Shadow Detection Algorithms for Traffic Flow Analysis: a Comparative Study

Andrea Prati, Ivana Mikić, Costantino Grana and Mohan M. Trivedi

Abstract— **Shadow detection is critical for robust and reliable vision-based systems for traffic flow analysis. In this paper we discuss various shadow detection approaches and compare two critically. The goal of these algorithms is to prevent moving shadows being misclassified as moving objects (or parts of them), thus avoiding the merging of two or more objects into one and improving the accuracy of object localization. The environment considered is an outdoor highway scene with multiple lanes observed by a single fixed camera. The important features of shadow detection algorithms and the parameter set-up are analyzed and discussed. A critical evaluation of the results both in terms of accuracy and in terms of computational complexity are outlined. Finally, possible integration of the two approaches into a robust shadow detector is presented as future direction of our research**

Keywords— **Traffic analysis, background suppression, shadow detection**

(a) Vehicle detection w/o shadow suppression (b) Vehicle detection with shadow suppression

Fig. 1. Example of vehicle undersegmentation due to shadows. Note without shadow suppression, two vehicles on the right part (a) can not be properly separated.

I. Introduction

DESIGN of vision-based systems for traffic analysis is an important and challenging problem of great practical value. In the context of Intelligent Transportation Systems, the information added by image processing techniques is very useful and it comes at a very low computational load. Many works on ITS aim at helping traffic flow management by providing information on how many vehicles are in the scene. Moreover, incident detection, queue detection and measurement, intersection management, and many other applications could exploit such information provided by visual tasks. All the above mentioned ITS applications aim, as first step, at detecting vehicles in the scene in order to count them, build up a behaviour-based database to identify erroneous or unauthorized operations or simply track them [1].

This task can be achieved by means of motion detection technique. and background suppression is the more commonly used. However, neither motion segmentation nor change detection methods can distinguish between moving objects and moving shadows.

Moving shadows cause the erroneous segmentation of objects in the scene. An example of undersegmentation and its correction by means of shadow suppression is shown in Fig. 1. To solve this problem, moving shadows have to be detected explicitly to prevent them being misunderstood as moving objects or their parts.

The paper is organized as follows. In Section II, the rationale of shadow detection approach is depicted and related work on shadow detection techniques are reported and briefly described. In Section III, the two systems under comparison are described. Their novelty, relevant issues and parameter set-up are discussed. Section IV presents our experimental results and in Section V we critically compare the two approaches, outlining the differences and the similarities. Finally, Section VI outlines the conclusions and the possible future directions of this work.

Andrea Prati and Costantino Grana are with the Dipartimento di Scienze dell'Ingegneria, Università di Modena e Reggio Emilia, Italy

Ivana Mikić and Mohan M. Trivedi are with Department of Electrical and Computer Engineering, University of California, San Diego

II. Shadows Detection and Related Work

It is not difficult for human eyes to distinguish shadows from objects. However, identifying shadows by computer is a challenging research problem. Shadows occur when objects totally or partially occlude direct light from a light source. According to the classification reported in [2], shadows are composed of two parts: *self-shadow* and *cast shadow*. The former is the part of the object which is not illuminated by the light source. The last one is the area projected on the scene by the object and is further classified in *umbra* and *penumbra*. The umbra corresponds to the area where the direct light is totally blocked by the object, whereas in the penumbra area it is partially blocked. This detailed classification has also been reported in [3].

In order to systematically develop and evaluate various shadow detector, it is useful to identify the following three important quality measures: *good detection* (low error probability to detect correct shadow points should occur), *good discrimination* (the probability to identify wrong points as shadow should be low, i.e. low false alarms rate) and *good localization* (the points marked as shadows should be as near as possible to the real position of the shadow point).

Several methods for identifying shadows have been developed in recent years: for aerial image analysis [4], for estimating objects shape [2] and for object detection enhancement [3][5][6], with particular relevance to ITS [7][8]. Most of the papers, and also the two system considered in the next section, do not examine the self-shadow (expect in [2]) and typically they concentrate the attention on umbra, considering the penumbra as a particular case of umbra. Moreover, in ITS there is little interest (typically) on detecting still cast shadows (due, for example, to trees in the scene), but only on *moving cast shadows*.

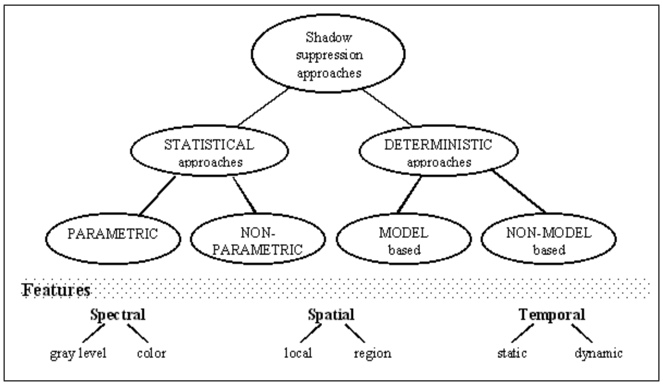

Fig. 2. Shadow detector approaches taxonomy. First the approaches are classified by means of their statistical or non-statistical approach. Further classification for parametric or model-based choose is done. Finally, every approach can be identified by the spectral, spatial and temporal features they exploit.

In Fig. 2 we show a possible taxonomy of existing approaches on shadow detection present in literature. Examples of a deterministic model-based approach are reported in [4][7][9] (exploiting gray level, local and static features), while approaches in [2][3] can be classified as deterministic non-model based (both use region and static gray level features). In the statistical approaches [8] is an example of the parametric approach and [5][6] of the non-parametric approach. All these [8][5][6] use color, region and dynamic features. In this paper, we focus on two particular approaches: a deterministic non-model based approach with color, local and static features, and a statistical parametric approach with color, region and dynamic features.

In [3], authors classify the previous work by means of their assumptions and state that four assumptions can be mainly found:

1. The light source is assumed to be strong.
2. The background and the camera are assumed to be static and the background textured.
3. The background is assumed to be planar.
4. The extent of the light source is sufficiently large to cause the formation of penumbra [2].

Even if some papers do not take into account these assumptions, most of them are based on them.

In [3], these assumptions are exploited to assert some statements. First, shadow points can be detected from the scene by means of frame difference. In fact, they demonstrate that difference between current frame and a reference frame is high in presence of shadows. Thus, typically the first step computed for shadow detection is the difference between the current frame and a reference image, that can be the previous frame, as in [3], or a reference frame, typically named *background model* [7][5][10]. Shadow points are part of the thresholded difference mask.

Second, the assumption 2 is used to distinguish between shadow and object in changed image regions (extracted using the above-mentioned difference). The textured background allows to extract edges in two consecutive images and to classify them as static or moving. In [3], this is achieved by classifying as static a pixel belonging to an edge if in a local neighborhood of it the energy in high frequencies of the frame difference is low.

Third, shadows smooth the background they cover. In fact, the assumption 3 says that the background is planar and the ligth source is distant from it. Starting from that, the authors demonstrate that this implies that shadow regions are characterized by low variance (in a local neighborhood of their pixel) of the ratio between the appearance of the pixels in current frame and that in the reference frame.

Frame differencing typically detects not only shadows, but also foreground points. The papers in literature mainly differ by means of how they distinguish between those points. In [8] a statistical approach is used to describe the distribution of shadow points by means of a-priori probability and a-posteriori estimation. Some other works [7][4][9] use the 3-D model information of both the scene and the objects to estimate position and size of cast shadows. Finally, many works locally exploit pixel appearance change due to cast shadows [2][6][5][3]. A possible approach is to compute the ratio between the luminance of pixels after and before shadow appears. This is actually what some authors exploit in their works [5].

III. Description of the Systems

In this section, two ITS management systems with shadow detection capabilities are presented: the former based on a deterministic non-model based approach, the last on a statistical parametric approach.

The first one has been developed in the Imagelab laboratory in the Dipartimento di Scienze dell'Ingegneria at the Università di Modena e Reggio Emilia, Modena, Italy. It is called **SAKBOT** (Statistical And Knowledge-Based Object Tracker) and it aims to be a general-purpose objects detector and tracker for both indoor and outdoor environments. See [10][11] for details.

The second system is part of the **ATON** (Autonomous Transportation agents for On-scene Networked incident management) project [8][12][1]. It has been developed in the CVRR (Computer Vision and Robotics Research) laboratory in the Department of Electrical and Computer Engineering at the UCSD, San Diego, USA. It integrates the shadows knowledge into a complex and distributed multi-sensor environment for incident detection and management and for traffic flow analysis.

A. Deterministic non-model based approach

Sakbot performs object detection by means of a background suppression. One of its major novelties is the way the background model is computed. Please see [10][11] for details.

Many works in literature fail to distinguish between moving shadow points and moving object points. In fact, object points typically have a inter-frame difference greater than a threshold and it is not unusual that the ratio between their luminance and the luminance of the corresponding reference image point is similar to that between shadows and reference image. Even the variance of the ratio in a neighborhood could be low, if the object is particularly smooth. For this reason, Sakbot exploits the information on chrominance too.

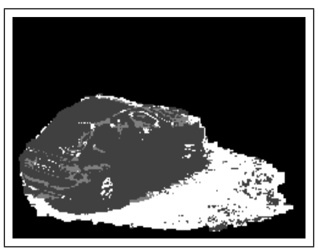

Fig. 3. Sakbot shadow detection. Light gray pixels would be detected as shadow using only luminance information, while white pixels are detected using the color information too. Note that the miscalssification of part of the car as shadow can be removed.

Using chrominance information increases the accuracy of the system, as it is possible to see in Fig. 3: black pixels are those classified as belonging to the background model, dark gray pixels those classified as foreground, light gray ones would be identified as shadows by means of only the luminance information, while white pixels are shadow points detecting using also the chrominance information. Removing light gray pixels from the shadow mask improves the accuracy by avoiding the misclassification as shadow of pixel belonging to the car. Nevertheless, this produces a wrong classification of some pixels of the shadow and does not eliminate all the wrong pixels of the car.

To do this, Sakbot first converts pixel information from the RGB color space to the HSV color space. HSV color space corresponds closely to the human perception of color [13] and it has revealed more accuracy to distinguish shadows. Then, it tries to estimate how the occlusion due to shadow changes the value of H, S and V. In Fig. 4, an example of shadow behaviour in HSV space is reported. Red arrows indicate three shadows passing over the point selected, whereas blue one indicates the passage of an object over it. The rationale is that cast shadow's occlusion darkens the background pixel and saturate its color. The first assumption is well evident in Fig. 4(a). The second assumption helps distinguish the object points from the shadow points. Another interesting point is that shadows often lower the saturation of the points.

Sakbot exploits these assumptions to define the shadow points mask SP_k as follows:

$$SP_k(x,y) = \begin{cases} 1 & if \quad \alpha \leq \frac{I_k^V(x,y)}{B_k^V(x,y)} \leq \beta \\ & \wedge \, (I_k^S(x,y) - B_k^S(x,y)) \leq \tau_S \\ & \wedge \, |I_k^H(x,y) - B_k^H(x,y)| \leq \tau_H \\ 0 & otherwise \end{cases} \tag{1}$$

where $I_k(x,y)$ and $B_k(x,y)$ are the pixel values at coordinate (x,y) in the input image (frame k) and in the background model (computed at frame k), respectively. The use of β allows to avoid identification as shadows of the points where the background was slightly changed by noise, whereas α takes into account the "power" of the shadow (see at the end of this subsection for further details). On channel S a threshold on the difference is performed. According to Fig. 4(b) and to the consideration above, this difference is usually negative for shadow points. On the channel H a threshold on the absolute difference produces better results.

In Fig. 5 some results of the shadow detector of Sakbot are presented. In Fig. 5(a) two shadows on the right are identified. This is correct because they are cast by two vehicles not visible in the scene. Consequently, Sakbot detects these shadows (Fig. 5(b)) and corrects this problem (Fig. 5(c)).

The values assigned to the four thresholds allow to adapt Sakbot to the environment conditions. As abovementioned, the value of β is useful for increasing the robustness of the system to noise. α takes into account how strong the light source is w.r.t. the reflectance and irradiance of the objects. Thus, stronger and higher the sun (in the outdoor scenes), the lower α should be chosen. The choice of the parameters τ_H and τ_S is less straightforward and, for now, is done empirically.

B. Statistical parametric approach

The moving shadow detection algorithm developed for the ATON project is described in detail in [8]. It is based on the detection of umbra of moving cast shadow in outdoor traffic video scene. It deals with non-modeled environments and it takes into account most of the assumptions described in Section II. Its key novelty is the use of three sources of information to help in detecting shadows and objects:

- *local* information based on appearance of the individual pixel
- *spatial* information based on the assumption that objects and shadows are compact regions in the scene
- *temporal* information that states that object and shadow position can be predicted from previous frames

The local information is exploited assuming that if $\mathbf{v} = [\mathbf{R}, \mathbf{G}, \mathbf{B}]^\mathbf{T}$ is the value of the pixel not shadowed, a shadow changes the color appearance by means of an approximated linear transformation $\bar{\mathbf{v}} = \mathbf{D}\mathbf{v}$, where $\mathbf{D} = diag(\mathbf{d_R}, \mathbf{d_G}, \mathbf{d_B}) = diag(\mathbf{0.48}, \mathbf{0.47}, \mathbf{0.51})$ is a diagonal matrix obtained by experimental evaluation and on the assumptions described in Section II. Moreover, in traffic video scenes it has been observed that shadowed surfaces appear bluer. Given the means and the variances for the

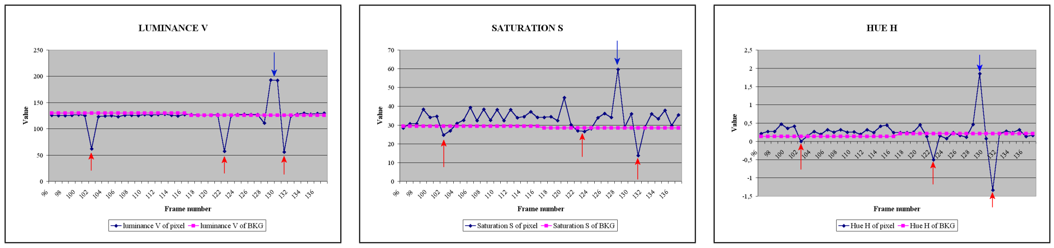

(a) Luminance (b) Saturation (c) Hue

Fig. 4. HSV color space components change due to shadows. Red arrows indicate the passage of three shadows over the point. Blue arrow indicates the passage of a car over it.

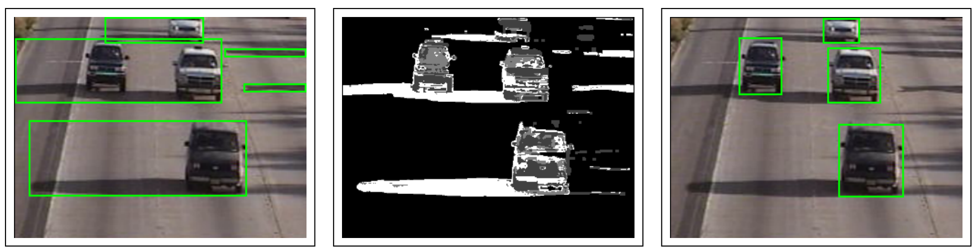

(a) Detection w/o shadow suppression (b) Shadows detection (c) Detection with shadow suppression

Fig. 5. Examples of Sakbot's shadow detector. Note that shadows are included in the moving objects if no shadows suppression is applied (a) and that, using the shadow detection (b), Sakbot removes the shadows from the objects (c).

three color channel for reference points, it is straightforward to derive the corresponding values under shadows as $\mu^i_{SH} = \mu^i_{IL} d_i$ and $\sigma^i_{SH} = \sigma^i_{IL} d_i$, with $i \in R, G, B$.

Pixels in current frame are classified by ATON into three classes: background, foreground and shadow. A fading memory estimator calculates background means and variances for all pixel locations and derived statistics for the same pixel when shadowed are obtained using the equations above. Then, ATON starts the segmentation by comparing the feature vector for each pixel to the mean at that location in the background model. If not significantly different, pixel is classified as background. Otherwise, ATON assigns to the pixel the a-priori probabilities p_{BG}, p_{FG} and p_{SH} of belonging to background, foreground and shadow classes, respectively. Values chosen are $p_{BG} = 0.3$, $p_{FG} = 0.4$ and $p_{SH} = 0.4$. A pixel is then classified maximizing the a-posteriori probability of the class membership:

$$p(C_i/\mathbf{v}) = \frac{p(\mathbf{v}/C_i)p(C_i)}{\sum_{\mathbf{j=R,G,B}} p(\mathbf{v}/C_j)p(C_j)} \quad (2)$$

ATON improves the detection performance imposing spatial constraints. It performs an iterative probabilistic relaxation to propagate neighborhood information. This is achieved by assigning the class membership probabilities in the next iteration using the results of the previous step on the neighborhood pixels. The scheme (reported in [8]) converges quickly, and there is no noticeable change beyond the second iteration. After this, a post-processing is performed: small gaps in foreground regions are eliminated by scanning horizontally and vertically the image and then a morphological opening is computed.

For now, ATON does not exploit the temporal information, achieving good experimental results even without it. In Fig. 6 an example of ATON's shadow detector is reported: in Fig. 6(b) red pixels are classified as foreground, blue as shadow and green as background.

The values of $\mathbf{D}$ matrix and the starting a-priori probabilities are parameters to be chosen for a certain video sequence, and should be recomputed for different light and environment conditions. However, these are the only parameters to be set (a part for the threshold on the difference between the mean of the background and the current pixel, but this is almost independent of the sequence).

IV. Experimental Results

In this section, the two systems are compared by means of their *accuracy*. This issue takes into account both the goodness features of good detection and good discrimination, emphasized at the begin of Section II. The first one can be achieved by minimizing the *false negatives*, i.e. the shadow points classified as background/foreground, while to obtain a good discrimination, the *false positives*, i.e. the foreground/background points detected as shadows, should be minimized. In Table I, these two types of evaluation are reported for both systems, before and after their respective improvements.

To compare the systems some frames have been manually segmented in order to identify shadows, foreground and

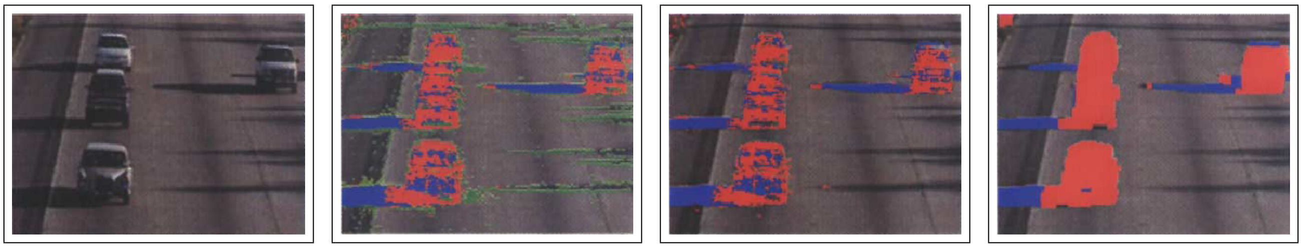

(a) Original image frame (b) Classification results (c) Classification results w/o background pixels (d) Final result after post-processing

Fig. 6. Example of ATON's shadow detector. (a) reports the input image. (b) Classification results exploiting only local information. Red pixels are classified as foreground, blue as shadows and green as background. (c) is the same of (b) with background pixels not shown and (d) is the result after spatial exploitation and post-processing.

	FN	FP	Shadow Accuracy%	Overall Accuracy%
Sakbot w/o color	1508	3190	60.59%	91.65%
Sakbot with color	1806	1778	**68.56%**	93.10%
ATON w/o post-proc.	2529	1527	57.37%	92.09%
ATON with post-proc.	2654	1159	**60.67%**	92.65%

TABLE I

Experimental measurements of accuracy. Both the systems achieve an higher accuracy by introducing the respective improvements (compare row 1 and 3 with row 2 and 4. The overall accuracy is more than the 90%.

background regions. In the test set different situations have been considered (dark/ligth cars, multiple cars or single car, occlusions or not) to stress the Sakbot and ATON systems. The input sequences have been collected from the highway I-5 near San Diego, CA.

In Table I, the **FN** and **FP** columns report the false negatives and the false positive in recognizing the shadows, respectively. The column **Shadow Accuracy%** reports the accuracy in percentage computed as the ratio between the total number of pixels correctly identified as shadows or foreground and the total number of them present in the manually segmented images. This value takes into account both the false negatives and the false positive and it is the main aim of this comparison. However, in the last column of Table I the **Overall Accuracy%** of the system is reported and it considers the background points too.

From the experimental results, the improvements achieved using color information in the Sakbot system and the post-processing module in the ATON system are clearly demonstrated. For Sakbot, the color addition heavily decreases the amount of **FP** (i.e. decrease the misclassification of car points as shadow), but slightly increases the amount of **FN**. However, the shadow accuracy is increased.

Also in ATON the improvement of post-processing decreases the amount of **FP** and increases at the same time the number of **FN**. This is due to the fact that the post-processing mainly acts on the foreground regions, trying to fill small gaps in them. This actually removes shadow points assigned to foreground regions, thus reducing false positive. For the same reasons, however, this dilatation of foreground regions assigns some shadow points to foreground class, thus the number of false negative increases slightly. Also in this case, the accuracy is increased.

Even if it is the only quantitative way to evaluate the systems, the accuracy is not enough: the capability of shadow suppression to improve the segmentation and tracking systems is the more important issue! If the shadow detection is followed by region labeling, we can achieved an high accuracy (but with sparse shadow pixels) and, nevertheless, we could be unable to cope with undersegmentation. For example, using the Sakbot parameters that produces the results in Table I leads (in some frames) to a fragmented shadow, unable to separate correctly two vehicles. Slightly changing the Sakbot parameters allows to accurately segment the scene, and the shadow accuracy increases from $68,56\%$ to $69,37\%$. However, the **FP** raise from 1778 to 2274, i.e. the vehicles are eroded. In conclusion, there is a trade-off between accuracy and effectiveness of the system and this must be taken into account during parameter selection for both systems.

An accuracy at *object level* would be a good issue to evaluate the effectiveness of the system. Unfortunately, this is typically a qualitative evaluation, difficult to do without introducing any uncertainty due to subjective evaluation. For this reason, this work does not consider this type of evaluation.

V. A Critical Analysis of Comparison

Sakbot and ATON have some differences. In general, Sakbot uses a *deterministic non-model based local* approach, whereas ATON a *statistical parametric spatial* one. This implies a different approach on the detection of foreground and shadow points. In fact, Sakbot uses a classical background suppression approach and obtains foreground points with a difference between images. Moreover, it does not exploit any spatial information and does not perform any post-processing on shadow detection results. ATON describes the distribution of classes of pixels with a Gaussian and classifies each pixel by means of its class membership.

In both systems, a background model is implicitly computed and the way to compute it differs only in the statistical function used to describe the background behaviour (Sakbot computes the background model as a temporal median over last frames, whereas ATON uses the Gaussian distribution). Background modeling is very reliable w.r.t. noise sensitivity and responsiveness to changes. Nevertheless, the statistical and knowledge-based background updating of Sakbot [10][11] proves an higher responsiveness to changes in the environment, whereas the statistical smoothness of the ATON classifier shows an higher robustness to noise.

The shadow detectors have some similar features: both use the color information, both assert that the assumptions described in Section II hold, and both do not exploit any model of scene and objects. However, the approach to shadow detection is different. First, they use two different color spaces. Second, the local pixel by pixel computation is smoothed by the Gaussian in ATON (statistical decision), whereas Sakbot makes an on-off (deterministic) decision. This implies a more compact detection of shadows for ATON, consolidated by the spatial exploitation.

Referring to Table I, the two systems achieve quite comparable results in terms of accuracy. However, it is straightforward to see that ATON finds more false negatives, whereas Sakbot achieves an higher number of false positives. This means that ATON better achieves a *good discrimination* and Sakbot better fits with the *good detection* requirement. Even if this evaluation is not exhaustive, it seems possible to assert that both the systems operate very well as shadow detectors, but both have some drawbacks.

Finally, the computational load of both systems is negligible. In fact, a color space conversion, a ratio and two differences have to be computed for each pixels of the image for Sakbot. On the other hand, ATON achieves local classification by simply computing the ratio in eq. 2 , and the spatial improvement by two iterations of the local classification and a computationally simple post-processing. On a Pentium III 600 MHz, the average computational time is of 69,38 msec/frame for ATON and of 44,10 msec/frame for Sakbot (not including the background updating system).

VI. Conclusions and Future Directions

In this paper we presented two system for ITS management able to handle shadows to improve vehicle detection and tracking. Then, we described the two systems and the motivations that drive their approaches. Finally, we compared them by means of their accuracy in shadow detection and by their computational load and tried to outline their advantages and their drawbacks. They both achieve a shadow detection accuracy of more than 60% and more than 90% for correctly classifying each pixel in background, foreground or shadows. We have concluded that ATON system better distinguishes between moving shadows and moving objects, whereas Sakbot better detects moving shadows.

Some future directions for the single system could be: improving parameter set-up of ATON in order to better adapt them to the input video sequence and trying to explain the hue/saturation behaviour of shadowed points for Sakbot. But the more attractive direction is to study possible integration of the two systems, combining the advantages and discarding the drawbacks. The possibilities can be at least two: to include spatial information and post-processing into the Sakbot system or to try ATON in the HSV color space. The first one will correct the false positive problem of that system, while the second one could improve the detection of shadows, decreasing the false negative. Finally, a more complete comparison between shadow detection approaches in literature is a very interesting future direction that we will try to explore.

References

[1] M.M. Trivedi, S. Bhonsle, and A. Gupta, “Database architecture for autonomous transportation agents for on-scene networked incident management (aton),” in *Proceedings of Int'l Conference on Pattern Recognition*, Sept. 2000.

[2] C. Jiang and M.O. Ward, “Shadow identification,” *Proceedings of IEEE Int'l Conference on Computer Vision and Pattern Recognition*, pp. 606–612, 1992.

[3] J. Stauder, R. Mech, and J. Ostermann, “Detection of moving cast shadows for object segmentation,” *IEEE Transactions on Multimedia*, vol. 1, no. 1, pp. 65–76, Mar. 1999.

[4] R.B. Irvin and D.M. McKeown Jr., “Methods for exploiting the relationship between buildings and their shadows in aerial imagery,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 19, pp. 1564–1575, 1989.

[5] A. Elgammal, D. Harwood, and L.S. Davis, “Non-parametric model for background subtraction,” in *Proceedings of IEEE ICCV'99 FRAME-RATE Workshop*, 1999.

[6] T. Horprasert, D. Harwood, and L.S. Davis, “A statistical approach for real-time robust background subtraction and shadow detection,” in *Proceedings of IEEE ICCV'99 FRAME-RATE Workshop*, 1999.

[7] M. Kilger, “A shadow handler in a video-based real-time traffic monitoring system,” *Proceedings of IEEE Workshop on Applications of Computer Vision*, pp. 11–18, 1992.

[8] I. Mikic, P. Cosman, G. Kogut, and M.M. Trivedi, “Moving shadow and object detection in traffic scenes,” in *Proceedings of Int'l Conference on Pattern Recognition*, Sept. 2000.

[9] D. Koller, K. Daniilidis, and H.H. Nagel, “Model-based object tracking in monocular image sequences of road traffic scenes,” *International Journal of Computer Vision*, vol. 10, pp. 257–281, 1993.

[10] R. Cucchiara, C. Grana, M. Piccardi, and A. Prati, “Statistical and knowledge-based moving object detection in traffic scene,” in *Proceedings of IEEE Int'l Conference on Intelligent Transportation Systems*, Oct. 2000, pp. 27–32.

[11] A. Prati, R. Cucchiara, C. Grana, and M. Piccardi, “Robust background update for moving visual objects segmentation,” in *submitted to IEEE Int'l Conference on Image Processing*, 2001.

[12] M.M. Trivedi, I. Mikic, and G. Kogut, “Distributed video networks for incident detection and management,” in *Proceedings of IEEE Int'l Conference on Intelligent Transportation Systems*, Oct. 2000, pp. 155–160.

[13] N. Herodotou, K.N. Plataniotis, and A.N. Venetsanopoulos, “A color segmentation scheme for object-based video coding,” in *Proceedings of the IEEE Symposium on Advances in Digital Filtering and Signal Processing*, 1998, pp. 25–29.