

Vehicle Logo Recognition Using a SIFT-Based Enhanced Matching Scheme

Apostolos P. Psyllos, Christos-Nikolaos E. Anagnostopoulos, *Member, IEEE*, and Eleftherios Kayafas, *Member, IEEE*

Abstract—In this paper, a new algorithm for vehicle logo recognition on the basis of an enhanced scale-invariant feature transform (SIFT)-based feature-matching scheme is proposed. This algorithm is assessed on a set of 1200 logo images that belong to ten distinctive vehicle manufacturers. A series of experiments are conducted, splitting the 1200 images to a training set and a testing set, respectively. It is shown that the enhanced matching approach proposed in this paper boosts the recognition accuracy compared with the standard SIFT-based feature-matching method. The reported results indicate a high recognition rate in vehicle logos and a fast processing time, making it suitable for real-time applications.

Index Terms—Image matching, manufacturer recognition, vehicles.

I. INTRODUCTION—RELATED WORK

MANY vision-based intelligent transport systems for detecting, tracking, or recognizing vehicles in image sequences have been cited in the literature. Vehicle-type classification is a task that has been adequately addressed [1]–[6]. However, compared with vehicle-type classification, vehicle manufacturer recognition (VMR) is an area with fewer published systems and methods.

Image matching is a fundamental problem in computer vision that takes place in many image processing applications in a variety of fields, including image retrieval for security enforcement and robot navigation. A common approach is to locate characteristic image features (or keypoints) from the images and compare them through descriptors of these features. Earlier research for characteristic keypoints includes the Harris corner detector [10] and keypoints invariant to rotation and translation [11], [12]. The majority of the approaches concerning VMR described in the literature focus on scale-invariant feature transform (SIFT) and image matching using vehicle images properly segmented above the license plate area. SIFTs were introduced by Lowe [7], and they are invariant to rotation, translation, and scale variation between images and partially invariant to affine distortion, illumination variance, and noise.

Manuscript received May 20, 2009; revised October 20, 2009; accepted January 23, 2010. Date of publication February 25, 2010; date of current version May 25, 2010. The Associate Editor for this paper was M. M. Trivedi.

A. P. Psyllos and E. Kayafas are with the Department of Electrical and Computer Engineering, National Technical University of Athens, 15780 Athens, Greece (e-mail: psyllos@central.ntua.gr; kayafas@cs.ntua.gr).

C.-N. E. Anagnostopoulos is with the Department of Cultural Technology and Communication, University of the Aegean, 81100 Mytilene, Greece (e-mail: canag@ct.aegean.gr).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TITS.2010.2042714

Research related to affine invariant features has been published by Brown and Lowe [8] and Mikolajczyk and Schmidt [9]. Dlagnekov and Belongie [13] utilized SIFT features. In their work, rear-view vehicle images were used, attaining 89.5% recognition accuracy. The vehicle manufacturer and model were treated as a single class and recognized simultaneously. However, it was reported that the system does not have real-time performance, and no results for recognition speed were published. Èonos [14] deals with a vehicle model recognition problem of frontal-view images, proposing a SIFT-based descriptor for feature extraction with 90% recognition rate. The drawback of the latter approach is the extensive computational cost—about 10 s is reported.

Petrovic and Cootes [15] presented an interesting approach for VMR from frontal-view vehicle images that displayed a 93% recognition accuracy. The vehicle manufacture and model were treated as a single class and recognized simultaneously, and no results for recognition speed were reported. A comparative knowledge-acquisition system appears in [16], consisting of several object recognition modules that represent a car image viewed from the rear, such as a window, tail lights, and so on, based on color recognition. This approach has the drawback of being sensitive to lighting conditions.

Csurka *et al.* [17] proposed a method of “bag-of-keypoints” that is based on vector quantization of affine invariant descriptors of image patches. They present results for simultaneously classifying seven semantic visual categories. Opelt *et al.* [18] used a boundary-fragment model for object category classification (e.g., airplanes versus animals) and based only in the object’s boundary recovered by an edge detector. Leibe *et al.* [19] presented an implicit shape model for combined object categorization and segmentation. Maji and Malik [20] used a probabilistic Hough transform and a support vector machine classifier for object categorization.

The VMR problem is addressed through vehicle logo recognition (VLR), and the recognition task requires the successful extraction of the small logo area from the original vehicle image. Usually, this process involves a license plate location (LPL) module, followed by coarse-to-fine methods to identify the logo area using symmetry and/or edge statistics in the image. Then, logo recognition is performed through either neural networks [21]–[23] or template matching. Wang *et al.* [24] presented a method for logo recognition using template matching and edge-orientation histograms with good results.

In this paper, a reliable VLR schema is proposed, which detects and extracts the points of interest (SIFT keypoints) in a logo image and performs content-based image retrieval from a

logo-image database. SIFT keypoints are invariant to scale and rotation and even partially invariant to illumination differences. Additionally, if the object is partially occluded, then the keypoints obtained from the observed part of the object are capable of recognizing the object, which is particularly useful in the case of vehicle recognition. SIFT recognition might fail in cases with poor illumination conditions (excessive, nonuniform, or poor lighting, shadows), weather, dirt, high occlusion from other objects, and camera wide (perspective) view angle (low contrast), which results in poorly detected features. The logo has been detected from a frontal-view vehicle image using an LPL module to locate the position of vehicle license plates. A symmetry axis module is employed, optionally, to verify and centrally adjust the license plate, if detected in a nonsymmetric position, followed by a special image processing technique called phase congruency calculation, which detects the logo boundaries.

The contribution of this paper mainly lies in the effective use of many different views of the same model database features to describe a detected query feature, increasing the possibilities for correct recognition and making the recognition process more robust in outdoor conditions. Another novelty of this paper is the dynamic size of the feature database. After a successful recognition, the matched features are inserted into the detected match database model, enhancing the views for a model and further improving recognition robustness. In addition, the use of relatively simple image processing techniques for license plate recognition (LPR), vehicle mask, and logo-image detection and segmentation results in fast overall recognition time, thus making it suitable for real-time applications.

The rest of this paper is organized as follows: In Section II, vehicle logo detection and segmentation is described in brief, and in Section III, the analysis of the proposed logo-recognition procedure is presented. In Section IV, the experimental procedure is described, and the recognition performance results are presented, and Section V concludes this paper and refers to future extensions of this research.

II. VEHICLE LOGO DETECTION

A captured image is converted to greyscale and fed to the LPL module, and the result is a rectangular area (or areas) that includes the candidate plate (or plates) and the coordinates of the upper left and lower right corners, respectively. Using the coordinates of the detected plate, the vehicle mask—an area containing the headlights, the radiator grille, and the manufacturer logo—can be segmented. The size of the vehicle mask depends on the dimensions of the license plate, and this segmentation method can be properly applied in cases with a close or a distant view from the vehicle. However, this process is not valid for vehicles having nonsymmetric frontal LPL. Therefore, to correct such inconsistencies and to extract a symmetric vehicle mask, a symmetry axis identification module is used based on the gradient orientation histogram method [22], [23]. The highest peaks of orientation histogram correspond to the angles of possible symmetry axes.

The whole process is quite fast (about 100 ms on average) and is described analytically in [25], where it is reported that

the above modules, when applied in a large image set of 1334 images captured in various illumination conditions, achieve a success rate of 96.5% for plate segmentation and 89.1% success in the recognition of the entire plate content.

It should be noted that in the case where the vehicle mask has not been detected and segmented correctly, this would affect the next step of logo detection, and therefore, logo recognition will eventually fail. Such failure could also be attributed to poor illumination conditions, weather, dirt, partial occlusion from other objects, and camera wide view angle.

To locate the manufacturer logo contained in the image mask, a method based on phase congruency feature map (PCFM) calculation (introduced by Kovese [26]) is implemented and provides an illumination and contrast invariant measure of feature significance. The horizontal projection of PCFM is a unique and representative measurement for each of the samples used and provides useful information concerning the characteristic parts of the vehicle mask, such as headlights, radiator grille, and manufacturer logo. PCFM derivative values that exceed a proper threshold have a higher possibility to contain the boundaries of salient features from the segmented frontal area. The segmentation algorithm searches near the middle of this area to extract the manufacturer logo. PCFM calculation is the slowest step of the logo-detection process (~ 250 ms).

III. VEHICLE LOGO RECOGNITION

A. Feature Detection and Description

SIFT [7] is the state of the art in the field of image recognition and is used in a wide range of content-based image-retrieval applications. It exploits the idea of replacing images by a set of scale and orientation-invariant keypoint (feature) descriptors using gradient orientation histograms. The invariant features are detected and extracted, exploring the scale-space structure of an image [27]. Features are localized and filtered, keeping only those that are likely to remain stable over affine transformations, have adequate contrast, and are not along edges. The presence of keypoints not lying in the edges and having adequate contrast is ensured with the appropriate selection of two parameters h and ε in the Difference of Gaussians (DoG) function, as described analytically in [7]. The former parameter defines the appropriate threshold to eliminate low contrast keypoints. The latter parameter gives the ratio between the largest and smallest magnitude eigenvalues of a matrix containing curvature information of a DoG function. In this paper, these parameters are set to the values of $h = 0.01$ and $\varepsilon = 28$. The keypoint descriptor is created by sampling the magnitudes and orientations of the image gradient in a patch around the detected feature, resulting in a 128-D vector of direction histograms.

B. Merged Features From Multiple Images (MFM)

To enhance the recognition process, a group of images describing similar scenes of the same object (i.e., a manufacturer logo) are employed instead of using a single image. This process detects all SIFT-based features from a set of N images, and one is selected as a reference. The reference image is chosen by an expert as the best representing the set (best projected

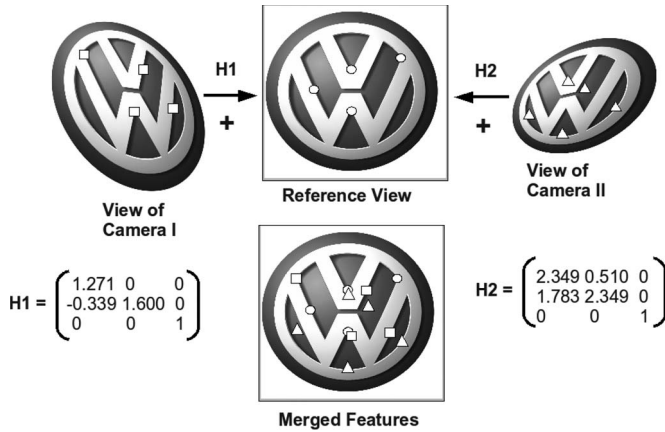


Fig. 1. Features from different camera views being merged into a reference view using homography matrices $H1$ and $H2$.

and centered, uniformly illuminated, etc., if possible). The remaining $N - 1$ plain images of the set are transformed to the coordinate system of the reference image calculating their homographies using the random sample consensus (RANSAC) method [28].

In brief, the exact homography is calculated for four randomly selected feature pairs from the reference and plain images, respectively. The inliers are estimated based on feature locations within a fixed error margin. The process is repeated until a satisfactory number of inliers is found (in this paper, the criterion is at least 40% inliers). At the end of this merging process, an object database with fused features is formed, where the descriptors are taken from each of the $N - 1$ plain images, and their locations are transformed to the reference image coordinates (see Fig. 1). Increasing the features related with a database model (even if new keypoints are projected to the same position in the reference image) intuitively gives a better approximation and description of that model, which in turn increases the possibilities of a correct descriptor match.

This multiple-image matching procedure will be referred to as merged feature matching (MFM), and the single image matching will be referred to as single feature matching (SFM). It should be noted that after a successful recognition, the matching keypoints are inserted in the detected match database model using the same procedure with MFM, further enriching the views for a model and thus improving recognition robustness.

C. Feature Matching

For each feature i in the query image, the descriptor is used to search for its nearest neighbor (NN) matches among all the features from all the images j contained in a feature database. The NNs are selected by satisfying a minimum L2 Euclidean distance threshold γ criterion for the descriptor vectors Q_i and DB_{ij} for the query and database image j , respectively, i.e.,

$$NN_i = \text{cardinality} \{ \arg (\|Q_i - DB_{ij}\| < \gamma) \}. \quad (1)$$

Therefore, the number of NNs in the database for each feature (keypoint) (NN_i) depends on threshold (γ) selection. In

$$\tau = \frac{NN_i}{F_Q \cdot F_{DB}} \quad (2)$$

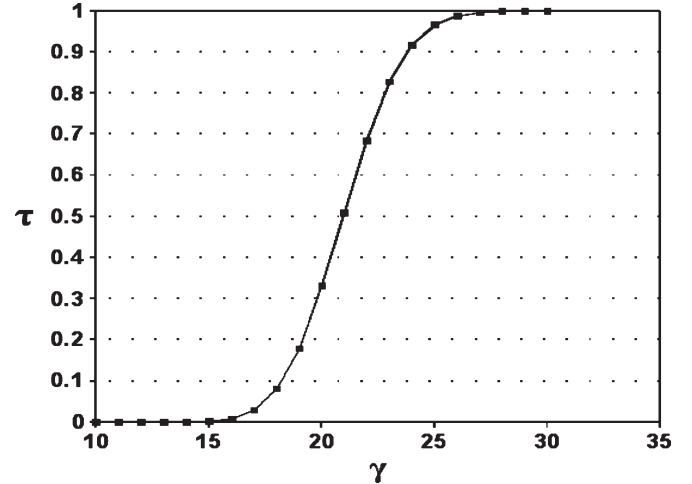


Fig. 2. Reduced NN parameter (τ) plotted versus distance threshold (γ).

a reduced NN parameter τ is used, where NN_i is calculated from (1) for each feature i in the query image. F_Q is the number of features detected in the query image, and F_{DB} is the total number of features in the database.

In Fig. 2, the reduced NN parameter (τ) is plotted against threshold (γ), and an optimum threshold for NN matching is selected as the turning point of this plot. A fast KD-Tree data structure is created using the feature descriptor database with search time complexity that is almost linear with database size [29], [30].

D. Feature Clustering and Geometric Validation

This step determines which of the NN matches found in the previous step are reasonable, verifying whether the query image represents a logo image stored in the database. NN descriptors are clustered using a generalized Hough transform (GHT) [31], [32] to define the parameters for a similarity transformation between the query and database features. GHT identifies clusters of features with a consistent similarity transformation (position, orientation, and scale) by using each feature to vote for all the logo images that are consistent with the feature. When clusters of features are found to vote for the same pose of a vehicle logo, the probability of a correct match is considerably higher than for any other feature. The database logo image with the highest number of votes is considered to be the most similar with the query image. Next, the coordinates of the keypoints in the query and the matched database image are checked for geometrical consistency. The RANSAC method (see the MFM procedure in Section III-B) is applied for all GHT clusters found, and the affine transformation with the maximum number of inliers (or minimum number of outliers) is estimated. Lowe [7] reports that it is possible to have reliable recognition with as few as three feature correct matches.

IV. EXPERIMENTAL PART

A. Database and Query Set

The proposed method is assessed on a set of 1200 frontal area images that have been successfully segmented using

TABLE I
LOGO DETECTION RATE

Manufacturer	True	False
Alfa Romeo	73	7
Audi	77	3
BMW	76	4
Citroen	79	1
Fiat	79	1
Peugeot	77	3
Renault	79	1
Seat	79	1
Toyota	78	2
VW	79	1
Average (%)	97	3

our previous work described in [22], [23], and [25] from the Medialab LPR Database [33]. The segmented logo image dimensions are scaled to 100×100 pixels. From the total of 1200 logo images, 400 have been selected to form the logo database. These images correspond to ten classes of selected vehicle manufacturers, as seen in Table I. Each class contains 40 images, and for each one, the keypoints have been detected, and the descriptors have been stored. Then, a reference image was selectively chosen by an expert, and the remaining 39 samples were registered according to that reference view using the homography calculated by the RANSAC algorithm. Only areas belonging to the common parts of the images were selected, and the keypoint descriptors were rereferenced to the new position, scale, and orientation. This way, the number of keypoints for every manufacturer logo is substantially increased by merging the keypoints, thus making the recognition process more robust in illumination conditions. Finally, a database containing merged keypoints for every manufacturer is created. The query set for our experiments is the remaining 800 frontal area images. For those images, the logo has been detected and segmented automatically, and for every image, the keypoints and their respective descriptors have been calculated following the SIFT procedure, as described in Section II.

B. Logo Detection

The license plates were located, and using their position and size, the vehicle masks were segmented and forwarded to the module of symmetry axis locator and PCFM calculation. The accuracy of vehicle logo segmentation was tested experimentally, validating automatic logo segmentation with a human expert. The results are shown in Table I, where the performance of the logo-segmentation process is 97% on average. The logo-detection process requires about 380 ms.

C. Logo Recognition

Query image descriptors were matched against database descriptors using a properly tuned threshold value, which corresponds to the turning point of NN number found versus threshold. For Euclidean NN search and match, a KD-Tree data structure was employed, having 128 dimensions and 40 000 nodes (400 database images \times approximately

100 keypoints per image). Table II shows the confusion matrix of the verification process, where the main diagonal displays the matched Euclidean distance NN percentage. It is evident that the matched NN keypoints are a fraction of the total NN keypoints detected per logo image. The database image whose descriptors get the maximum number of votes in Hough transformation array is considered to be the one most similar to the query image and is then checked using RANSAC for geometric consistency. The number of GHT-clustered keypoints is on average 63% of the total keypoints detected and contains on average 32% of the inliers (see Table III). The recognition statistics are given analytically in Table IV, where the combined logo-detection and logo-recognition success rates are 88% and 91% for SFM and MFM, respectively.

An example of invariant feature matches between query and database images is shown in Fig. 3. The average recognition speed is about 850 ms for SFM and 1020 ms for multiple feature matching (MFM). It should be noted that, usually, after the filtering process, MFM produces approximately 15%–30% additional keypoints compared with SFM.

The total recognition performance is evaluated, using 1-Precision versus Recall, as introduced in [34], for a range of experimental conditions, varying the threshold value (γ) used for NN matching. The results from these experiments using the single image matching and MFM scheme are plotted together for comparison in Fig. 4.

V. CONCLUSION

This paper has provided a description of a practical application for vehicle make recognition based on the well-established image-recognition method of SIFT. The enhanced logo-recognition system (MFM) demonstrated good performance, yielding 91% overall recognition success rate (97% logo segmentation \times 94% logo recognition), when applied to a database of already-segmented logo images. The SFM system has shown an 88% overall recognition success rate (97% logo segmentation \times 91% logo recognition). The proposed scheme also performs better compared with our previous work, in which a probabilistic neural network was used [22], [23] with an 84% overall recognition rate. The VLR speed is rather fast, with a combined detection and recognition time of about 1400 ms (380 ms + 1020 ms) for MFM, which is suitable for real-time applications. For SFM, the total logo detection and recognition time is 1230 ms (380 ms + 850 ms).

To further boost performance and robustness, there has been ongoing research to combine logo recognition with vehicle model recognition matching the model of a logo-recognized vehicle manufacturer. A short-time video that demonstrates some preliminary results in real-time execution can be downloaded from [35]. Moreover, another extension of this paper might be the possibility of dealing with a wider range of viewpoints or 3-D recognition, as well as recognition of more complex scenes, including many vehicles and under a wider variety of illumination conditions.

As an implementation example, a parking lot entrance wireless IP camera captures vehicle frontal view images when a change is detected (a vehicle moves in/out). The vehicle

TABLE II
AVERAGED QUERY AND DATABASE-MATCHED KEYPOINTS

Query	Database Matches (% of total matched NN keypoints per manufacturer)									
	Alfa Romeo	Audi	BMW	Citroen	Fiat	Peugeot	Renault	Seat	Toyota	VW
Alfa Romeo	39	5	3	8	9	8	3	11	5	9
Audi	9	45	1	2	4	15	2	1	3	18
BMW	18	14	32	0	7	4	7	11	0	7
Citroen	8	8	0	27	2	19	4	13	6	13
Fiat	16	15	3	2	23	15	10	3	3	10
Peugeot	10	8	7	10	11	27	9	1	8	9
Renault	11	3	1	12	14	6	30	9	4	10
Seat	10	4	2	11	12	8	9	24	7	13
Toyota	5	5	1	9	10	8	3	2	41	16
VW	9	5	7	10	9	4	6	7	12	31

TABLE III
KP SIZE DETECTED VERSUS HOUGH CLUSTER SIZE AND RANSAC INLIERS

Query Image	NN	Hough Cluster Size	RANSAC Inliers
Alfa Romeo	51	42	18
Audi	43	30	14
BMW	30	24	9
Citroen	23	17	7
Fiat	39	22	15
Peugeot	45	20	14
Renault	47	40	12
Seat	56	16	16
Toyota	37	26	12
VW	22	12	10
Total	393	249	127
Average on Total NN (%)	-	63%	32%

TABLE IV
LOGO RECOGNITION RATE AND COMBINED DETECTION RECOGNITION RATE

Manufacturer	SFM Recognition Rate		MFM Recognition Rate		Combined Detection & Recognition SFM (percentage %)		Combined Detection & Recognition MFM (percentage %)	
	True	False	True	False	True	False	True	False
Alfa Romeo	72	8	76	4	82%	18%	87%	13%
Audi	68	12	70	10	81%	19%	84%	16%
BMW	73	7	77	3	87%	13%	91%	9%
Citroen	78	2	80	0	96%	4%	99%	1%
Fiat	68	12	68	12	84%	16%	84%	16%
Peugeot	72	8	75	5	87%	13%	90%	10%
Renault	75	5	76	4	92%	8%	94%	6%
Seat	70	10	73	7	86%	14%	90%	10%
Toyota	76	4	78	2	93%	7%	95%	5%
VW	79	1	79	1	97%	3%	98%	2%
Average (%)	731/800	69/800	752/800	48/800	88%	12%	91%	9%
	(91%)	(9%)	(94%)	(6%)				

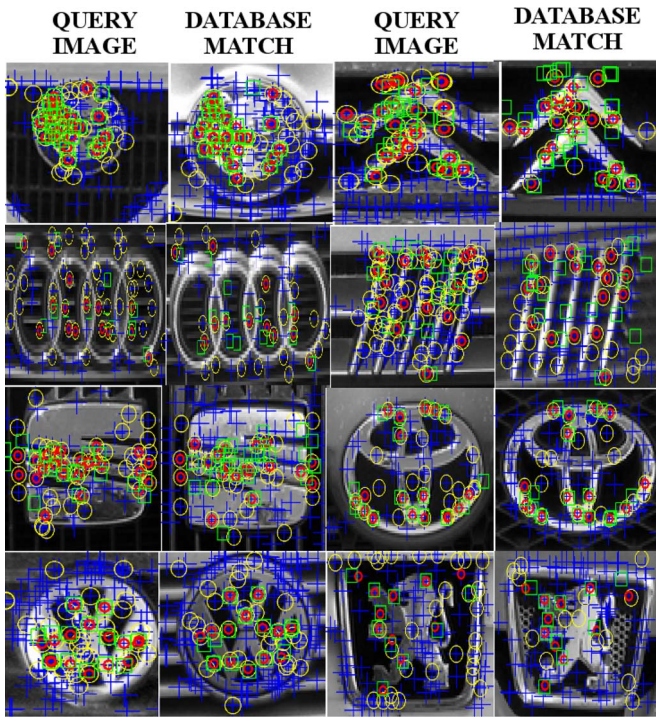


Fig. 3. Keypoint matching and geometric validation for query and database images. Total detected keypoints are shown in blue, NN keypoints are in yellow, keypoints belonging to the maximum GHT cluster are shown in red, and affine transform RANSAC inliers are surrounded by a green rectangle.

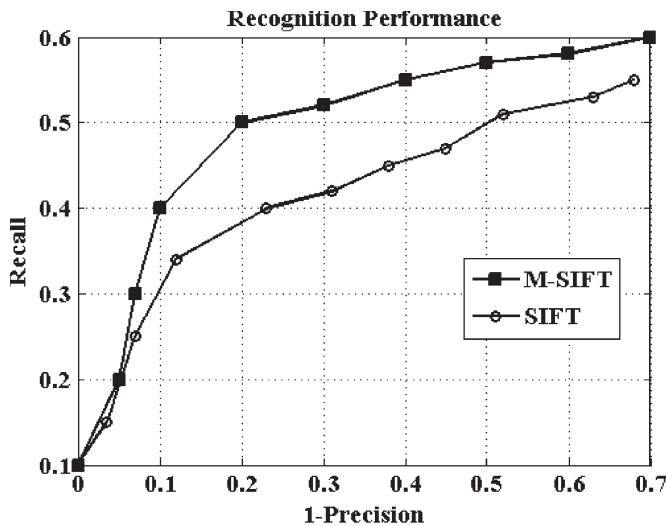


Fig. 4. Logo recognition performance for MFM and standard SFM, respectively. 1-Precision and Recall are calculated by varying the threshold for NN matching.

logo-detection module is then activated, which in turn forwards the detected logo to the logo-recognition module, using a dynamic-updated central logo database. A module for vehicle model recognition can be used together with a color-recognition module to obtain more information about the vehicle. Our vision in the case of practical application is the use of an embedded system where an optimized code can be executed in parallel using the graphics processing unit (GPU) architecture for even faster response. An embedded Linux is also a potential candidate for a low-cost implementation of a smart camera, which includes the whole recognition process.

REFERENCES

- [1] M. Weber, M. Welling, and P. Perona, *Unsupervised Learning of Models for Recognition*. Berlin, Germany: Springer-Verlag, 2000, pp. 18–32.
- [2] T. Kato, Y. Ninomiya, and I. Masaki, “Preceding vehicle recognition based on learning from sample images,” *IEEE Trans. Intell. Transp. Syst.*, vol. 3, no. 4, pp. 252–260, Dec. 2002.
- [3] A. H. S. Lai and N. H. C. Yung, “Vehicle-type identification through automated virtual loop assignment and block-based direction-biased motion estimation,” *IEEE Trans. Intell. Transp. Syst.*, vol. 1, no. 2, pp. 86–97, Jun. 2000.
- [4] A. H. S. Lai, G. S. K. Fung, and N. H. C. Yung, “Vehicle type classification from visual-based dimension estimation,” in *Proc. IEEE Intell. Transp. Syst. Conf.*, Oakland, CA, 2001, pp. 201–206.
- [5] S. Gupte, O. Masoud, R. F. K. Martin, and N. P. Papanikolopoulos, “Detection and classification of vehicles,” *IEEE Trans. Intell. Transp. Syst.*, vol. 3, no. 1, pp. 37–47, Mar. 2002.
- [6] J. W. Hsieh, S. H. Yu, Y. S. Chen, and W. F. Hu, “Automatic traffic surveillance system for vehicle tracking and classification,” *IEEE Trans. Intell. Transp. Syst.*, vol. 7, no. 2, pp. 175–187, Jun. 2006.
- [7] D. Lowe, “Object recognition from local scale-invariant features,” in *Proc. Int. Conf. Comput. Vis.*, Corfu, Greece, Sep. 1999, pp. 1150–1157.
- [8] M. Brown and D. Lowe, “Invariant features from interest point groups,” in *Proc. BMVC*, Cardiff, U.K., Sep. 2002, pp. 656–665.
- [9] K. Mikolajczyk and C. Schmid, “An affine invariant interest point detector,” in *Proc. Eur. Conf. Comput. Vis.*, 2002, vol. 1, pp. 128–142.
- [10] C. Harris and M. Stephens, “A combined corner and edge detector,” in *Proc. 4th Alvey Vis. Conf.*, Manchester, U.K., 1988, pp. 147–151.
- [11] S. Siggelkow, “Feature histograms for content-based image retrieval,” Ph.D. dissertation, Albert-Ludwigs-University Freiburg, Freiburg im Breisgau, Germany, Dec. 2002.
- [12] H. Schulz-Mirbach, “Invariant features for gray scale images,” in *Proc. DAGM Symp.*, 1995, pp. 1–14.
- [13] L. Dlagnekov and S. Belongie, “Recognizing cars,” University of California at San Diego, La Jolla, CA, Tech. Rep. CS2005-0833, 2005.
- [14] M. Čonos, “Recognition of vehicle make from a frontal view,” M.S. thesis, Fac. Elect. Eng., Czech Tech. Univ., Prague, Czech Republic, 2006.
- [15] V. S. Petrovic and T. F. Cootes, “Analysis of features for rigid structure vehicle type recognition,” in *Proc. Brit. Mach. Vis. Conf.*, 2004, vol. 2, pp. 587–596.
- [16] H. Maemoto, S. Okuma, and Y. Yano, “Parametric vehicle recognition using knowledge acquisition system,” in *Proc. IEEE Int. Conf. Syst., Man, Cybern.*, 2004, vol. 4, pp. 3982–3987.
- [17] G. Csürka, C. Bray, C. Dance, and L. Fan, “Visual categorization with bags of keypoints,” in *Proc. ECCV Workshop Stat. Learn. Comput. Vis.*, 2004, pp. 59–74.
- [18] A. Opelt, A. Pinz, and A. Zisserman, “A boundary-fragment-model for object detection,” in *Proc. ECCV*, 2006, vol. II, pp. 575–588.
- [19] B. Leibe, A. Leonardis, and B. Schiele, “An implicit shape model for combined object categorization and segmentation,” in *Towards Category-Level Object Recognition*. Berlin, Germany: Springer-Verlag, 2006, pp. 496–510.
- [20] S. Maji and J. Malik, “Object detection using a max-margin Hough transform,” in *Proc. CVPR*, 2009, pp. 1038–1045.
- [21] H. J. Lee, “Neural network approach to identify model of vehicles,” in *ISNN 2006*. Berlin, Germany: Springer-Verlag, 2006, pp. 66–72.
- [22] A. Psyllos, C. N. Anagnostopoulos, E. Kayafas, and V. Loumos, “Image processing and artificial neural networks for vehicle make and model recognition,” in *Proc. 10th Int. Conf. Appl. Adv. Technol. Transp.*, Athens, Greece, May 28–30, 2008.
- [23] A. Psyllos, C. N. Anagnostopoulos, and E. Kayafas, “Vehicle authentication from digital image measurements,” in *Proc. 16th IMEKO TC4 Symp., 13th Workshop ADC Model. Test.*, Florence, Italy, Sep. 22–24, 2008.
- [24] Y. Wang, Z. Liu, and F. Xiao, “A fast coarse-to-fine vehicle logo detection and recognition method,” in *Proc. IEEE Int. Conf. Robot. Biomimetics*, Sanya, China, Dec. 2007, pp. 691–696.
- [25] C. Anagnostopoulos, I. Anagnostopoulos, V. Loumos, and E. Kayafas, “A license plate-recognition algorithm for intelligent transportation system applications,” *IEEE Trans. Intell. Transp. Syst.*, no. 3, pp. 377–392, Sep. 2006.
- [26] P. D. Kovsi, “Image features from phase congruency,” in *Videre: A Journal of Computer Vision Research*. Cambridge, MA: MIT Press, 1999, pp. 1–27.
- [27] T. Lindeberg, “Scale-space theory: A basic tool for analyzing structures at different scales,” *J. Appl. Stat.*, vol. 21, no. 1/2, pp. 224–270, 1994.
- [28] M. Fischler and R. Bolles, “Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography,” *Commun. ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981.

- [29] S. Arya and D. Mount, "Approximate nearest neighbor searching," in *Proc. 4th Annu. ACM-SIAM SODA*, 1993, pp. 271–280.
- [30] J. Friedman, J. Bentley, and R. Finkel, "An algorithm for finding best matches in logarithmic expected time," *ACM Trans. Math. Softw.*, vol. 3, no. 3, pp. 209–226, Sep. 1977.
- [31] D. Ballard, "Generalizing the Hough transform to detect arbitrary patterns," *Pattern Recognit.*, vol. 3, no. 2, pp. 111–122, 1981.
- [32] E. Grimson, *Object Recognition by Computer: The Role of Geometric Constraints*. Cambridge, MA: MIT Press, 1990.
- [33] Images Database, Feb. 2009. [Online]. Available: <http://www.medialab.ntua.gr/research/LPRdatabase.html>
- [34] Y. Ke and R. Sukthankar, "PCA-SIFT: A more distinctive representation for local image descriptors," in *Proc. Conf. Comput. Vis. Pattern Recog.*, 2004, vol. 2, pp. 506–513.
- [35] Video Demo, Oct. 2009. [Online]. Available: <http://www.medialab.ntua.gr/research/LPRdatabase/sdemo.mpg>



Apostolos P. Psyllos was born in Athens, Greece, in 1966. He received the Diploma and Ph.D. degrees in chemical engineering from the National Technical University of Athens (NTUA) in 1988 and 1993, respectively.

Since 2006, he has been with the Department of Electrical and Computer Engineering, NTUA, where he started as a Ph.D. student with the Multimedia Technology Laboratory. His research interests are image and signal processing, computer vision, remote sensing, and artificial intelligence. He has

published three papers in conferences in the above subjects.

Dr. Psyllos is a member of the Greek Chamber of Engineers and the Greek Informatics Society.



Christos-Nikolaos E. Anagnostopoulos (M'04) was born in Athens, Greece, in 1975. He received the Diploma in mechanical engineering and the Ph.D. degree from the National Technical University of Athens in 1998 and 2002, respectively.

Since 2008, he has been an Assistant Professor with the Department of Cultural Technology and Communication, University of the Aegean, Mytilene, Greece. His research interests are image processing, computer vision, neural networks, and artificial intelligence. He has published more than 100 papers

in journals and conferences in the above subjects, as well as other related fields in informatics.

Dr. Anagnostopoulos is a member of the Greek Chamber of Engineers.



Eleftherios Kayafas (M'97) received the B.Sc. degree from Athens University, Athens, Greece, in 1970 and the M.Sc. and Ph.D. degrees in electrical engineering from the University of Salford, Salford, U.K., in 1975 and 1978, respectively.

From 1973 to 1974, he was with the Civil Aviation Service. From 1978 to 1979, he was the Hellenic Aerospace Industry. Since 1979, he has been with the National Technical University of Athens, where he was a Lecturer with the Electrical Engineering Department, an Assistant Professor in 1987, an As-

sociate Professor in 1992, and finally a Professor of applied electronics in 1996. His research interests are applied electronics, multimedia applications, and multimedia communication systems. He has published more than 140 papers in journals and conferences in the above subjects.

Prof. Kayafas is a member of the Greek Chamber of Engineers and a member of the TC-4 Committee of the International Measurement Confederation.