2012 IEEE International Conference on Vehicular Electronics and Safety
July 24-27, 2012. Istanbul, Turkey

M-SIFT: A new method for Vehicle Logo Recognition

Apostolos Psyllos, Christos-Nikolaos Anagnostopoulos, Eleftherios Kayafas

Abstract - In this paper, a new algorithm for Vehicle Logo Recognition is proposed, on the basis of an enhanced Scale Invariant Feature Transform (Merge-SIFT or M-SIFT). The algorithm is assessed on a set of 1500 logo images that belong to 10 distinctive vehicle manufacturers. A series of experiments are conducted, splitting the 1500 images to a training set (database) and to a testing set (query). It is shown that the M-SIFT approach, which is proposed in this paper, boosts the recognition accuracy compared to the standard SIFT method. The reported results indicate an average of 94.6% true recognition rate in vehicle logos, while the processing time remains low (~0.8sec).

I. Introduction

Many vision-based Intelligent Transport Systems for detecting, tracking or recognizing vehicles in image sequences are referred in the literature. Vehicle type classification is a task that has been adequately addressed in the literature [1-6]. However, Vehicle Manufacturer Recognition (VMR) is a crucial subject, with relatively limited research reported in the field. This is due to the fact that manufacturer recognition is an inherently hard problem due to the wide variety in the appearance of vehicles. This task becomes even more difficult since image acquisition is performed in outdoor environment where illumination is ambient.

Scale Invariant Feature Transforms (SIFT) were introduced by Lowe [7-9] and they are invariant to rotation, translation and scale variation between images and partially invariant to affine distortion, illumination variance and noise. Research related to fully invariant features has been published by Brown and Lowe [10], Mikolajczyk and Schmidt [11]. Image matching is a fundamental problem in computer vision which occurs in many computer vision applications regarding a variety of fields, including image retrieval for security enforcement and robot navigation. A common approach to accurate image matching is known as the invariant keypoint or the extraction of the point of interest from the images compared. It involves identifying the points that can be reliably extracted from different images of the same object or the same category of objects. Earlier research into invariant keypoints includes the Harris corner detector [12] and the keypoints invariant to rotation and translation [13], [14]. Dlangenkov and Belongie [15] utilized Scale Invariant Feature Transform (SIFT) features, using a vehicle database of rear-view vehicle images. Čonos [16] deals with a vehicle model recognition problem of frontal view images, proposing a SIFT-based descriptor for feature extraction. The drawback of the latter approach is the extensive computational cost reported in the respective paper. Petrovic and Cootes [17] presented an interesting approach based on edge gradient and match refinement algorithm for vehicle model recognition and verification. A comparative knowledge acquisition system, consisting of several object recognition modules for the rear view of vehicles appears in Maemoto et al. [18]. The VMR problem is resolved through vehicle logo recognition (VLR) and the recognition task requires the successful extraction of the small logo area from the original vehicle image. Usually, the process involves a license plate location 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 either performed through neural networks [19-21] or template matching. Wang et al. [23] presented a method for logo recognition using template matching and edge orientation histograms with good results.

In this paper, a SIFT-based Vehicle Logo Recognition schema is described, whose aim is to obtain reliable Vehicle Make Recognition from the logo image. An enhanced SIFT matching module detects and extracts the points of interest (keypoints) in the query logo image, describes them and matches them with keypoints stored in a logo image database. The whole process is fine-tuned by clustering the matched keypoints, using Generalised Hough Transform [24], and geometric verification by means of an affine transformation. The proposed method is assessed on a database containing 1500 logo images, which were processed using our previous work described in [21], [22], [25] using the Medialab Vehicle Database [26].

II. Vehicle Logo Recognition

SIFT is the state-of-the-art in the field of image recognition and is used in a wide range of image retrieval

Manuscript received June 1, 2012.

A. Psyllos is with NTUA, Sc. of Applied Mathematical and Physical Science, Athens, Greece (email: psyllos@central.ntua.gr).

C. N. Anagnostopoulos is with University of the Aegean, Mytilene, Greece. (Corresponding author: +302251036624; fax: +302251036637; e-mail: canag@ct.aegean.gr

E. Kayafas is with NTUA, Sc. of Electrical & Computer Engineering, Athens, Greece (email: kayafas@cs.ntua.gr).

applications. It is based on the idea of representing images by a set of invariant keypoint descriptors using gradient orientation histograms. The invariant features (points of interest), or just keypoints, are distinguished as local peaks in the image scale-space representation and filtered to preserve only those that are likely to remain stable over transformations. For each invariant feature a SIFT keypoint descriptor is calculated having the following properties: a) smooth changes in location, orientation and scale do not cause radical changes in the feature vector, b) it is fairly compact, expressing an orientation histogram from a patch of pixels in the neighborhood of the invariant feature, using a 128 dimensional vector, c) it is resilient to affine deformations, such as those caused by weak perspective effects (for small angles of deviation), and thus it could be efficient for logo recognition in non-controlled conditions. In this paper, in the initial stage of computation, scale and orientation invariant features are detected and extracted from the query image, exploring the scale-space structure of the image as described in [7] and [27]. In the next sections, our contribution in developing an enhanced Scale Invariant Feature Transform (Merge-SIFT or M-SIFT) is described, pointing out the feature matching and hypothesis validation procedures in the candidate points.

A. Feature Matching

For each feature i in the query image, the descriptor is used to find its Nearest-Neighbor (NN) matches among all the features stored from all the images in a database called "NN database" (see Figure 1). The nearest neighbors are identified minimizing an L2 Euclidean metric:

$$\|Q_i - D_i(j)\| = \sqrt{\sum_{k=1}^{N}(q_k - d_k)^2} \qquad (1)$$

where Q_i (i.e. $q_1,q_2...,q_N$) is the ith descriptor for the query image, $D_i(j)$ (i.e $D_1,D_2...,D_N$) is the ith descriptor for image j and N=128. The total number of nearest neighbors is calculated from Equation (1), where cardinality is the number of elements in a set and γ is an appropriate threshold value.

$$NN_i = cardinality(\arg\{\|Q_i - D_i(j)\| < \gamma\}) \qquad (2)$$

Therefore, the number of NNs in the database for each feature (keypoint) depends on the threshold selection. In Figure 2, the parameter τ is plotted versus threshold (γ), where τ is described by the equation (3):

$$\tau = \frac{NN_i}{KP_q \cdot KP_{db}} \qquad (3)$$

where NN_i is the number of nearest neighbors found from Equation (2) for each feature i in the query image, KP_q is the number of keypoints in the query image and KP_{db} is the total number of keypoints in NN database.

B. Hypothesis Validation

1) Feature Clustering

A typical logo image contains approximately 100 features, numerous of which are not important to the matching procedure. This is due to the fact that some features correspond to ambient illumination reflections, shadows or noise. To maximize the performance of VLR for small, ill-captured or noised logos, logo identification should occur with the fewest possible feature matches. It was found that it is possible to have reliable recognition with as few as 3 feature matches. This step determines the reasonable NN matches found in the previous step and verifies whether the query image represents a logo image stored in the database. Nearest Neighbor features are clustered using Generalised Hough Transform (GHT) [24], [28], in order to define the parameters for a similarity transformation between the query and the database features. Then, an approximate affine transformation follows, which maps each NN from the query image onto the matched NN feature from the database image.

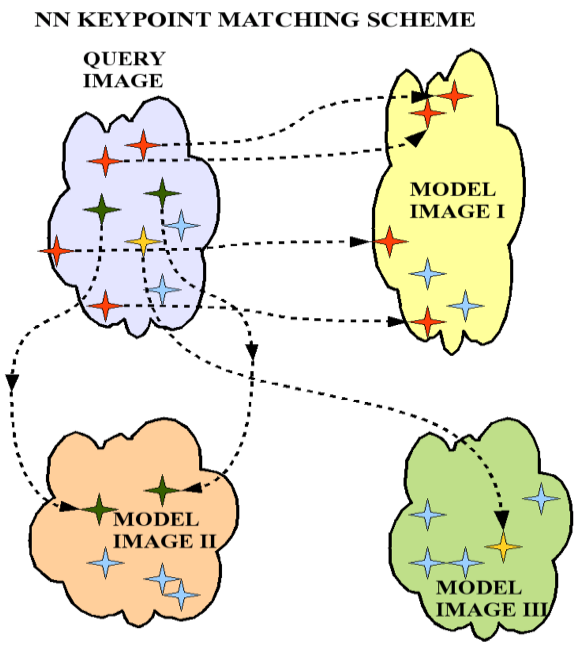

Fig. 1. Keypoint Nearest - Neighbour Matching Scheme.

For every invariant feature in the database there is a record that contains the image from which the feature comes from as well as the feature parameters position, orientation, scale and the 128-dimensional descriptor. Query Q and database D feature vector similarity transformation geometry is shown in Figure 3. In this procedure, λ is a normalization coefficient in order to achieve a reasonable size for the vector length compared to the image size (default λ=10).

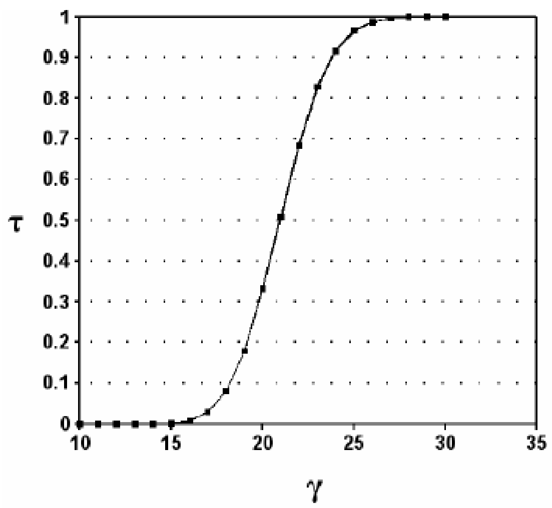

Fig. 2. Reduced Nearest Neighbour parameter (τ) plotted versus distance threshold (γ).

$$(x'_q, y'_q) = (x_q + \lambda s_q \cos\theta_q, y_q + \lambda s_q \sin\theta_q) \qquad (4)$$

$$(x'_d, y'_d) = (x_d + \lambda s_d \cos\theta_d, y_d + \lambda s_d \sin\theta_d) \qquad (5)$$

In order to be consistent with a similarity transformation, points (x'_q, y'_q), (x'_d, y'_d), (x_q, y_q), (x_d, y_d) in Equations (4) and (5) should satisfy the matrix relationship (homography) shown in Equation (6).

$$\begin{bmatrix} x_q & -y_q & 1 & 0 \\ y_q & x_q & 0 & 1 \\ x_q & -y_q & 1 & 0 \\ y_q & x_q & 0 & 1 \end{bmatrix} \begin{bmatrix} \Lambda\cos \\ \Lambda\sin \\ T_x \\ T_y \end{bmatrix} = \begin{bmatrix} x_d \\ y_d \\ x'_d \\ y'_d \end{bmatrix} \qquad (6)$$

where Λ is the scale change, φ is the rotation angle and T_x, T_y are the translation values of the similarity transformation. GHT identifies clusters of features with a consistent interpretation by using each feature to vote for all the logo images that are consistent with the feature. Therefore, a four dimensioznal hash table (accumulator) entry can be created, predicting T_x, T_y, Λ and φ parameters from the match hypothesis, through the similarity transformation, described in Equation (6). T_x and T_y get values between [-a, +a] and [-b +b], where axb is the image dimension. In addition, Λ and φ get values in [2-4 ,– 2+4] and [-π,+π] rads, respectively. Bins size of Δφ= π/6 rads, ΔΛ=2, ΔTx =a/4 and Δ Ty =b/4 are used for orientation, scale and translations, respectively. The above bin sizes are rather broad and allow clustering even in the case of substantial geometric distortion, due to a change in 3D point of view. This prediction has large error bounds, as the similarity transformation implied by these 4 parameters is only an approximation to the full 6 degrees-of-freedom pose-space for a 3D object and also it does not account for any nonrigid deformations. When clusters of features are found to vote for the same pose of a vehicle logo, the probability of the interpretation being correct is much higher than for any single feature. Therefore, when the NN distance between a pair of query and database feature is less than a certain threshold, this pair participates in a Hough similarity transformation and a vote is added to the corresponding cell in the Hough array. The database logo image with the highest number of votes is considered to be the most similar with the query image and it is forwarded to the geometrical consistency check with affine transformation (homography).

2) *Geometrical Verification*

In this step, the locations of the keypoints in the query and the database image are verified for geometrical consistency. Those keypoints that fit well to a homography transformation are called inliers, while those that are inconsistent to it are called outliers. A keypoint is considered to be an inlier or outlier according to its conformity to an affine geometrical transformation, which takes into account small differences in translation, rotation and scaling between the two images. Equation (7) defines the affine transformation for a set of query and data points.

$$\begin{bmatrix} x_q^1 & y_q^1 & 0 & 0 & 1 & 0 \\ 0 & 0 & x_q^1 & y_q^1 & 0 & 1 \\ x_q^2 & y_q^2 & 0 & 0 & 1 & 0 \\ 0 & 0 & x_q^2 & y_q^2 & 0 & 1 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ x_q^N & y_q^N & 0 & 0 & 1 & 0 \\ 0 & 0 & x_q^N & y_q^N & 0 & 1 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ V_x \\ V_y \end{bmatrix} = \begin{bmatrix} x_d^1 \\ y_d^1 \\ x_d^2 \\ y_d^2 \\ \cdots \\ x_d^N \\ x_d^N \end{bmatrix} \qquad (7)$$

In the above equation, a_i (i = 1, 2, 3, 4) are the warping parameters and V_x, V_y are the translation parameters of the affine transformation. If Equation (7) is considered as a system of A*b=c, then solving the respective equations, there is a least square solution for the parameters b as seen in Equation (8). This equation minimizes the sum of the squared distances from the query to the database image keypoint positions.

$$b = [A^T A]^{-1} A^T c \qquad (18)$$

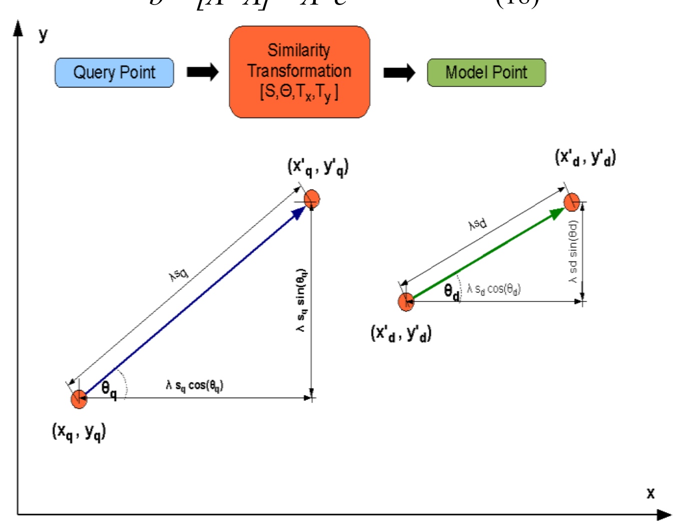

Fig. 3. Query and model keypoint vector similarity transformation geometry.

In order to select which candidate points should be used for affine transformation, the RANSAC method [29] is applied to a pair of images by randomly choosing three pairs of matched points of interest, since this is the minimum required for an affine transformation. These three pairs of points of interest are fitted to an affine transformation, which is then applied to the rest of the points of interest (keypoints) of the two images.

Inliers are identified by selecting those feature points in one image that are within a certain threshold distance from its matched feature of the other image and minimizing the fitting error. This process was repeated until the inlier count reaches a maximum or the fitting error is below a fixed value. The transformation with the highest inlier count is labeled as the best one to map the query image onto the image of the database. The total fitting error, ρ, between query and logo database images under affine transformation is:

$$\rho = \sqrt{\frac{\|Ab - c\|^2}{w - 2}} \qquad (9)$$

where w are the number of keypoints in matrix A, so if ρ in eq. (9) exceeds a threshold, we reject the respective match. Here we have used a threshold of value $\rho = 0.05*L_x$ where L_x is the maximum dimension of the query image. The above procedure was repeated, for all clusters found in the Generalised Hough Transform and the database image with the maximum number of inliers was selected. When two or more matches have the same number of inlier points, in order to choose the best match, the database image with the minimum number of outliers is selected. However, if the maximum clusters size is less than three points, the unconditional match hypothesis is adopted. In this case, the model with the maximum value of feature matches in the query image decides for the logo recognition. The whole process for feature detection, matching and hypothesis validation is shown in Figure 4, while an indicative example of the NN matches is given in Figure 5.

C. Experimental Part

1) Database set

The proposed method is assessed on a set of 1500 logo images that have been successfully segmented using our previous work described in [21], [22] using the Medialab LPR Database [26]. The Medialab LPR Database contains frontal views of vehicles, both stationary and in motion, captured by a Nikon Coolpix 885 adjusted to acquire 1024x768 pixel images. From the total of 1500 logo images, 400 have been selected to form the logo database as in [22]. These images correspond to 10 classes of selected vehicle manufacturers. 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 re-referenced to the new position, scale and orientation. In 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.

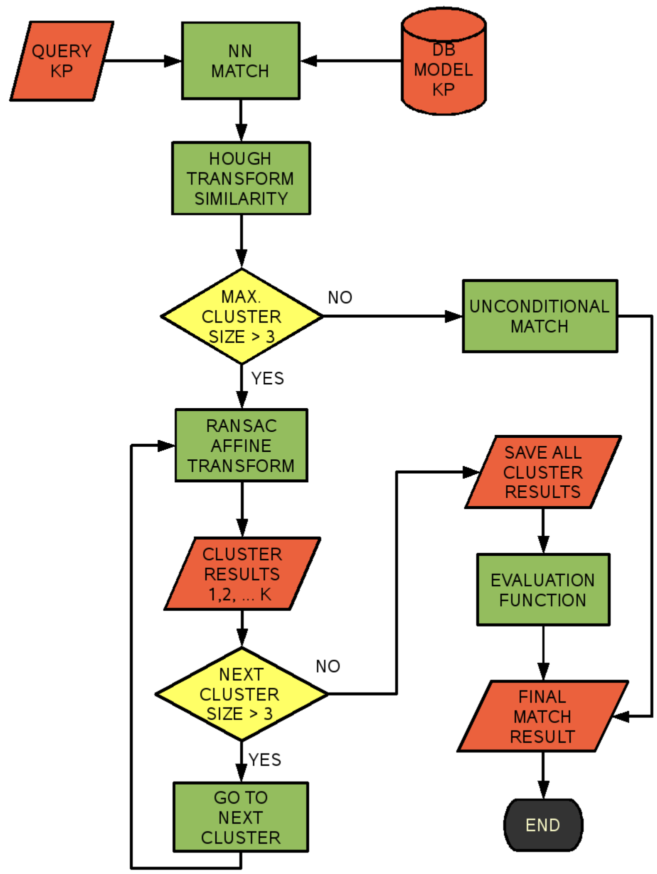

Fig.4. SIFT Match Hypothesis Testing Flow Chart.

2) Query set

The remaining 1100 logo images form the query set for our experiments. For every image, the keypoints and their respective descriptors have been calculated following the SIFT procedure. Then, these descriptors were matched against the descriptors of the database set, with the parameter γ set to γ=21. These values correspond to the turning point in the curve shown in Figure 2. For Euclidean NN search, a KD-Tree data structure is employed with search time complexity shown in Equation 10.

$$O(kN^{\left(\frac{k-1}{k}\right)}) \qquad (10)$$

where k is the tree dimension and N the number of tree nodes [30],[31]. For the experiments in this paper, the typical values are: k = 128 and N ~ 40000 (400 database images with approximately 100 keypoints per image). The image in the database 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 with RANSAC for geometric consistency. The recognition statistics are given analytically in Table I. There is an overall 94.6% correct recognition and 5.4% false recognition. The number of Hough-clustered keypoints is approximately 62% of the total keypoints detected, while the number of inliers is about 50% of the Hough-clustered, as presented in Table II. The recognition speed is shown for comparison with standard SIFT and M-SIFT, in Table III.

A correct match occurs when the two keypoints correspond to the same physical location and a false match when the two keypoints come from different physical locations. The total number of correspondences for the given dataset is known a priori and has been calculated by estimating the homography between each pair of images, for an affine transformation. Total correspondences are the total ground truth matches which are greater that the total matches (false + true) found. Using 1-Precision versus Recall, we can evaluate the total recognition performance, for a range of experimental conditions, varying the threshold value γ, for NN matching. The results from these experiments using the standard SIFT features and the enhanced (merged) ones, are plotted together for comparison in Figure 6.

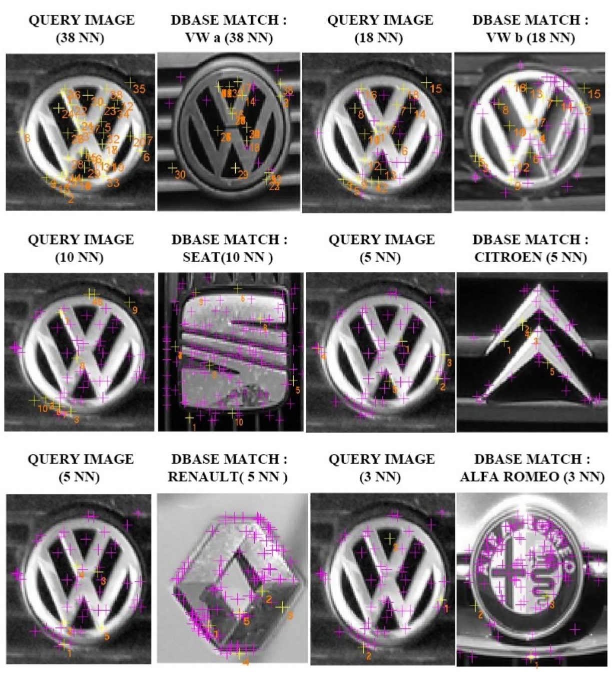

Fig. 5. Database Model NN matches for a sample query image.

III. CONCLUSIONS

This paper provides a description of a practical application for vehicle make recognition, based on and extending well known image processing routines, such as SIFT. Specifically, a modification of SIFT has been used, where the keypoints from similar objects were merged after homography transformation to the same coordinates and scales. With this technique, the number of keypoints for every model was substantially increased, thus making the recognition process more robust in outdoor conditions, such as low-contrast, inadequate lighting, reflections, cloudy weather etc.

TABLE I
M-SIFT RECOGNITION RATE

Manufacturer	True	False
Alfa Romeo	105	5
Audi	104	6
Bmw	105	5
Citroen	103	7
Fiat	102	8
Peugeot	106	4
Renault	103	7
Seat	104	6
Toyota	106	4
Volkswagen	103	7
Total	**1041**	**59**
Average (%)	**94.6%**	**5.4%**

TABLE II
KP SIZE DETECTED VS. HOUGH CLUSTER SIZE & RANSAC INLIERS

Query Image	Total 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**

TABLE III
VEHICLE LOGO RECOGNITION RATE

Methodology	Total Recognition Time (ms)
SIFT	850
M-SIFT	1020

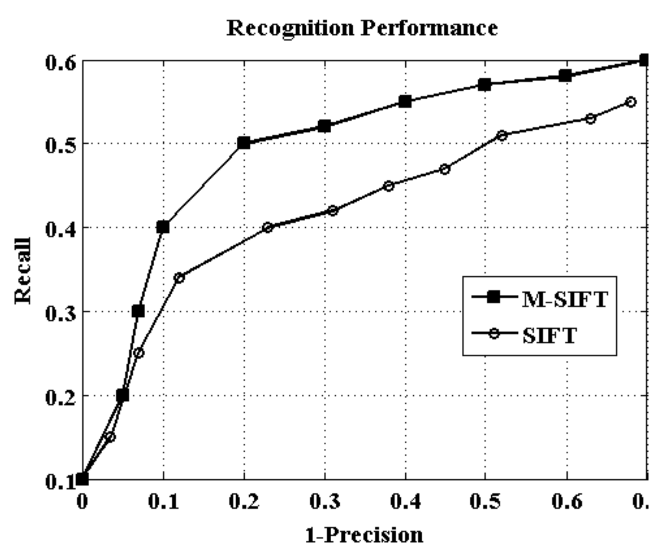

Fig. 6. Logo recognition performance for merged SIFT features and standard SIFT, respectively. 1-Precision and Recall calculated with various thresholds for NN matching.

The logo recognition system demonstrated good performance, yielding approximately 95% recognition rate, when applied to a database of already segmented logo images. One thousand and one hundred (1100) query logo images were processed with the enhanced SIFT-based approach, which gives better recognition accuracy compared to the standard SIFT procedure. The proposed experiments performs better compared to our previous work, in which a Probabilistic Neural Network (PNN) was used [20], [21] with 87% true recognition rate and extends the work described in [22].

The recognition speed is rather fast (~1 sec), which is suitable for real-time applications. To further boost performance and robustness, there is an ongoing research to combine logo recognition with vehicle model recognition (e.g. Audi A3). A short time video that demonstrates some preliminary results in real-time execution can be downloaded from [33]. Moreover, another extension of this work 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.

REFERENCES

[1] Weber M., Welling M., Perona P. “Unsupervised Learning of Models for Recognition”, Lecture Notes in Computer Science 1842, Springer-Verlag, pp. 18–32, 2000.

[2] Kato T. , Ninomiya Y., and Masaki I., “Preceding vehicle recognition based on learning from sample images”, IEEE Transactions on Intelligent Transportation Systems, vol. 3(4), pp. 252–260, 2002.

[3] Lai A.H.S., Yung N.H.C., “Vehicle-type identification through automated virtual loop assignment and block-based direction-biased motion estimation”, IEEE Trans. Intelligent Transportation Systems, 1(2), pp. 86-97, 2000.

[4] Lai A.H.S., Fung G. S. K. and Yung N.H.C., “Vehicle Type Classification from Visual-Based Dimension Estimation”, 2001 IEEE Intelligent Transportation Systems Conference, Oakland (CA), USA, pp. 201-206, 2001.

[5] S. Gupte, O. Masoud, R. F. K. Martin, N. P. Papanikolopoulos, “Detection and Classification of Vehicles”, IEEE Trans. on Intelligent Transportation Systems, Vol. 3, No. 1, March 2002, pp. 37-47.

[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. on Intelligent Transportation Systems, Vol. 7, No. 2, June 2006, pp. 175- 187.

[7] Lowe D., “Object recognition from local scale-invariant features”, International Conference on Computer Vision, Corfu, Greece (September 1999), pp. 1150-1157

[8] Lowe D.,, “Local feature view clustering for 3D object recognition”, 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'01), vol. 1, pp. 682-688, 2001.

[9] Lowe D., “Distinctive image features from scale-invariant keypoints”, International Journal of Computer Vision, vol. 2(60), pp. 91–110, 2004.

[10] Brown M. and Lowe D.,”Invariant features from interest point groups”, In British Machine Vision Conference, BMVC 2002, Cardiff, Wales (September 2002), pp. 656-665.

[11] Mikolajczyk K. and Schmid C. “An Affine Invariant Interest Point Detector”, European Conference on Computer Vision , vol 1, pp. 128 - 142, 2002.

[12] Harris C. and Stephens M., “A combined corner and edge detector”, Fourth Alvey Vision Conference, Manchester, UK, pp. 147-151, 1988.

[13] Siggelkow S., “Feature Histograms for Content-Based Image Retrieval”, PhD Thesis, Albert-Ludwigs-University Frieiburg , December 2002.

[14] Schulz-Mirbach H., ”Invariant Features for Gray Scale Images”, DAGM Symposium I, pp.1-14,1995.

[15] Dlagnekov L., Belongie S., “Recognizing Cars”, University of California San Diego, Tech. Rep. CS2005-0833, 2005.

[16] Čonos M., “Recognition of vehicle make from a frontal view”, Diploma Thesis, Faculty of Electrical Engineering, Czech Technical University, 2006.

[17] Petrovic V. S. and Cootes T. F., "Analysis of Features for Rigid Structure Vehicle Type Recognition", British Machine Vision Conference, vol. 2, pp. 587-596, 2004.

[18] Maemoto H., Okuma S., Yano Y., “Parametric Vehicle Recognition Using Knowledge Acquisition System”, IEEE International Conference on Systems, Man and Cybernetics, vol. 4, pp. 3982 – 3987, 2004.

[19] Lee H. J., “Neural Network Approach to Identify Model of Vehicles”, ISNN 2006, Lecture Notes in Computer Science 3973, 2006, Springer-Verlag, pp. 66–72.

[20] Psyllos A., Anagnostopoulos C.N., Kayafas E., Loumos V. “Image Processing & Artificial Neural Networks for Vehicle Make and Model Recognition”, 10th International Conference on Applications of Advanced Technologies in Transportation, May 28 - 30, Athens, Greece, 2008.

[21] Psyllos A., Anagnostopoulos C.,N., Kayafas E., “Vehicle Authentication from Digital Image Measurements”, 16th IMEKO TC4 Symposium & 13th Workshop on ADC Modeling and Testing Florence, Italy, September 22 - 24, 2008.

[22] Psyllos A., Anagnostopoulos C.,N., Kayafas, “Vehicle Logo Recognition using a SIFT-based Enhanced Matching Scheme” IEEE Trans on Intelligent Transportation Systems, vol. 11, no. 2, 2010, pp. 322-328.

[23] Wang Y., Liu Z., Xiao F., “A fast coarse-to-fine vehicle logo detection and recognition method” , IEEE International Conference on Robotics and Biomimetics, pp. 691-696, Sanya, China, Dec. 2007

[24] Ballard, D., “Generalizing the Hough transform to detect arbitrary patterns,” Pattern Recognition, 3(2), pp. 111-122,1981.

[25] Anagnostopoulos C., Anagnostopoulos I., Loumos V., Kayafas E. (2006), "A license plate recognition algorithm for intelligent transportation system applications", IEEE Transactions on Intelligent Transportation Systems, 7 (3), 377-392.

[26] The Medialab License Plate Recognition Database: http://www.medialab.ntua.gr/research/LPRdatabase.html last accessed March 2012.

[27] Lindeberg T., “Scale-space theory: A basic tool for analysing structures at different scales”, Journal of Applied Statistics, 21(2), pp. 224-270.

[28] Grimson, E., “Object Recognition by Computer: The Role of Geometric Constraints”, The MIT Press: Cambridge, MA, 1990.

[29] Fischler M., Bolles R.,“Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography”, Comm. of the ACM, vol. (24), pp. 381–395, 1981.

[30] Arya S., Mount D., “Approximate Nearest Neighbor Searching”, Proceedings 4th Ann. ACM-SIAM Symposium on Discrete Algorithms} (SODA), pp. 271—280, 1993.

[31] Friedman J., Bentley J., Finkel, R.,”An algorithm for finding best matches in logarithmic expected time”, ACM Transactions on Mathematics Software,vol. 3(3), pp. 209 – 226, 1977.

[32] Y. Ke, R., Sukthankar, “PCA-SIFT: A More Distinctive Representation for Local Image Descriptors”, Conference on Computer Vision and Pattern Recognition, vol .2 pp. 506-513, 2004.

[33] Video Demo: http://www.medialab.ntua.gr/research/LPRdatabase/sdemo.mpg, last accessed March 2012.