriptor for the left region is changed not to penalize negative-slope gradients in the lower rightmost cells. Observe that some other edges appear in the upper right part of the vehicle, however, these are usually close-to-horizontal, therefore they are supported by the original descriptor. In turn, the tilted edges in the upper part of the vehicle structure also change: the left edge is even more tilted due to the perspective effect, whereas the right edge turns almost vertical. The threshold t_b is accordingly relaxed to $135°$ in the upper-left cells, while only horizontal and vertical edges are expected in the upper-right cells.

Furthermore, the descriptor also addresses the special cases produced by continuous slope structures such as lane markings and guardrails. Indeed, although most of those are appropriately handled by the original descriptor (they are diagonal due to the perspective, hence their distance f_1 is large), they can lead to trouble in road sides since they are almost horizontal. For instance, when $\beta = 18$, a lane marking of angle $14°$ results in several cells of $f_1^c = 1$ throughout the image, which is relatively low and could be classified as vehicle. So as to tackle this, the central cells of the image penalize all negative-slope gradients with neither horizontal nor vertical dominant angle to have a maximum difference $f_1^c = f_1^{max}$, which depends on the configuration ($f_1^{max} = 2, 3,$ and 4 for $\beta = 8, 12$ and 18, respectively; see Fig. 4). This is only applied in the central and not in the outer (25% leftmost and rightmost) cells, since vehicle instances often contain heterogeneous background elements in the latter. For the same reason, the information content is deemed to be more trustworthy in the central cells, therefore a weighting function is applied to the image that doubles the importance of the central cells with respect to the outer cells in the final computation of f_1.

The descriptor for the right region is completely analogous to the left pattern, only mirrored. As for the far region, similar reasoning is followed. In this case, tilted edges are expected both in the upper left and right parts of the vehicle, hence $t_r = 25°$ and $t_l = 155°$, as in the front close/middle range. In turn, this range includes the left and right far regions, therefore in order to avoid conflicts due to nearly-horizontal lane markings or guardrails, penalization $f_1^c = f_1^{max}$ is performed for all cells featuring neither vertical nor horizontal dominant orientation. Furthermore, a weighting function similar to that explained for the left region is applied favoring the central cells.

V. EXPERIMENTS AND DISCUSSION

A Bayesian classifier is used to evaluate the performance of the above defined gradient-based descriptor. Since there is no prior information about the probability that the samples belong to one class or the other, the decision rule reduces to selecting the class with the largest likelihood. Those are defined to be normal, as usually done in multivariate
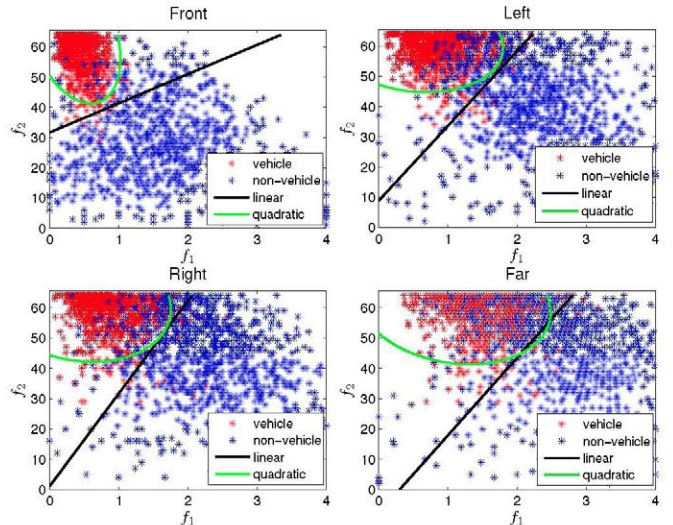

Fig. 6. Vehicle (red) and non-vehicle (blue) data distribution in the gradient-based descriptor space. The corresponding linear and quadratic decision surfaces are also shown.

density modeling. Two different classifiers have been used. The first assumes that the covariance matrices of the vehicle and non-vehicle classes are equal, which results in a linear decision surface [14]. The second classifier allows different covariance matrices for each class, thus involving quadratic decision surfaces. Examples of the data distribution and the corresponding decision surfaces are illustrated in Fig. 6 for $s = 8$ and $\beta = 18$ for all the image regions.

In order to find the testing error, experiments are performed in the same manner as with the symmetry feature, i.e., 50% holdout cross-validation is used 5 times. Different experiments are carried out according to the parameters of the descriptor: the size of the cells, s, the number of orientations, β (both as defined in Section IV), and the proportion of pixels to consider a cell significant gradient-wise, t_p. Accuracy rates as a function of s and β are given in Table III, and the t_p yielding the best performance, $t_p^{(m)}$, is specified for each case (experiments are performed by varying t_p between 5% and 50%). The first conclusion is that the quadratic classifier largely outperforms the linear classifier regardless of the remaining parameter configuration (almost 2% accuracy gain in average). In addition, performance is enhanced by increasing the orientation binning, especially for the left, right and far regions, and by increasing the cell size s. The best results are thus obtained for $s = 16$ and $\beta = 18$. The parameter $t_p^{(m)}$ is heuristically set for each configuration, and typically ranges between 5% and 10% for $s = 16$ (a larger threshold is too stringent: many cells containing edge pixels are below it, as many homogeneous patches are also included on account of the large size of the cells). For the best configuration, the accuracy is 92.48%, which is significantly better than that of the most widely used explicit descriptor, i.e. symmetry.

TABLE III

PERFORMANCE OF GRADIENT-BASED DESCRIPTOR IN TERMS OF ACCURACY AS A FUNCTION OF THE PARAMETERS s, β AND $t_p^{(m)}$

$s = 4$	$\beta = 8$				$\beta = 12$				$\beta = 18$			
	linear		quadratic		linear		quadratic		linear		quadratic	
	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$
Front	95.46	15	96.90	15	95.52	30	97.14	10	95.40	15	97.32	30
Left	84.50	5	86.36	5	87.24	15	89.96	30	89.52	10	91.34	30
Right	84.04	20	86.70	5	87.54	20	88.92	15	88.96	10	90.34	5
Far	80.60	15	83.50	15	84.08	15	86.22	35	86.06	45	87.92	45
Mean	86.15		86.37		88.60		90.56		89.99		91.73	

$s = 8$	$\beta = 8$				$\beta = 12$				$\beta = 18$			
	linear		quadratic		linear		quadratic		linear		quadratic	
	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$
Front	95.50	15	97.08	15	95.44	15	97.46	15	95.72	20	97.32	20
Left	85.00	10	88.24	10	88.94	5	90.44	5	89.60	5	91.42	10
Right	85.50	5	88.08	5	89.22	5	89.94	5	90.18	5	91.10	5
Far	81.32	5	82.30	10	85.34	15	86.66	15	86.94	5	88.66	20
Mean	86.93		88.93		89.74		91.13		90.61		92.13	

$s = 16$	$\beta = 8$				$\beta = 12$				$\beta = 18$			
	linear		quadratic		linear		quadratic		linear		quadratic	
	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$	Rate	$t_p^{(m)}(\%)$
Front	94.90	20	96.68	15	94.94	25	97.00	5	94.92	25	96.94	10
Left	85.30	10	87.72	5	88.72	5	91.46	5	89.98	5	91.98	10
Right	85.50	5	89.58	5	90.46	5	91.60	5	90.94	5	91.76	5
Far	82.60	5	84.20	5	86.10	10	87.22	20	88.06	5	89.22	10
Mean	87.08		89.55		90.06		91.82		90.98		92.48	

VI. CONCLUSIONS

The contributions of this paper are two-fold. First, a study has been presented that analyzes the performance of the most widely used feature for vehicle verification, i.e., symmetry. In particular, this has been introduced into a two-class Bayesian classification framework. Appropriate modeling of the likelihood of each class has been carried out by finding the best-fitting distributions via Kolmogorov-Smirnov tests. The performance of the symmetry feature has been shown to have an upper limit of around 80% for all image regions. As a second contribution, a new HOG-like gradient-based descriptor has been proposed for vehicle verification. This has been tested under a similar Bayesian classification framework, and proven to yield much higher separability between classes than symmetry. In particular, the optimal configuration reaches a correct classification rate of 92.48% in average, which is as high as 96.94% for the front close/middle range. The effectiveness of the proposed descriptor is thus proven, and hence we propose its use for vehicle verification either as a single cue or within a multi-feature scheme.

VII. ACKNOWLEDGMENTS

REFERENCES

[1] L. Gao, C. Li, T. Fang and Z. Xiong, Vehicle Detection Based on Color and Edge Information, *Image Analysis and Recognition, Lecture Notes in Computer Science, Int. Conf. on Image Analysis and Recognition,* vol. 5112, 2008, pp. 142–150.

[2] W. Liu, X. Wen, B. Duan, H. Yuan and N. Wang, "Rear Vehicle Detection and Tracking for Lane Change Assist", *in IEEE Intelligent Vehicles Symposium,* Istanbul, Turkey, 2007, pp. 252–257.

[3] G. Y. Song, K. Y. Lee and J. W. Lee, "Vehicle Detection by Edge-Based Candidate Generation and Appearance-based Classification", *in IEEE Intelligent Vehicles Symposium,* Eindhoven, The Netherlands, 2008, pp. 428 - 433.

[4] J. Arróspide, L. Salgado, M. Nieto and R. Mohedano, Homography-based Ground Plane Detection Using a Single On-board Camera, *IET Intelligent Transport Systems,* vol. 4, no. 2, 2010, pp. 149– 160.

[5] C. Hoffmann, T. Dang and C. Stiller, "Vehicle Detection Fusing 2D Visual Features", *in IEEE Intelligent Vehicles Symposium,* Parma, Italy, 2004, pp. 280–285.

[6] A. Broggi, P. Cerri and P. C. Antonello, "Multi-Resolution Vehicle Detection Using Artificial Vision", *in IEEE Intelligent Vehicles Symposium,* Parma, Italy, 2004, pp. 310–314.

[7] T. Liu, N. Zheng, L. Zhao, H. Cheng, "Learning Based Symmetric Features Selection for Vehicle Detection", *in IEEE Intelligent Vehicles Symposium,* Las Vegas, USA, 2005, pp. 124–129.

[8] G. Cheng and X. Chen, "A Vehicle Detection Approach Based on Multi-Features Fusion in the Fisheye Images", *in Int. Conf. Computer Research and Development,* Shanghai, China, 2011, pp. 1–5.

[9] T. K. ten Kate *et al.,* "Mid-range and Distant Vehicle Detection with a Mobile Camera", *in IEEE Intelligent Vehicles Symposium,* Parma, Italy, 2004, pp. 72–77.

[10] T. Zielke, M. Brauckmann, W. v. Seelen, Intensity and Edge-based Symmetry Detection with an Application to Car-following, *CVGIP: Image Understanding,* vol. 58, no. 2, 1993, pp. 177–190.

[11] J. Arróspide, L. Salgado and M. Nieto, Video Analysis-based Vehicle Detection and Tracking Using an MCMC Sampling Framework, *EURASIP Journal on Advances in Signal Processing,* vol. 2012, no. 2, 2012.

[12] The GTI Vehicle Image Database, Grupo de Tratamiento de Imágenes (GTI), Polytechnic University of Madrid, available at http://www.gti.ssr.upm.es/data/

[13] N. Dalal and B. Triggs, "Histograms of Oriented Gradients for Human Detection", *in IEEE Conf. on Computer Vision and Pattern Recognition,* San Diego, USA, vol. 1, 2005, pp. 886–893.

[14] C. M. Bishop, *Pattern Classification, Second Edition, Ch. 2: Bayesian Decision Theory,* John Wiley & Sons; 2000.