

# Vehicle logo recognition in traffic images using HOG features and SVM

D. F. Llorca, R. Arroyo, M. A. Sotelo

**Abstract**—In this paper a new vehicle logo recognition approach is presented using Histograms of Oriented Gradients (HOG) and Support Vector Machines (SVM). The system is specifically devised to work with images supplied by traffic cameras where the logos appear with low resolution. A sliding-window technique combined with a majority voting scheme are used to provide the estimated car manufacturer. The proposed approach is assessed on a set of 3.579 vehicle images, captured by two different traffic cameras that belong to 27 distinctive vehicle manufacturers. The reported results show an overall recognition rate of 92.59%, which supports the use of the system for real applications.

**Index Terms**—Vehicle Manufacturer Recognition, Vehicle Logo Recognition, Traffic Images, HOG features, SVM, sliding window, majority vote.

## I. INTRODUCTION

Nowadays, the ability of recognizing the vehicle manufacturer (car make) by the standard License Plate Recognition (LPR) systems in the context of Intelligent Transportation Systems (ITS) applications, is getting more and more important. Current traffic surveillance, speed control and access control systems rely on the correct detection of the license plate to identify the monitored vehicle. However, as discussed in [1], LPR systems fail, and most of the recognition errors are not due to missing recognition power, but to segmentation errors related with the high variability of the environmental conditions.

Recognition errors can be the cause of improper fines, if these ones are directly sent to drivers without manual supervision of the results provided by the LPR systems. In addition, LPR systems are unable to detect fake car plate numbers. Automatic vehicle identification systems increase their reliability and robustness by including more details about the monitored vehicle, such as the vehicle color [2], the plate color [3], the car make [4], the car model, etc. If the detected license plate and vehicle manufacturer are correlated with data stored on police or homeland security databases, LPR errors and suspicious vehicles can be automatically identified without manual supervision.

As stated in [5], Vehicle Manufacturer Recognition (VMR) is a subject with relatively limited research in the field. Most of the approaches concerning VMR described in the literature are based on the recognition of the vehicle logo, which is a clear indicator of the car make. Thus, the VMR problem can be simplified to a multi-classification problem of all the vehicle logos of the different vehicles manufacturers.

D. F. Llorca, R. Arroyo, M. A. Sotelo are with the Computer Engineering Department, Polytechnic School, University of Alcalá, Madrid, Spain. email: llorca, sotelo@aut.uah.es.

In this paper a vehicle manufacturer recognition system based on the classification of the vehicle logo is proposed. Logo detection is assisted by a previously developed license plate recognition stage. A sliding window technique is then applied in a region of interest (ROI) defined above the detected license plate. Local Binary Patterns (LBP), Scale-Invariant Feature Transform (SIFT) and Histograms of Oriented Gradients (HOG) have been studied as features to represent the vehicle logo. A multi-class Support Vector Machine (SVM) is then used to classify all the regions provided by the sliding window stage. Finally, a majority vote approach is implemented to estimate the logo using the binary outputs given by the SVM. Key differences with previous approaches are the use of HOG features in combination with a multi-class SVM to deal with logo recognition, the evaluation concerning a texture-based descriptor such as LBP, and the use of a sliding window technique combined with a majority vote approach. In addition, our VMR system has been specially devised to be used with traffic images, like the ones depicted in Fig. 1. Whereas other approaches recognize logos with high resolution regions, our proposal works with images in which the logos appear contained in a low-medium resolution region. The proposed approach is assessed on a set of 3.579 vehicle images, captured by two different traffic cameras (see Fig. 2), that belong to 27 distinctive vehicle manufacturers.



Fig. 1. Examples of traffic images used for vehicle logo recognition.

## II. RELATED WORK

The VRM is usually carried out through Vehicle Logo Recognition (VLR). The procedure usually involves a License Plate Recognition (LPR) module, followed by a coarse-to-fine approach to identify the logo ROI. Finally, some logo recognition method is applied. In [6], logo recognition is addressed using template matching and edge-orientation histograms. The use of SIFT descriptor with neural networks for logo recognition was proposed in [7], [5].

Most of the approaches available in the literature uses SIFT descriptor. Thus, in [4] a set of 1200 logo images corresponding to 10 distinctive vehicle manufactures are




Fig. 3. Overall architecture of the proposed VRM approach.




Fig. 2. Prototypes used to record the traffic sequences.

used to assess a SIFT-based approach. An enhanced feature matching method that merges SIFT points from different logo samples provides an accuracy of 97%. However, the samples used in [4] do not correspond with naturalistic traffic images. Close-up logo images are available which allows ROI sizes of around  $100 \times 100$  pixels ( $\approx \times 3$  the size we use).

Finally, in [8], where Fourier shape descriptors and inner structure mean square error analysis are used, it is stated that large variations in brightness, vehicle features in the foreground and specular reflections render the SIFT approach practically useless.

### III. SYSTEM DESCRIPTION

The overall architecture of the proposed VMR system is depicted in Fig. 3. As can be observed, the recognition module is triggered by previous stages if these ones detect the motion and the license plate of the vehicle that appears on the images. LPR module is only applied if we detect motion on the traffic images. Once the logo recognition module is switched on, a sliding window technique is applied to select ROIs where the logo can be located. A fast implementation of the HOG features is run for each previously selected ROI, and these feature vectors are classified by means of a multi-class SVM. The estimated logo that will correspond to the vehicle manufacturer is finally obtained using a majority voting scheme which takes all the binary outputs of the SVM classifiers for each one of the pre-selected ROIs.

The proposed majority voting scheme is defined in a multi-frame fashion, i.e., the SVM binary outputs used to recognize the logo correspond to all the frames in which the same vehicle appears. Depending on the vehicle speed, the

camera captures between 2-7 frames corresponding to the same vehicle with full visibility of its logo. This approach is graphically described in the Finite State Machine (FSM) depicted in Fig. 4. The *No Car* state means that neither motion nor a license plate appear on the images. Once we detect the first appearance of a vehicle license plate on the images, we change to *Car* state, in which the logo recognition module is run. We accumulate all the binary outputs of the multi-class SVM. The majority voting approach is only applied once the license plate is lost.




Fig. 4. Finite State Machine (FSM).

#### A. Vehicle Logo Localization

The first stage uses a frame-to-frame differencing approach to detect changes between consecutive frames corresponding to the traffic sequence captured by the camera. Thus, if no vehicles appear on the images, the CPU load remains low. Once the motion module detects a considerable change between two adjacent frames, a LPR system, previously developed by our research group, is executed providing both the text and the location of the license plate. Then, the following assumption is made: in most cases, the vehicle logo will appear in a region located just above the license plate. This assumption is not valid for certain types of vehicles in which the license plate is not located in the center but in one of the sides of the vehicle frontal area. Accordingly, we apply a sliding widow technique using different region sizes (square windows), sliding the windows along the vertical axis that separates the license plate in two equal sized regions (see Fig. 5).

It is worth to mention that other approaches are specifically devised to provide only one region fitted to the vehicle logo [6], [9], [10]. In our case, we supply a set of regions that are fed to the classifier. Although this procedure increase the CPU load since HOG features and SVM classification have to be run for more than one region, we expect better results since we do not rely on the output of only one classification per image. A similar approach was successfully applied by




Fig. 5. Overall view of the sliding window approach.

the authors in the context of pedestrian [11] and pavement [12] recognition.

### B. Vehicle Logo Recognition

As suggested by [4], the natural approach to perform vehicle logo recognition is to use SIFT features due to its invariance to perspective, rotation and shift [13]. Although we have not implemented the whole method suggested by [4], we quickly found that SIFT features are insufficient for low resolution images captured under different lighting conditions (see Fig. 6). This problem was also referenced by [8] with high resolution images under everyday lighting conditions. In this way, we finally discarded the use of SIFT features.




Fig. 6. Insufficient SIFT performance for two images containing the same car manufacturer logo under different lighting conditions.

We consider to use the approach of Dalal and Triggs [14] to model the local shape and appearance of the vehicle logos using well-normalized dense histograms of oriented gradients (HOG). Local gradients are binned in accordance with orientation, weighted depending on their magnitude, within a spatial grids of cells with overlapping blockwise contrast normalization. Within each overlapping block of cells, a feature vector is obtained by sampling the histograms from the contributing spatial cells. The feature vectors for all overlapping blocks are concatenated to produce the final feature vector which is fed to the classifier. The average gradient image for some of the vehicle logos is depicted in Fig. 7 and the overview of the HOG/SVM architecture is showed in Fig. 8.

Our choice of HOG parameters is as follows. The ROIs provided by the logo localization stage are resized to a resolution of  $32 \times 32$  pixels. We use fine scale gradients  $((-1, 0, 1)$  masks with smoothing), fine orientation binning (8 bins) and  $2 \times 2$  blocks of either  $8 \times 8$  pixels cells. Finally, overlapping blocks contrast normalization ( $L_2$ -norm) is applied. The concatenated feature vector size is 288.

Finally, we have tested the use of Local Binary Patterns [15] to determine if the texture of the logos contains relevant information when accomplishing vehicle logo classification. The reported results showed that the system performance were not improved either working alone or in combination with HOG features. Accordingly, we have discarded the use of texture features such as LBP to perform vehicle logo recognition.




Fig. 7. Average gradient image for some of the vehicle logos used in this work.

Once we have all the HOG features corresponding to all the regions given by the sliding window approach for all the frames in which one specific vehicle appears (see Fig. 4), these feature vectors are then classified by means of a multi-class SVM [16] based on the standard *one-versus-one* approach to reduce the single multi-class problem into multiple binary classification problems. The maximum-margin hyper-plane which best separates all the classes is computed by solving an optimization problem. A linear kernel function has been chosen as it achieves the best accuracy when comparing with radial basis kernel function. In addition, the linear kernel function provides an on-line application function that involves a dot-product with the weight vector independent of the support vectors. Finally, a majority voting scheme is used to assign the final result as the class which has the largest number of votes.

## IV. EXPERIMENTS

A digital camera with a resolution of  $640 \times 480$  pixels and a focal length of  $50mm$  was placed at two different road bridges pointing to one specific lane of a highway (see Fig. 2). The images were captured on two different days, under different lighting conditions (from sunny to cloudy). The sequences comprise a total of 3.579 vehicle images that belong to 27 distinctive vehicle manufacturers. Ground truth corresponding to the license plate number and its position as well as the logo and its position was obtained by manual labeling the corresponding bounding boxes in the camera images. The number of samples for each car manufacturer




Fig. 8. Overview of the HOG/SVM architecture.

TABLE I  
NUMBER OF SAMPLES FOR EACH VEHICLE MANUFACTURER.

Manufacturer	# samples / # vehicles	Manufacturer	# samples / # vehicles
Alfa Romeo	17 / 5	Man	24 / 6
Audi	232 / 64	Mazda	25 / 10
BMW	142 / 47	Mercedes	293 / 85
Chevrolet	19 / 6	Mini	7 / 2
Citroen	276 / 107	Nissan	107 / 35
Fiat (old)	17 / 4	Opel	257 / 100
Fiat (new)	123 / 47	Peugot	285 / 105
Ford	337 / 123	Renault	325 / 120
Honda	46 / 16	Seat	254 / 93
Iveco	10 / 3	Skoda	69 / 27
Jaguar	12 / 3	Toyota	137 / 52
Jeep	17 / 4	Volkswagen	366 / 118
Kia	33 / 11	Volvo	142 / 45
Lancia	7 / 1	<b>TOTAL</b>	<b>3.579 / 1.239</b>

are shown in Table I. Some examples of the 27 vehicle logos are depicted in Fig. 9.

Although the distribution is somewhat unbalanced, it is related with the Spanish motor vehicle census. As can be observed, some of the manufacturers correspond with trucks (Man and Iveco). Furthermore, different logos that correspond to the same car manufacturer are independently labeled (e.g., Fiat).




Fig. 9. Examples of the 27 vehicle manufacturers.

#### A. Off-line classification results

The performance of our HOG/SVM classifier is firstly evaluated in an off-line fashion using the manually labeled logo regions. All the samples are scaled to  $32 \times 32$  pixels and HOG features are obtained. We compared two SVM kernels: linear and RBF. According to the data distribution shown in Table I we select 2/3 of the samples for each one of the vehicle logos for the training phase, leaving 1/3 for the test. Table II compares the performance of the

HOG/rbfSVM classifier with the HOG/linSVM for each one of the vehicle manufacturers. The overall accuracy is 95.88% and 92.87% for the linear and RBF kernels respectively. It is worth to mention that logos with an accuracy of 0.00% correspond with classes in which the number of training samples is below 12, a value which is clearly insufficient for generalizing.

#### B. Single-frame logo recognition results

In order to evaluate the proposed logo recognition approach, including both the vehicle logo localization stage and the majority voting scheme, we trained a SVM classifier using the samples that correspond to 2/3 of the sequence images, i.e., 2.386 logo samples maintaining the distribution showed in Table I. The rest of the images (1/3) were used in a single-frame fashion to test the system performance. According to the off-line classification results, we directly selected a linear kernel function. A second experiment was carried out by training a linear SVM classifier with a higher amount of samples. For each manually labeled logo cutout corresponding to the training images we automatically created 162 samples by horizontal mirroring, geometric jittering and size varying (see Fig. 10). Accordingly, the number of training samples increased up to 386.532.




Fig. 10. Example of some additional training samples created for each manually labeled logo.

The reported results are depicted in Table III. The overall accuracy is 81.87% for the standard case and 88.23% for the system trained with multiple samples. These results proved the benefits of training the classifier with samples containing background pixels that do not correspond to the logo but to the vehicle structure. This can be explained by the fact that the sliding window approach used when selecting ROIs provides samples with background information.

TABLE II  
OFF-LINE VEHICLE LOGO RECOGNITION RATES

Manufacturer	Accuracy (%) rbfSVM / linSVM	Manufacturer	Accuracy (%) rbfSVM / linSVM	Manufacturer	Accuracy (%) rbfSVM / linSVM
Alfa Romeo	0.00% / 100.00%	Iveco	0.00% / 100.00%	Nissan	68.57% / 77.14%
Audi	100.00% / 100.00%	Jaguar	0.00% / 0.00%	Opel	98.84% / 96.51%
BMW	100.00% / 95.74%	Jeep	0.00% / 100.00%	Peugot	100.00% / 100.00%
Chevrolet	0.00% / 100.00%	Kia	10.00% / 90.00%	Renault	93.40% / 91.17%
Citroen	98.94% / 100.00%	Lancia	0.00% / 0.00%	Seat	98.44% / 100.00%
Fiat (old)	0.00% / 0.00%	Man	100.00% / 100.00%	Skoda	100.00% / 95.45%
Fiat (new)	81.58% / 81.58%	Mazda	100.00% / 100.00%	Toyota	74.42% / 81.40%
Ford	100.00% / 100.00%	Mercedes	100.00% / 100.00%	Volkswagen	99.17% / 98.33%
Honda	100.00% / 83.33%	Mini	0.00% / 0.00%	Volvo	78.72% / 85.11%

TABLE III  
SINGLE-FRAME VEHICLE LOGO RECOGNITION RATES

Manufacturer	Accuracy (%) standard / multiple	Manufacturer	Accuracy (%) standard / multiple	Manufacturer	Accuracy (%) standard / multiple
Alfa Romeo	0.00% / 0.00%	Iveco	0.00% / 100.00%	Nissan	45.71% / 65.71%
Audi	90.79% / 96.05%	Jaguar	0.00% / 0.00%	Opel	82.56% / 88.33%
BMW	80.85% / 85.17%	Jeep	50.00% / 66.67%	Peugot	94.74% / 97.89%
Chevrolet	40.00% / 80.00%	Kia	50.00% / 70.00%	Renault	76.42% / 83.02%
Citroen	95.74% / 100.00%	Lancia	0.00% / 0.00%	Seat	92.77% / 96.39%
Fiat (old)	0.00% / 0.00%	Man	0.00% / 42.86%	Skoda	68.18% / 90.91%
Fiat (new)	73.68% / 73.68%	Mazda	66.67% / 100.00%	Toyota	60.47% / 67.44%
Ford	87.27% / 93.64%	Mercedes	83.96% / 90.57%	Volkswagen	97.50% / 97.50%
Honda	83.33% / 83.33%	Mini	0.00% / 0.00%	Volvo	55.33% / 70.21%

We remark the fact that some of the mistakes are directly related with the nature of our vehicle logo localization stage. This is the case of Alfa Romeo, that reported an accuracy of 100% in the off-line experiment, with a linear SVM. However, when using the whole image, vehicle logo localization fails since the license plate of most of these vehicles is not centered. Some examples are depicted in Fig. 11. Other errors appeared due to the high distance between the license plate and the logo that provides too many ROIs with useless information. That is the case of most of the trucks (see Fig. 12).



Fig. 11. Some examples of vehicle logo localization errors for the Alfa Romeo case.



Fig. 12. Some examples of vehicle logo recognition errors in trucks due to the high distance between the license plate and the logo.

### C. On-line vehicle logo recognition results

In the last experiment, we tested the system in real conditions, i.e., considering the fact that one vehicle is captured by the camera obtaining between 2 and 7 useful images depending on the vehicle speed. We applied the structure shown in Fig. 3 and the FSM described in 4. The estimated logo that will correspond to the vehicle manufacturer is here obtained using the majority voting scheme applied over all the binary outputs provided by the classifier for all the available images for one specific vehicle. The classifier used in this experiment was the linear SVM trained with multiple samples. Table IV depicts the overall performance per each car manufacturer using the same number of samples of previous experiments (1/3) including consecutive frames corresponding to the same vehicle. Note that in this case the recognition rates are related with the number of detected vehicles. The final global performance we obtained was 92.59% (375/405). We can observe a correlation between the number of samples used for training and test, and the performance. Most of the vehicle manufacturers with more than 30 vehicles used for training provides recognition rates greater than 90%. Some exceptions are Volvo and Toyota, but still their recognition rates are around 70%. Fig. 13 depicts some of the vehicle recognition results.

## V. CONCLUSION

This paper presented a HOG/SVM framework for vehicle logo recognition using images captured by traffic cameras. Previous approaches were mainly based on the use of SIFT features, which are not useful in scenarios where the vehicle logos are not available with high resolution. The gradient distribution of the logos has been proved as an effective

TABLE IV  
ON-LINE VEHICLE LOGO RECOGNITION RATES

Manufacturer	Correct / Mistaken Accuracy (%)	Manufacturer	Correct / Mistaken Accuracy (%)	Manufacturer	Correct / Mistaken Accuracy (%)
Alfa Romeo	0/1 (0.00%)	Iveco	0/1 (0.00%)	Nissan	8/11 (72.73%)
Audi	21/21 (100.00%)	Jaguar	0/1 (0.00%)	Opel	30/33 (90.91%)
BMW	14/15 (93.33%)	Jeep	1/1 (100.00%)	Peugot	35/35 (100.00%)
Chevrolet	2/2 (100.00%)	Kia	3/3 (100.00%)	Renault	36/40 (90.00%)
Citroen	35/35 (100.00%)	Lancia	N/A	Seat	31/31 (100.00%)
Fiat (old)	0/1 (0.00%)	Man	2/2 (100.00%)	Skoda	8/9 (88.89%)
Fiat (new)	13/15 (86.67%)	Mazda	3/3 (100.00%)	Toyota	12/17 (70.59%)
Ford	41/41 (100.00%)	Mercedes	26/28 (92.86%)	Volkswagen	39/39 (100.00%)
Honda	5/5 (100.00%)	Mini	N/A	Volvo	10/15 (66.67%)



Fig. 13. Vehicle recognition examples.

descriptor for logo classification, whereas other features such as texture-based (LBP) do not provide better results. The sliding window technique provides a set of ROIs that are feed to the classifier. A majority voting approach reports good results, specially when multiple images are available for one specific vehicle. Results provided by the linear kernel are better than by means of radial basis functions. This result is reinforced by the higher speed of linear kernel when used as an on-line classifier. In addition, the increase of the number of training samples by creating new samples for each image of a vehicle, reported a considerable performance improvement.

The system is evaluated on a set of 3.579 vehicle images that belong to 27 distinctive vehicle manufacturers, captured by two different traffic cameras. The observed performance validates the proposed approach for being used in real traffic applications. Future works include the integration of the vehicle logo recognition approach in a hierarchical vehicle model recognition system.

## VI. ACKNOWLEDGMENTS

This work was supported by the Spanish Ministry of Science and Innovation under Research Grant ONDA-FP TRA2011-27712-C02-02.

## REFERENCES

- [1] C.-N. E. Anagnostopoulos, I. E. Anagnostopoulos, I. D. Psoroulas, V. Loumos, and E. Kayafas, "License plate recognition from still images and video sequences: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 9, pp. 377–391, 2008.
- [2] X. Z. Li, G. M. Zhang, J. Fang, J. A. Wu, and Z. M. Cui, "Vehicle color recognition using vector matching of template," in *International Symposium Electronic Commerce and Security*, 2010, pp. 189–193.
- [3] W. Feng, L. Zhi-fang, and M. Li-chun, "Color recognition of license plates using fuzzy logic and learning approach," *Jouranal of Optoelectronics.Laser*, vol. 20, pp. 84–88, 2009.
- [4] A. P. Psyllos, C.-N. E. Anagnostopoulos, and E. Kayafas, "Vehicle logo recognition using a sift-based enhanced matching scheme," *IEEE Trans. on Intell. Transp. Sys.*, vol. 11, pp. 322–328, 2010.
- [5] A. Psyllos, C. N. Anagnostopoulos, and E. Kayafas, "Vehicle model recognition from frontal view image measurements," *Comput. Stand. Interfaces*, vol. 33, no. 2, pp. 142–151, 2011.
- [6] W. Yunqiong, L. Zhifang, and Z. Fei, "A fast coarse-to-fine vehicle logo detection and recognition method," in *IEEE International Conference on Robotics and Biomimetics*, 2007.
- [7] A. Psyllos, C. N. Anagnostopoulos, and E. Kayafas, "Vehicle authentication from digital image measurements," in *16th IMEKO TC4 Symp., 13th Workshop ADC Model. Test.*, 2007, pp. 691–696.
- [8] T. Burkhard, A. J. Minich, and C. Li, "Vehicle logo recognition and classification: Feature descriptors vs. shape descriptors," Stanford University, EE368 Final Project, 2011.
- [9] W. Li, "A novel approach for vehicle-logo location based on edge detection and morphological filter," in *Second International Symposium on Electronic Commerce and Security*, 2009, pp. 343–345.
- [10] Z. Yong, Y. Hong-Yu, and Y. Zhi-Sheng, "A method of fast vehicle-logo location," *Journal of Si Chuan University: Natural Science Edition*, vol. 41, no. 6, pp. 1167–1171, 2004.
- [11] I. P. Alonso, D. F. Llorca, M. A. Sotelo, L. M. Bergasa, P. R. de Toro, J. Nuevo, M. Ocana, and M. A. G. Garrido, "Combination of feature extraction methods for svm pedestrian detection," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 2, pp. 292–307, 2007.
- [12] M. Gavilán, D. Balcones, O. Martín, D. F. Llorca, M. A. Sotelo, I. Parra, M. O. na, P. Aliseda, P. Yarza, and A. Amirola, "Adaptive road crack detection system by pavement classification," *Sensors*, vol. 11, pp. 9628–9657, 2011.
- [13] D. G. Lowe, "Object recognition from local scale-invariant features," in *Int. Conference on Computer Vision*, 1999, pp. 1150–1157.
- [14] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *International Conference on Computer Vision & Pattern Recognition*, 2005, pp. 886–893.
- [15] T. Ojala, M. Pietikäinen, and D. Harwood, "A comparative study of texture measures with classification based on featured distributions," *Pattern Recognition*, vol. 29, pp. 51–59, 1996.
- [16] C. C. Chang and C. J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011, software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.