IMPROVED RAINDROP DETECTION USING COMBINED SHAPE AND SALIENCY DESCRIPTORS WITH SCENE CONTEXT ISOLATION

Dereck D. Webster*, Toby P. Breckon†

*Cranfield University, Bedfordshire, UK   †Durham University, Durham, UK

ABSTRACT

The presence of raindrop induced image distortion has a significant negative impact on the performance of a wide range of all-weather visual sensing applications including within the increasingly import contexts of visual surveillance and vehicle autonomy. A key part of this problem is robust raindrop detection such that the potential for performance degradation in effected image regions can be identified. Here we address the problem of raindrop detection in colour video imagery using an extended feature descriptor comprising localised shape, saliency and texture information isolated from the overall scene context. This is verified within a bag of visual words feature encoding framework using Support Vector Machine and Random Forest classification to achieve notable 86% detection accuracy with minimal false positives compared to prior work. Our approach is evaluated under a range of environmental conditions typical of all-weather automotive visual sensing applications.

Index Terms— rain detection, raindrop distortion, all-weather computer vision, automotive vision

1. INTRODUCTION

Despite their potential impact across a wide range of computer vision applications, prior work in the detection and removal of raindrops in video imagery is limited [1,4]. Furthermore, as the range of automotive applications such as speed sign detection [3], road feature detection [4] headlight detection [5] and road marking recognition [6] are becoming more prevalent in the automotive industry, all-weather operation is becoming a key topic for future vehicle autonomy. In addition, the impact of adverse weather conditions notably affects other common visual surveillance tasks [7,8]. However, despite the potential impact of raindrop distortion upon such approaches, work on raindrop removal remains in its infancy [1,9].

In Fig. 1 we see two typical examples of rain induced image distortion within an automotive setting. Fig. 1A shows a typical example of ill-defined droplets that are barely visible within the overall context of the image but still effectively occlude the localised image region, causing significant distortion. Such examples are difficult to detect due to low contrast differences with their surroundings and the lack of any clearly defined boundaries. By contrast, Fig. 1B shows well-defined, stereotypical raindrops which are prominent within the image context. These examples are characterized by inverted scene
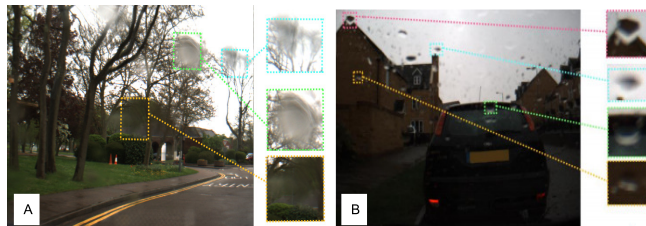

Fig. 1. Raindrop variation in terms of shape and colour.

illumination (dark top / light bottom) and exhibit clearly defined boundaries. From this variation (Fig. 1 A-B) we can readily appreciate the broad challenge of automated raindrop detection and removal.

Prior work on explicit raindrop detection generally follows two forms:- a) understanding of the photometric properties of raindrops [2] and b) feature based detection and removal [10]. The basis for photometric raindrop understanding leads to the ideal raindrop model [2]. Potential raindrop regions of a video image can thus be compared to the model forming a viable detection methodology [11]. Notably, variation within raindrops (e.g. Fig. 1) determines that not all such instances comply with any such model, compromising effective detection [2]. Furthermore, such approaches are characterised by high computational requirements limiting the applicability to real-time visual sensing applications [2,11]. The second strand of raindrop detection work focuses on a range of feature driven detection approaches [9], [12]–[14], [28]. The work of [12] uses stereo vision for raindrop detection whereby raindrops are detected via stereo matching, however atypical raindrops or poor conditions can cause failure [12]. Later work [15], using a spatial-temporal approach, utilizes consecutive video frames to analyse raindrops from a panning camera model in a visual surveillance context. However, this approach [15] requires prior knowledge of the camera dynamics and a relatively motion free scene between panning operations.

Within the automotive context, the work of [14] investigates windshield raindrop detection with the use of a specific hardware set up with the aim to detect raindrops through a reflection and focusing principle. A lens arrangement brings the raindrops into focus and blurring the background with a bespoke LED arrangement providing necessary illumination for nocturnal use. Raindrops are then detected by distinguishing sharp, in-focus regions in the image [14] but only over a limited field of view. Here, we base our work on the premise of [9] which investigates the idea that raindrops

are salient regions in the field of view and puts forward an adapted saliency map for detecting these regions within the *XYZ* colour space [16]. A multi scale approach derives colour and texture saliency feature maps over all three colour channels (for first stage candidate detection). These are combined with rudimentary shape information from a basic Hough circle transform [19,9] as an input to Support Vector Machine (SVM) [9] classification (for second stage verification). Although [9] produces strong detection results (precision / recall = 0.79 / 0.76, outperforming [2, 28]) it has been found to suffer from notably high false positives within this study (see - Table 2). In general, prior work in this domain is limited by its computational efficiency [2,11], sensing configuration [14] and/or overall detection accuracy [9,10, 12-15]. In this work, we consider a novel extension to the second stage of raindrop verification, that combines both the additional use of shape descriptors [21] and the isolation of raindrop features from the surrounding the scene context. This is shown, via a bag of visual words feature encoding, for both support vector machine (SVM) and random forests (RF) classification to comparatively outperform the current state of the art [2, 9, 28].

2. RAINDROP FEATURES

We extend the feature formulation of [9] to consider additional shape features from the candidate raindrop regions in addition to the isolation of colour and texture saliency features from the scene context. This will result in a stronger shape prior and reduced contextual bias within the subsequent raindrop classification model (Section 3).

2.1. Colour and Texture Saliency Features

Following the approach of [9], we construct saliency features based on a *XYZ* colour space transform of the original image (shown to improve raindrop discrimination in [13]). This results in three, single channel images which are passed through a multi-scale Gaussian pyramid [22], where the nth scale corresponds to sub-sampling by a factor of $2n$ for five scales, $n = \{0..4\}$. Intensity saliency maps, I, are then created for each *XYZ* colour channel by first detecting local intensity differences (i.e. Difference of Gaussian (DoG), contrast saliency) between the fine and coarser scale images. The DoG, denoted $\ominus$, are computed by simple pixel-to-pixel subtraction, between scales $c \in \{0, 1, 2\}$, and corresponding pixels from scales, $s = c + d$ (for $d \in \{2, 3, 4\}$, $s \leq 4$), after prior interpolation of the coarser resolution of scale s to the higher image dimensions of scale c. For each of the three *XYZ* colour channels we compute these differences as a set of maps for six different pairs of (c, s) values, $\{I(0, 2), I(0, 3), I(0, 4), I(1, 3), I(1, 4), I(2, 4)\}$ for $c \in \{0, 1, 2\}$ and $s = c + d$ with $d \in \{2, 3, 4\}$, $s \leq 4$ such that $I(c, s) = \|c \ominus s\|$. Subsequently, we obtain a global intensity saliency map, $\hat{I}$, as the summation of all six saliency maps, over all three XYZ colour channels, $i \in \{X, Y, Z\}$, as follows:

$$\hat{I} = \sum_{i}^{i \in \{X,Y,Z\}} \sum_{c=0}^{2} \sum_{s=c+d}^{s \leq 4} I_i(c,s) \ : \ d \in \{2,3,4\} \quad (1)$$

Separate texture saliency maps, T, are created following the same methodology but using a multi-scale pyramid constructed using a Laplacian of Gaussian (LoG) filter [18] over five scales (nth scale: sub-sampling by $2n$, $n = \{0..4\}$ as above). Here cross-scale summation, denoted $\oplus$, is computed to accumulate the LoG filter response of each pixel, for scales $c \in \{0, 1, 2\}$, and corresponding pixels from scales, $s = c + d$ (for $d \in \{2, 3, 4\}$, $s \leq 4$), with interpolation as before. Again a set of maps for six different pairs of (c, s) values, $\{T(0, 2), T(0, 3), T(0, 4), T(1, 3), T(1, 4), T(2, 4)\}$, are computed such that $T(c, s) = \|c \oplus s\|$. We similarly obtain a global texture saliency map, $\hat{T}$, from these texture maps over all colour channels, $i \in \{X, Y, Z\}$, as follows:

$$\hat{T} = \sum_{i}^{i \in \{X,Y,Z\}} \sum_{c=0}^{2} \sum_{s=c+d}^{s \leq 4} T_i(c,s) \ : \ d \in \{2,3,4\} \quad (2)$$

This results in a set of 18 intensity saliency maps, I_i (from 6 (c, s) scale pairings over three XYZ colour channels), and similarly 18 texture saliency maps, T_i, over the same. Each image pixel can now be represented by a 36-dimensional $\{\{\vec{I}\}, \{\vec{T}\}\}$ feature vector over which a k-means clustering [24] is used to construct a code word vocabulary for subsequent feature encoding of candidate raindrop regions via vector quantization (i.e. standard bag of visual words feature histogram encoding in dimensionality k [25]). The overall image can now be spatially represented by global saliency maps, $\hat{I}$ and $\hat{T}$.

2.1.1. Raindrop Candidates

Candidate raindrop regions are initially detected based on a weighted combination of these global saliency maps, $\hat{I}$ (Eqn. 1) and $\hat{T}$ (Eqn. 2), combined with an additional measure of shape saliency, $\hat{S}$, provided as the summed Hough circle detection accumulator maps [23] for circles of radii, $r = \{2..7\}$, following [9]. These form a global image saliency map, S_{map}, across all three saliency modalities (DoG intensity, LoG texture and Hough shape) as $S_{map} = \alpha\hat{I} + \beta\hat{T} + \gamma\hat{S}$ (where $\alpha = 0.2232$, $\beta = 0.3157$ and $\gamma = 0.4611$ following from the Ada-boost optimization of [9]). This saliency map, S_{map}, is subsequently thresholded for combined saliency exceeding τ_S and morphologically refined following [9]. As illustrated in Fig. 2, this results in the derivation of a global image saliency map, S_{map} (Fig. 2B), from a given example image (Fig. 2A) from which, via thresholding (Fig. 2C), we arrive at a bounded set of salient regions as candidate raindrops (i.e. first stage candidate detection - Fig. 2D). This set of candidate raindrops (e.g. Fig. 2D) typically contains a
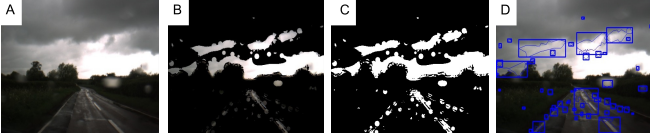
Fig. 2. Intermediate stages and output from the saliency-driven initial raindrop candidate detection stage, S_{map}.

high number of false candidates (i.e. non-raindrops) - a second stage of verification is subsequently employed to filter this set [9].

2.2. Shape Features

In order to improve the secondary verification of these raindrop candidate regions (Fig. 2D), we extend our feature descriptor used for final classification (from [9]) to include additional shape information. We use the seminal set of seven invariant Hu shape moments [21], $\phi_{i\in\{0..6\}}$, extracted from the droplet shape outline present within each candidate raindrop region for this purpose. These are calculated from the droplet contour via the structural analysis approach of [26] operating on the droplet outline, contained within each candidate region (e.g. 2D), obtained by thresholding S_{map} as outlined in Section 2.1.1 (e.g. Fig. 2C).

2.3. Scene Context Isolation

In a further extension to [9], which considers the entire bounding box of each candidate for saliency feature extraction (e.g. Fig. 2D), we compare performance using only features extracted from the interior region of the raindrop contour (as extracted in Section 2.2). Hereby the interior features of the raindrop are isolated from the varying scene context - a significant departure from prior feature extraction for this task [9,28].

3. CLASSIFICATION

From the bag of visual words feature encoding of our 36-dimensional per-pixel features, $\{\{\vec{I}\}, \{T\}\}$ (Section 2.1) and our additional Hu moment based shape descriptor ($\phi_{i\in\{0..6\}}$, Section 2.2) we have an overall feature descriptor length of $k + 7$ where k is the number of visual code words used in our earlier bag of visual words vocabulary. SVM and Random Forest (RF) classifiers are trained on these feature descriptors over a corpus of positive and negative raindrop candidates isolated by our first stage candidate detection (Section 2.1.1). SVM are trained using grid search over a Radial Basis Function (RBF), linear and polynomial kernel parameter spaces, $\{SVM_{RBF}, SVM_{Linear}, SVM_{Poly}\}$. RF [27] are trained over a set of varying configuration parameters including a max tree depth d, minimum sample count, sc and maximal tree count, t. Results from the best performing parameter and kernel configurations are presented in Section 4.

	TP%	TN%	FP%	FN%	P	R	A
RF, k=20, sc=25, d=24, t=100	60.24	**77.93**	**22.07**	34.91	**0.73**	0.65	**0.72**
RF, k=40, sc=20, d=18, t=1300	63.45	70.13	29.87	31.44	0.68	0.69	0.69
RF, k=60, sc=5, d=14, t=1900	**63.56**	71.08	28.92	**31.32**	0.69	**0.69**	0.70
RF, k=80, sc=5, d=22, t=1700	62.57	71.93	28.07	32.39	0.69	0.68	0.70
SVM, $k = 20$	**79.41**	13.24	86.76	**14.19**	0.48	**0.86**	0.48
SVM, $k = 40$	63.35	62.22	37.78	31.55	0.63	0.68	**0.65**
SVM, $k = 60$	66.56	52.51	47.49	28.08	0.58	0.72	0.62
SVM, $k = 80$	45.87	**76.09**	**23.91**	50.44	**0.66**	0.50	0.63

Table 1. Experimental results using [9].

4. RESULTS

Within the raindrop feature descriptor and classification approach outlined we consider the comparison of True Positives (TP), True Negatives (TN), False Positives (FP), False Negatives (FN) together with the Precision (P) Recall (R) and accuracy (A) against three variants:- the prior saliency descriptor approach of [9] (Table 1), the same but with use of the aforementioned scene context isolation (Table 3) and finally each of these with the addition of the Hu moments shape descriptor (Table 2 and Table 4). All test video imagery was gathered using a forward facing digital camera (1024×768 image resolution) mounted behind a car windscreen under a variety of road environments and weather (rainy) conditions (e.g. Fig. 1). Classifier training was performed using cross-validation over approximately 4500 positive raindrop image examples and 4000 negative non-raindrop image examples extracted from this source video. Testing was performed over an independent set of approximately 8000 images (4000 each of raindrop / non-raindrop).

For Table 1 (baseline, [9]) we can see a favourable TP outcome for $k = 20$ (code words) and the SVM classification (TP = 79.41%) but we also suffer notable high false positive outcomes (FP = 86.76%) with precision/recall similar to that originally reported in [9] (P/R: ~0.7; A: 0.7). The TN is low at 13.24% and similarly FN is 14.19%. For Table 1 we can see a general trend where increasing k has a detrimental effect on the performance of the SVM classification performance (falling TP). By contrast, RF classification performance of TP is seen to increase by 3% with decreasing TN by 6% under the same conditions.

Results shown in Table 2 show the impact of adding the additional shape prior into the classification model (Section 2.2). These results show a general increase in both TP and TN detection rates for both SVM and RF classifiers. The SVM in this case showing a TP of 69% for $k = 80$ with a decrease in FP to 7%. TN increased to 91.6% but FN similarly increased to 25.64% for $k = 80$. The performance of both classifiers plateau at $k = 40$ and $k = 60$, although overall accuracy

	TP%	TN%	FP%	FN%	P	R	A
RF, k=20, sc=25, d=14, t=1800	69.63	71.96	28.04	**24.76**	0.71	**0.75**	0.74
RF, k=40, sc=5, d=24, t=1900	**68.89**	74.89	25.11	25.56	0.73	0.74	0.75
RF, k=60, sc=20, d=22, t=1800	68.54	**77.47**	22.53	25.94	0.75	0.74	**0.76**
RF, k=80, sc=25, d=22, t=1800	67.55	77.30	**22.70**	27.01	**0.75**	0.73	0.75
SVM, k = 20	**69.03**	91.60	8.40	**25.41**	0.89	**0.75**	0.83
SVM, k = 40	68.43	92.87	7.13	26.06	**0.91**	0.74	**0.84**
SVM, k = 60	67.90	92.87	7.13	26.63	0.90	0.73	0.83
SVM, k = 80	68.82	**93.04**	**6.96**	25.64	**0.91**	0.74	0.84

Table 2. [9] with additional shape features (Section 2.2).

	TP%	TN%	FP%	FN%	P	R	A
RF, k=20 sc=5, d=10, t=1600	78.43	71.61	**28.39**	15.26	**0.73**	0.85	**0.78**
RF, k=40, sc=15, d=14, t=1200	**79.45**	71.12	28.88	**14.15**	**0.73**	**0.86**	**0.78**
RF, k=60, sc=20, d=10, t=1400	78.88	**71.33**	28.67	14.77	**0.73**	0.85	**0.78**
RF, k=80, sc=5, d=12, t=1800	79.24	68.82	31.18	14.38	0.72	0.86	0.77
SVM, k = 20	69.84	72.60	27.40	24.53	0.72	0.75	0.74
SVM, k = 40	66.07	79.45	20.55	28.62	0.76	0.71	**0.76**
SVM, k = 60	**73.41**	72.18	27.82	**20.68**	0.73	**0.79**	**0.76**
SVM, k = 80	56.57	**82.63**	**17.37**	38.88	**0.77**	0.61	0.72

Table 3. [9] with scene context isolation (Section 2.3).

(A) remains stable. Overall we can see superior accuracy (A) for the addition of shape information (Table 2) outperforming those of the prior state of the art of [9] (Table 1).

Table 3 shows us the impact of isolating the raindrop saliency features from the environmental context within the approach of [9] (without additional shape information). We observe that TP increases and a general increase in overall accuracy (A) (compared to Table 1).

The results shown in Table 4 show the impact of adding the additional shape prior in addition to scene context isolation. Results show the greatest increase in TP and TN detection rates for both classifiers against the prior state of the art [9] (Table 1). The SVM in this case shows a TP of 79% at $k = 80$ and a TN of 93% and the RF classifier shows a TP of 80% at $k = 80$ (overall accuracy ~0.86). Both classifiers show a significant improvements compared to prior art [9] (Table 1) with the additional use of our novel shape descriptor information (Section 2.2) and scene context isolation (Section 2.3). Overall it can be seen that the additional use of both results in the highest overall recall, precision and accuracy (Table 4, SVM).

Figure 3 A-B shows examples of successful raindrop verification using the combined approach of context isolation and

	TP%	TN%	FP%	FN%	P	R	A
RF, k=20, sc=10, d=8, t=1900	77.05	**79.38**	**20.62**	16.75	**0.79**	0.83	0.83
RF, k=40, sc=15, d=20, t=1900	78.64	77.44	22.56	15.03	0.78	0.85	0.85
RF, k=60, sc=5, d=24, t=1200	79.70	77.19	22.81	13.89	0.78	0.86	0.86
RF, k=80, sc=15, d=16, t=1200	**80.26**	72.99	27.01	**13.28**	0.75	**0.87**	**0.87**
SVM, k = 20	68.47	92.30	7.70	26.02	0.90	0.74	0.74
SVM, k = 40	76.87	92.83	7.17	16.94	0.91	0.83	0.83
SVM, k = 60	78.85	92.90	7.10	14.80	**0.92**	0.85	0.85
SVM, k = 80	**79.17**	**93.15**	**6.85**	14.46	**0.92**	**0.86**	**0.86**

Table 4. [9] with both context isolation (Section 2.3) and additional shape information (Section 2.2).
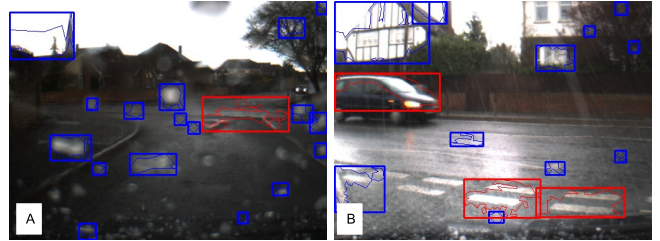

Fig. 3. Example output from classification (second stage raindrop verification) *[Blue = raindrop; Red = non-raindrop].*

additional shape feature information (Section 2.2 / 2.3). Figure 3A-B shows candidate raindrops (from initial detection, Section 2.1.1) that are now correctly verified as raindrops (blue border) and those correctly identified as non-raindrops (red border). Some mis-classification remains (e.g. false positive, Figure 3B - top left) and this remains an area for future work.

5. CONCLUSION

This work introduces two novel extensions for raindrop detection in video imagery - the use of additional shape priors in the classification model in the form of Hu moment based features and the use of scene context isolation for all features used in the secondary stage of raindrop verification. This produces a clear performance enhancement over the prior work of [9, 2, 28] with only a marginal impact on comparative computational efficiency. The use of Hu moments in a combined shape and texture feature descriptor with the addition of scene context isolation achieves a maximal 80% True Positive (TP) detection and 93% True Negative (TN) with marginal false positives and false negatives (Precision = 0.92 / Recall/Accuracy = 0.86). This contrasts sharply with the prior state of the art [2, 9, 15] which suffered from notably high false positive detection (as illustrated in this study).

Future work will consider further aspects of enhanced shape feature information and improved contour segmentation for the raindrop detection task.

6. REFERENCES

[1] K. Garg and S. K. Nayar, "Detection and removal of rain from videos," in Proc Comp. Vis.and Pattern Recog., pp. I–528–I–535 Vol.1, 2004.

[2] M. Roser and A. Geiger, "Video-based raindrop detection for improved image registration," in Proc Int. Conf. Comput. Vis. pp. 552–560, 2009.

[3] M. L. Eichner and T. P. Breckon, "Integrated Speed Limit Detection and Recognition from Real-Time Video," in Proc. IEEE Intell.Veh. Symp., pp. 626–631, 2008.

[4] M. Kowaliszyn and T. P. Breckon, "Automatic road feature detection and correlation for the correction of consumer satellite navigation system mapping," in Proc IET Road Transp. Inf. Control Conf. pp. 02–02, 2010.

[5] M. L. Eichner and T. P. Breckon, "Real-Time Video Analysis for Vehicle Lights Detection using Temporal Information", in Proc. 4th European Conference on Vis. Media Prod. pp. I–9,2007.

[6] A. Kheyrollahi and T. P. Breckon, "Automatic Real-time Road Marking Recognition Using a Feature Driven Approach", Mach. Vis. Appl., vol. 23, no. 1, pp. 123–133, 2012.

[7] T. P. Breckon, A. Gaszczak, J. Han, M. L. Eichner, and S. E. Barnes, "Multi-Modal Target Detection for Autonomous Wide Area Search and Surveillance," in Proc. SPIE Emer. Techn.s in Security and Def., vol. 8899, pp. 1–19,2013.

[8] T. P. Breckon, J. Han, and J. Richardson, "Consistency in Muti-modal Automated Target Detection using Temporally Filtered Reporting," in Proc. SPIE Electro-Optical Remote Sensing, vol. 8542, no. 85420L-1, pp. 23:1–23:12, 2012.

[9] Q. Wu, W. Zhang, and B. V. K. Vijaya Kumar, "Raindrop detection and removal using salient visual features", in Proc IEEE Int. Conf. Image Processing, pp. 941–944, 2012.

[10] S. You, R. T. Tan, R. Kawakami, and K. Ikeuchi, "Adherent Raindrop Detection and Removal in Video," in Proc Comp. Vis. and Pattern Recog., pp. 1035–1042, 2013.

[11] J. C. Halimeh and M. Roser, "Raindrop detection on car windshields using geometric-photometric environment construction and intensity-based correlation", in Proc. Int.Veh. Symp, pp. 610–615, 2009.

[12] A. Yamashita, M. Kuramoto, T. Kaneko, and K. T. Miura, "A virtual wiper - restoration of deteriorated images by using multiple cameras," in Proc Int Robots and Sys., vol. 4, pp. 3126–3131, 2003.

[13] A. Yamashita, I. Fukuchi, T. Kaneko, and K. T. Miura, "Removal of adherent noises from image sequences by spatio-temporal image processing," in Proc Robotics and Auto, pp. 2386–2391., 2008.

[14] S. Gormer, A. Kummert, S.-B. Park, and P. Egbert, "Vision-based rain sensing with an in-vehicle camera," in Proc. Int.Veh.Symp., pp. 279–284 2009.

[15] A. Yamashita, T. Harada, T. Kaneko, and K. T. Miura, "Removal of adherent noises from images of dynamic scenes by using a pan-tilt camera", in Proc Int. Robots and Sys., vol. 1, pp. 437–442 2004.

[16] L. Itti and C. Koch, "Computational modelling of visual attention.", in Proc Nat. Rev. Neurosci., vol. 2, no. 3, pp. 194–203, Mar. 2001.

[17] P. J. Burt, "Fast filter transform for image processing," Comput. Graph. image Process., vol. 16, no. 1, pp. 20-51, 1981.

[18] P. J. Burt and E. H. Adelson, "The Laplacian Pyramid as a Compact Image Code", in Proc Commun. IEEE Trans., vol. 31, no. 4, pp. 532–540, 1983.

[19] D. H. Ballard, "Generalizing the Hough transform to detect arbitrary shapes," Pattern Recog., vol. 13, no. 2, pp. 111–122, 1981.

[20] C. Cortes and V. Vapnik, in Proc "Support-vector networks," Mach. Learn., vol. 20, no. 3, pp. 273–297, 1995.

[21] M.-K. Hu, "Visual pattern recognition by moment invariants", in Proc Inf. Theory, IRE Trans., vol. 8, no. 2, pp. 179–187, 1962.

[22] R. Szeliski, Comp.Vis.: Algo. and Appl., 1st ed. New York, NY, USA: Springer-Verlag New York, Inc., 2010.

[23] P. E. Hart, "Use of the Hough Transformation to Detect Lines and Curves in Pictures," vol. 15, no. 1971, pp. 11–15, 1972.

[24] J. MacQueen, "Some methods for classification and analysis of multivariate observations," in Proc of the Fifth Berkeley Symp on Math. Stat. and Prob., Vol. 1: pp. 281–297, 1967.

[25] J. Sivic and A. Zisserman, "Video Google: a text retrieval approach to object matching in videos", in Proc Comp. Vis., pp. 1470–1477 vol.2 , 2003.

[26] S. Suzuki and K. Be, "Topological structural analysis of digitized binary images by border following", in Proc Comput. Vis, Graph. Image Process., vol. 30, no. 1, pp. 32–46, 1985.

[27] L. Breiman, "Random Forests," Mach. Learn., vol. 45, no. 1, pp. 5–32, 2001.

[28] H. Kurihata, T. Takahashi, and T. Miyahar, "Rainy weather recognition from in-vehicle camera images for driver assistance," in Intelligent Vehicles Symp., pp. 154–160, 2005.