

A Change-Detection Algorithm Based on Structure and Colour

Luigi Di Stefano, Stefano Mattoccia, Martino Mola
University of Bologna

Advanced Research Center on Electronic Systems (ARCES-DEIS)
Via Toffano 2/2, 40135, Bologna, Italy
{ldistefano, smattoccia, mmola}@deis.unibo.it

Abstract

The paper proposes a novel change-detection algorithm for automated videosurveillance applications. The algorithm is based on the idea of incorporating into the background model a set of simple low-level features capable of capturing effectively “structural” (i.e. robust with respect to illumination variations) information. Thanks to this approach, and unlike most conventional change-detection algorithms, the proposed algorithm is capable of handling correctly still and slow objects as well as of working properly throughout very long time spans. Moreover, the algorithm can naturally interact with the higher-level processing modules found in advanced video-based surveillance systems in order to allow for flexible and intelligent background maintenance.

1 Introduction

Automated videosurveillance systems analyse the incoming video stream in order to detect anomalous or dangerous situations and then trigger automatically an adequate counteraction, such as for example sending an alarm to an operator and/or recording the incoming images. Typically, these systems deploy one or more static cameras. Most systems currently available in the videosurveillance market do not try to achieve a high degree of scene understanding but restrict the analysis of the incoming images to the detection of relevant colour or gray-scale changes with respect to a reference image. However, an impressive amount of ongoing research aims at the development of advanced videosurveillance systems, that should be able to trigger a counteraction on the basis of the comprehension of the nature and behaviours of the objects appearing in the scene. Typically, the latter systems include as the very first image analysis step a change-detection algorithm aimed at segmenting out the interesting regions from a static background. Then, higher level processing modules, such as tracking, classifi-

cation and interpretation modules, process the output yield by the change detection algorithm in order to attain the required degree of scene understanding.

Most change-detection algorithms rely on the principle of background subtraction: a background model, consisting of a representation of the observed scene without moving objects, is compared against the current image in order to mark as foreground those pixels that exhibit a significant difference with respect to the corresponding background pixels. The main difficulty associated with change-detection is not the background subtraction step, but instead the maintenance of a background model that follows correctly the changes of the reference scene. These changes can be grouped into two major classes:

Illumination changes: sun-light intensity and inclination vary slowly, continuously and regularly during the day, thus varying the illumination conditions of outdoor scenes (and in many circumstances, also of indoor scenes) and causing *slow image changes*. *Sudden image changes* occur when the illumination intensity varies much more rapidly, for example, when a cloud appears in a sunny day or, in an indoor environment, when lights are switched.

Changes due to objects: these changes occur when an object is introduced into or removed from the reference scene. If a foreground object stops a decision should be taken on whether and when it is more appropriate to consider the object as part of the reference scene and consequently include its appearance into the background model. Similarly, if a background object starts moving, its absence in the previously occupied image region is detected by the subtraction step as a bogus blob (usually known as *ghost*). In this case, the elimination of the ghost is typically desirable and can be achieved by updating the background in the image region previously occupied by the object.

In this paper we address the change detection problem in the case of a static camera and propose a novel algorithm based on the synergic combination of colour and structural (i.e. robust with respect to illumination variations) information.

2 Related work

Among the change detection algorithms relying on the background subtraction principle, the statistical approach is by far the most widely adopted one [10, 5, 4, 8, 3, 6]. With this approach some features are selected to represent the background pixels (e.g. RGB components, hue, luminance) and modelled by a probability distribution. A pixel of the current image is then classified as foreground if the observed features are not coherent with the associated probability distribution. In statistical algorithms background maintenance consists typically in updating the parameters of the probability distributions based on the last observed images.

Unlike statistical approaches, in [9] a Wiener filter is used at each pixel to obtain a prediction of the new intensity value on the basis of the last intensity samples. If the observed value differs significantly from the predicted one, the pixel is classified as foreground. In [1] a similar strategy is used, but the prediction is achieved by means of a Kalman filter. In the latter two algorithms background updating is implicit since the output of the filters is strictly dependent upon the measures taken from the last images.

Similarly to our approach, in [7, 2, 6, 5] the image gradient is used to capture structural information for illumination-insensitive change detection. In [2] the background subtraction step relies on colour edges only, while in [7, 5] gradient and colour information is combined at pixel level to identify foreground points.

The method proposed in [6] is the most similar to our approach since the combination of colour and gradient information is performed at region level. A statistical grey-level gradient background model is maintained. Then, the edge points of the current image where the gradient differs from the background gradient are highlighted. An adaptive background subtraction step, that models the pixel colour as a mixture of Gaussians, produces connected components that are validated computing, on their boundaries, the percentage of pixels highlighted in the previous gradient-based subtraction step: only the components showing high percentages are finally classified as foreground. The objective of this strategy is to remove spurious regions caused by illumination changes and ghosts.

Conversely, our approach does not rely on a statistical background and colour gradient information is extracted at a lower resolution. The latter choice yield improved robustness to illumination changes and renders the algorithms as fast as to meet the real-time requirements of many applications. Moreover we propose a novel background updating scheme that explicitly exploits illumination-insensitivity and enables interaction with higher-level processing modules for intelligent maintenance.

3 The proposed approach

In the above mentioned algorithms a foreground object is immediately and gradually included into the background as soon as it stops moving so that, after a certain time interval, it will be perceived as background. Similarly, a ghost will be absorbed into the background according to the same dynamics. It is worth pointing out that, though not always explicitly highlighted by the authors, this common strategy allows the change-detection algorithm to recover from persistent false positive errors. These errors are due to background estimation errors, which are generated when illumination changes are not compensated adequately, as well as to ghosts. This capability is of fundamental importance for proper long-term functioning of the change-detection algorithm. However, this common strategy relies on a user-selectable time constant that determines the time needed for a pixel repetitively classified as foreground to become part of the background. If the time constant is fast false positives are absorbed very quickly into the background model but slowly moving objects may corrupt some region of the background, generating erroneous blobs. Conversely, if the time constant is slow, slowly moving objects are correctly detected but recovering from false positive errors takes a long time. It is clear from these considerations that the choice of the time constant implies a critical trade-off that often does not satisfy adequately all the different application requirements.

Ideally, if the background model were based on features invariant to illumination changes, the trade-off problem would be removed. In fact, in such a case, the background model should be updated only to accommodate the changes due to objects. Even though it seems impossible to find low-level features invariant to every kind of illumination changes, it is possible to devise low-level features resulting quite robust with respect to many illumination changes occurring in practise. If we take into account real features instead of ideal ones, we can observe that the more the adopted features will be robust the better the real algorithm will approximate the ideal behaviour (i.e. updating of the background model only to accommodate the changes due to objects).

Starting from these considerations, we have devised a set of very simple, low-level features, described in section 4, that have proven to be very robust with respect to brightness variations. By including these features into the background model we have obtained a novel change detection algorithm, described in section 5, that approximates quite satisfactorily the ideal behaviour outlined previously. To conclude this section we wish to point out an additional favourable feature of the proposed approach. We believe that the insertion/removal operations needed to update the background model according to the scene changes due to

objects could be carried out much more flexibly and intelligently by exploiting information concerning the nature and interactions of the objects appearing in the scene as well as context information. For example, if a car stops in a parking lot, it is reasonable to absorb it very quickly into the background since it is highly probable that it will remain still for a relatively long time. This choice will facilitate the detection of the driver as he gets off the parked car. On the other hand, if a car stops in front of a traffic light it is preferable not to include it rapidly into the background since, probably, the car will start again quite soon. Our approach naturally holds the potential to support such intelligent handling of the background model simply because the model is updated only to accommodate the scene changes due to objects. Unfortunately, the type of information required to actually carry out the intelligent handling of the background model cannot be extracted by the change detection algorithm itself. However, within an advanced videosurveillance framework, this information could be feed-back to the change detection module by the higher-level processing modules. Therefore, our change-detection algorithm has been designed to support this kind of interaction: it can accept as input a binary mask that controls the inclusion and removal of objects into the background model. In absence of any feed-back from higher-level processing modules, the algorithm updates autonomously the background model in order to accommodate the scene changes due to objects.

4 Image Structure

As discussed in the previous section, our approach relies on incorporating into the background model a set of features very robust with respect to illumination changes. To come-up with such desirable features we may observe that if we compare two snapshots of a scene taken by a static camera under different illumination conditions (e.g. one in the morning, the other in the afternoon), we find that, even though the colours locally can change significantly, the overall “structure” of the image is usually perceived as unchanged. For example, if in the first image we find a big rectangular red region (e.g. a wall) on the left, then a thin gray stripe in the center (e.g. a road) and an other big, irregularly shaped, green area on the right (e.g. trees), this will be the “structure” perceived in the second image too, although the actual red, gray and green can look quite different. Obviously, to incorporate some kind of structural information within the background model of a change detection algorithm we should be able to capture this information by means of *low-level* features. Moreover, since in practical applications the change-detection algorithm must often run at video-rate (or nearly video-rate), we would like the structural description to be as simple as possible. We could think out a way to capture structural information by

estimating the gradient components and then analysing the edges, as e.g. in [7, 2, 6, 5]. In practice, however, the exact position and intensity of edge points often varies too much as a function of the illumination conditions. Yet, starting from a similar “derivative” concept, we have devised a simple transformation yielding a representation of the original image, called *structure*, that turns out to be quite stable with respect to changes of the illumination intensity (i.e. to brightness variations). Given a $w \times h$ gray-level image, I , the first step of the transformation consists in obtaining a reduced-resolution image, $R[I]$. Let δ be a scale factor (in our experiments $\delta=10$), then $R[I]$ is a $\frac{w}{\delta} \times \frac{h}{\delta}$ image defined as:

$$R[I](x, y) = \frac{1}{\delta^2} \sum_{j=0}^{\delta-1} \sum_{k=0}^{\delta-1} I(\delta \times x + j, \delta \times y + k) \quad (1)$$

Thus, $R[I](x, y)$ represents the average gray-level computed on a $\delta \times \delta$ square window of the original image having the top-left corner at $(\delta \times x, \delta \times y)$. The next step of the transformation is to obtain two additional $\frac{w}{\delta} \times \frac{h}{\delta}$ images defined as:

$$D_x[I](x, y) = R[I](x + 1, y) - R[I](x, y) \quad (2)$$

$$D_y[I](x, y) = R[I](x, y + 1) - R[I](x, y), \quad (3)$$

with all the elements of the last column of $D_x[I]$ and last row of $D_y[I]$ set to zero. $D_x[I]$, $D_y[I]$ are simply the horizontal and vertical derivatives of the reduced-resolution image $R[I]$, and the pair $D_x[I]$, $D_y[I]$ forms the *structure* of the original gray-level image I . Roughly, the averages computation followed by a reduced-resolution differentiation resembles a bandpass filtering operation. It is worth pointing out that only the first step of the transformation works on a full resolution image, while successive operations require processing and storing significantly smaller images (i.e. reduced by a factor δ^2). To discuss the stability of the proposed representation we model an illumination change as a global, linear transformation of the intensities: $k_1 I + k_2$. It can be verified very easily that, under this assumption, the resulting structure change is given by $(k_1 - 1)D_x[I](x, y)$, $(k_1 - 1)D_y[I](x, y)$. Hence, where the original image is uniform the corresponding structure elements will be small and thus substantially unaffected by the illumination change. On the other hand, the structure sensitivity to the illumination change gets higher at larger structure elements, that correspond to windows of the original image located near or across intensity edges. Although this analysis derives from a simplified assumption, it provides a correct indication of the behaviour of the proposed representation under typical, real illumination changes, such as brightness changes. Moreover, as it will be shown later, we have found that in real video sequences the structure variations produced by illumination changes are usually much



Figure 1. Structure difference between two images.

smaller than those caused by true “structural” changes of the scene (i.e. those due to object activities). Consequently, even though the proposed representation is not fully unaffected by illumination changes, it exhibits a strong ability to discriminate between structural changes and illumination changes. To obtain a structural representation in the case of colour images we simply apply the described transformation to each of the three RGB channels. In this manner we obtain six $\frac{w}{\delta} \times \frac{h}{\delta}$ images that form the *structure* of the colour image I : $D_{x,r}[I]$, $D_{x,g}[I]$, $D_{x,b}[I]$, $D_{y,r}[I]$, $D_{y,g}[I]$, $D_{y,b}[I]$ (the second subscript identifies the colour channel). A simple way to compare the structures of two images, I_α and I_β , is to compute a *delta-structure* image, $\Delta S[I_\alpha, I_\beta]$, where each element, $\Delta S[I_\alpha, I_\beta](x, y)$, is defined as:

$$\max_{\substack{ch \in \{r, g, b\} \\ d \in \{x, y\}}} \left\{ |D_{d,ch}[I_\alpha](x, y) - D_{d,ch}[I_\beta](x, y)| \right\} \quad (4)$$

The top row of Figure 1 shows two images taken by a static camera at two different times and with two very different settings of the frame-grabber parameter that controls image brightness. The left image of the bottom row shows the *delta-structure* image (scaled-up to the original size) associated with the two images of the top row. As the elements of the *delta-structure* image gets larger the colours vary from blue to yellow (from dark to light if the paper is printed in gray-scale). Let’s consider background regions first: we can notice how the structure differences caused by the brightness change are very small in uniform areas and stronger where the image is richer of spatial variations (e.g. the portion of the background containing the parked cars). This is consistent with the previous theoretical considerations. Furthermore, we can notice that the structure differences due to the different foreground (the left image contain two persons in foreground while the right image do not contain foreground objects) are much stronger than those caused by the brightness change within background regions.

Hence, a thresholding operation on the *delta-structure* image (right image of the bottom row) allows for discriminating between the true scene changes and the structure variations caused by the brightness change.

Of course, the detection capability of the proposed structural representation is exalted for high-contrast objects, such as the dark people silhouetted against the brighter background in Figure 1. Objects with low contrast boundaries will yield less compelling edge information, particularly after greatly reducing the image resolution. However, it must be pointed out that in these cases structure differences between object and background in the internal part of the object’s silhouette can still be detected. Then, as described in detail in the next section, by combining structure and pixel-level information a single structure element detected as different can be enough to detect the whole object’s silhouette. For example in the last snapshot of Figure 2, we can see on the left a car with smooth boundaries and colour quite similar to the background. Nonetheless, as shown on the right, a few internal structure differences are enough to validate the two blobs detected at pixel level.

However, brightness variations are only a subset of all the possible image variations introduced by illumination changes. For example, during the day, the variation of the light inclination modifies slowly the position of the shadows cast by static object. These kind of variations cause a slow, regular movement of the edges associated with shadows and thus clearly affect the proposed structural representation. Similar consideration apply for some sudden illumination changes: when a cloud appears in a sunny day static shadows can completely disappear, thus changing significantly image *structure*. However it seems impossible for a low-level representation to possess the ability to distinguish between these of kind of illumination changes and the true structure changes due to objects. Therefore, there exists the possibility that our method treats as changes due to objects some structure modifications caused by illumination changes. Even so, our approach has shown good global performances. In fact these kind of variations are relatively occasional and generally produce local changes that affect only a limited area of the background. Moreover, as it will be shown later, slow structural variations, such as those due to the movement of the shadows cast by static objects, are filtered away by the overall change-detection process.

5 The overall Algorithm

A change detection process based solely on structure would detect changes with a very low resolution, yielding inaccurate blobs such as those shown in the bottom-right image of Figure 1. This is due to the scaling operation that reduces both structure dimensions by a factor δ with respect to the original image. Indeed, structure owes a great deal of

its robustness to the scaling operation. To overcome this problem we adopt a background model made out of two separate parts: the first is the structure of the scene without moving objects while the second is simply a colour image of the reference scene. The first part of the model will be referred to as background structure, and its components indicated as $BD_{x,r}$, $BD_{x,g}$, $BD_{x,b}$, $BD_{y,r}$, $BD_{y,g}$ and $BD_{y,b}$; the second part of the model will be referred to as background image, and its components indicated as B_r , B_g , B_b . Considering the properties of structure discussed in section 4, the background structure results largely unaffected by illumination changes, so that its variations can be ascribed to objects activities. The background image provides the algorithm with the capability of detecting changes at the highest possible resolution (i.e. the same resolution as the original image). This enables the algorithm to yield blobs following accurately the true objects silhouettes. Given the described background model, at each new frame the detection process operates at both the structure and image level; then, the information provided by the two detection operations is combined adequately to obtain the final output. Since structure variations can be largely ascribed to objects, the updating process of the background structure is focused on handling this type of changes, with a simple mechanism allowing a feedback from a higher-level processing module to control the process. On the other hand, since the background image is affected by all scene changes (illumination changes as well as changes due to objects) it needs to be updated continuously. Before starting the detection process, our algorithm activates a simple "bootstrap" process aimed at estimating the initial data to be included into the background model.

Initialization of the Background Model. This step consists mainly in estimating the initial background structure and lasts a user-selectable time period (we typically set it to 40 seconds). During this period, image structure is computed and stored at constant intervals (we typically take a new structure every second). A motion estimation operation is also performed so as not to include into the initial background structure the information associated with moving objects. When all the samples have been collected, the background structure is initialised by taking for each element of the six background structure components the value observed most frequently. Once the background structure is initialised, the background image is initialised by simply copying the current colour image into B_r , B_g , B_b . In fact, as it will be shown later, even though the current image contains moving objects, the correct background image will be established automatically after a short delay thanks to the information contained in the background structure.

Structure-Level Detection. We compare the structure of the current frame, I , with the background structure by building up two *delta-structure* images associated respectively with

the x and y directions:

$$\Delta S_x[I](x,y) = \max_{ch \in \{r,g,b\}} \left\{ |BD_{x,ch}(x,y) - D_{x,ch}[I](x,y)| \right\} \quad (5)$$

$$\Delta S_y[I](x,y) = \max_{ch \in \{r,g,b\}} \left\{ |BD_{y,ch}(x,y) - D_{y,ch}[I](x,y)| \right\} \quad (6)$$

Then, choosing a suitable threshold value, t_s , and recalling equation (2), we can observe that if $\Delta S_x[I] > t_s$ at structure element (x, y) , then a foreground object occupies (or a background object leaves) the image region associated with structure element (x, y) or $(x + 1, y)$, or both. Similarly, recalling equation (3), if $\Delta S_y[I] > t_s$ at structure element (x, y) , the structure change could be located at element (x, y) or $(x + 1, y)$, or both. Therefore, we define the $\frac{w}{8} \times \frac{h}{8}$ binary image S_{mask} containing the structure-level detection results as:

$$S_{mask}(x, y) = \begin{cases} 1, & \text{if } \left(\Delta S_x[I](x, y) > t_s \right) \vee \\ & \left(\Delta S_x[I](x - 1, y) > t_s \right) \vee \\ & \left(\Delta S_y[I](x, y) > t_s \right) \vee \\ & \left(\Delta S_y[I](x, y - 1) > t_s \right); \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

where the second OR condition is ignored at $x = 0$, and the fourth is ignored at $y = 0$. Our experiments have shown that the best results are typically achieved using low t_s values, so as to maximize the sensitivity of the structure-level detection process. In addition, our experiments have demonstrated also that once the threshold t_s has been set, it can be used throughout system's working time without the need for any adaptation. Hence, in our current implementation t_s can be chosen by the user at run-time, by simply looking at the output provided by the algorithm.

Image-Level Detection. In this case the detection step is simpler and consists in computing the difference between I and the background image. Hence, calling I_r, I_g, I_b the colour channels of the current frame and t_p a new threshold value, we define the $w \times h$ binary image I_{mask} containing the image-level detection results as:

$$I_{mask}(x,y) = \begin{cases} 1, & \text{if } \left(|I_r(x, y) - B_r(x, y)| > t_p \right) \vee \\ & \left(|I_g(x, y) - B_g(x, y)| > t_p \right) \vee \\ & \left(|I_b(x, y) - B_b(x, y)| > t_p \right); \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

Our experiments have shown that, similarly to t_s , t_p do not require any adaptation to enable long-term functioning of the system.

Combination of the Detection Results. Following structure and image level detection, the information contained in S_{mask} is used to decide whether or not each of the blobs detected in I_{mask} is valid. First some simple morphological

operations, such as erosions and dilations, are executed on I_{mask} to clean-out very small blobs. Then the validation is done by labelling the connected components of I_{mask} and erasing those not intersecting at least one structure-element marked with 1 in S_{mask} . The result of the combination step, representing the final output of the overall change-detection system, is a $w \times h$ binary image, $Mask$, that contains only the blobs associated with objects changing the image structure.

Updating of the Background Structure. If at a certain structure element the difference between the background structure and the current image structure is persistently above the threshold t_s , the system estimates the value that should be assigned to the element to absorb the change. The estimation process consists in observing the temporal behaviour of the above-threshold structure element until it exhibits a stable value (i.e. a value approximately constant for a time interval longer than the parameter τ). Only when this occurs, the stable value is considered as a reliable estimate and may be used to update the background structure. In fact, the updating to the estimated value must be enabled explicitly by a feedback coming from a higher-level module. The idea is that the higher-level module, analysing the output provided by the change-detection algorithm, specifies which objects are “not-interesting”, and thus can be inserted into the background structure so as not to be detected in the successive frames, and which ones are “interesting”, and thus should not be included into the background structure so as to be detected also in the next frames. To implement this feedback, the change-detection system can receive a binary image FB , as large as $Mask$, in which the higher-level module should redraw the blobs associated with “interesting” objects. Then, the structure elements intersecting at least one blob drawn in FB will not be updated, even though a reliable estimated value might be available. Conversely, if a structure element exhibiting a persistent change does not intersect any blob drawn in FB and a reliable estimated value is available, the estimated value is copied into the background structure. In absence of any feedback from higher-level modules (i.e. the algorithm receives always an empty FB image), the background structure updating turns out to be dependant only on the parameter τ . In this case, an above-threshold value approximately constant for a time interval longer than τ will be copied into the background structure. This means that, in absence of a feedback from a higher-level module, the objects remaining motionless for a time interval longer than τ will be included into the background structure and thus no longer detected.

Updating of the Background Image. Indicating the background image before and after the updating as B^t and B^{t+1} respectively, and considering the red channel, the updating

rule is given by:

$$B_r^{t+1}(x, y) = \begin{cases} B_r^t(x, y) + \eta, & \text{if } (Mask(x, y) = 0) \wedge \\ & (B_r^t(x, y) < I_r(x, y)); \\ B_r^t(x, y) - \eta, & \text{if } (Mask(x, y) = 0) \wedge \\ & (B_r^t(x, y) > I_r(x, y)); \\ B_r^t(x, y), & \text{otherwise.} \end{cases} \quad (9)$$

where η is a constant value. The same rule is used also for the green and blue channels. With this strategy, the updating of B does not need to be executed at each frame; for example in our experiments we chose $\eta = 3$ and the updating is done three times per second. If we increase the updating frequency or η , we augment the ability of the background image to follow sudden illumination changes. Occasionally, the updating process may erroneously inserts some object parts into the background (typically at objects’ boundaries). However, these errors won’t produce false positive into the final output thanks to the structure information enclosed in S_{mask} . In fact, generally, when I_{mask} contains a false positive caused by an error into the background image, the corresponding structure-element into the background structure turns out to be correct. Hence, no change will be detected in S_{mask} and the final combination step will filter-out the false positive. This will lead also to a fast recovering of the error in the background image. Eventually, we point out that if an object is inserted into or removed from the background structure, it will be automatically inserted into or removed from the background image. Recalling the final remarks of section 4, we explain here why slow structural variations misclassified as changes due to objects (e.g. those due to the movement of the shadows cast by static objects) are filtered away by the change detection process. In these situations the structure changes might result above-threshold, but generally the image changes are slow and affects a few neighbouring pixels at a time. Thus, the image changes are absorbed by the background image updating process and consequently no blob is produced in the final output. This implies also that a higher-level module won’t observe any blob in the image area affected by the change and hence won’t disable the background structure updating: as a result, the change will be correctly absorbed into the background structure.

6 Experimental Results

In this paper we report on the experiments aimed at evaluating the change-detection algorithm itself, i.e. apart from the particular higher-level module supplying the feedback information. Hence, in these experiments the feedback action was disabled, and τ was set to 100 seconds. The video sequences used to evaluate the algorithm were acquired with different cameras, in outdoor environments

and under various weather conditions (clear sky, overcast, changeable weather...). To evaluate the robustness with respect to long-term functioning, in all the experiments the algorithm ran continuously from morning to evening, with all the parameters set at initialisation time and left unchanged. These experiments showed that our algorithm can assure a good quality output all-day long, needing neither re-initialisations nor parameter adjustments. Objects silhouettes are detected quite accurately, taking into account that the current version of the algorithm do not include any shadow-removal operation. The output is characterized by a remarkably rare occurrence of false positives. Consequently, the background updating scheme conceived to handle the changes due to objects actually deals almost always with these class of changes, turning out very rarely as instrumental to the recovering from persistent false positives (recall section 3). This means that, in practise, the usual trade-off problem pointed out in section 3 is nearly absent in our algorithm. As for speed, our algorithm works at video rate (25 fps) with 24 bit RGB, 320x240 pixels images on a standard PC equipped with a 1GHz Pentium 4 CPU. In the remainder of this section we will discuss in more detail some experimental results that highlight the different, and better, behaviour of our change-algorithm with respect to the conventional algorithms affected by the trade-off problem. Figure 2 shows several frames of a sequence containing a car that reduces its speed, stops and successively goes away. Since, the colour of the car is similar to the background colour this sequence is also affected by the camouflage problem. The time elapsed between the first two snapshots is 4 seconds. The output yield by our algorithm, on the right of the figure, shows that even though the car moves very slowly, it doesn't taint the background since its blob is detected and no ghosts are generated. As can be seen in the third snapshot, the car rests almost motionless for 20 seconds and its blob is always correctly detected. Finally the car goes away and, as can be seen in the last snapshot, no ghost appears in the output even though during the stop the illumination undergoes a slight change. This is due to the robustness of the background structure with respect to illumination changes. As discussed in section 3, with a conventional algorithm and a fast time constant, the slow car would taint the background, causing the disappearance of the car's blob and the appearance of ghosts. Conversely, with a slow time constant, the car wouldn't taint the background, but the area covered by the car wouldn't be updated during the illumination change, so that when the car moves away false positives would be detected. Moreover, the slow time constant would lead to a slow elimination of these false positives.

Figure 3 shows several frames of a sequence in which an object remains motionless for a time interval longer than τ (i.e. 100 sec). In the first snapshot, the algorithm detects a

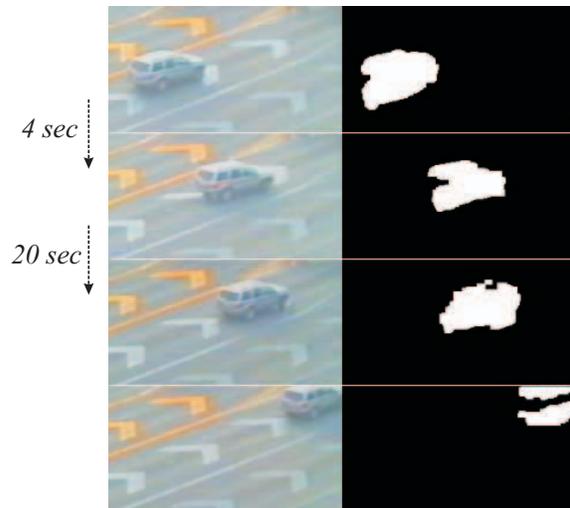


Figure 2. A sequence with a slow car.

passing motorcycle. The second snapshot has been taken nearly after one minute. As can be seen, a passing cloud caused a significant illumination change, but thanks to the robustness of background structure, no false positives are detected: the output contains a pedestrian and a car approaching to a parking-lot. The output associated with the third snapshot shows the parked car. After nearly two minutes (fourth snapshot), the blob of the parked car is still largely detected. The time elapsed between the fourth and the fifth snapshot is only 4 seconds. Now the car, which remained motionless for more than τ , is included in the background structure and its blob is no longer detected. We point out that, by exploiting the feedback action supported by the algorithm and using a lower τ , the parked car could be removed instantaneously "upon-demand" of a higher-level module. With a traditional algorithm, the blob of the parked car could be detected persistently only using a very slow time constant. Unfortunately, this would cause also a very slow recovery from the potential errors due to the various illumination changes present in the sequence. It is worth highlighting that the sequence contain a sudden change, which, for a conventional algorithm, is hard to compensate with a slow time constant. The last snapshot has been taken after nearly an hour. As can be noticed, the shadows cast by static objects have moved, changing the background structure. Since these are slow variations, as explained at the end of section 5, they are automatically filtered-out by the system and no false positives affect the output.

7 Conclusion

Conventional change-detection algorithms exhibit a somewhat subtle, beneficial interdependency between the updating of the background model to include/remove objects and the recovering from persistent false positives. This interdependency involves setting a parameter that determines the time needed for a still foreground pixel to be absorbed into the background. However, the choice of this parameter implies a trade-off between the ability to handle adequately still and slow objects and that of quickly recovering from false positive errors, the latter being fundamental to ensure proper long-term functioning of the algorithm.

Conversely, the novel change-detection algorithm proposed in this paper is largely unaffected by the above described trade-off and hence, as shown in the experimental results, can handle properly still and slow objects without any depletion of its ability to work correctly over very long time spans (e.g. all-day long). This has been achieved by incorporating into the background model a set of simple low-level features that captures effectively structural (i.e. robust with respect illumination changes) information concerning the observed scene. Since the background model contains a basic core quite stable with respect to illumination changes, it turns out that the background is updated mostly to include/remove objects and that the handling of these events is very rarely instrumental to the recovering from persistent false positives. Besides, since in principle our algorithm relies on a background model that needs to be updated only to handle “high-level” events, namely insertion/removal of objects, it naturally holds the potential to interact smoothly with higher-level processing module that, on the basis of scene comprehension, may control the background maintenance process flexibly and intelligently.

References

- [1] M. Boninsegna and A. Bozzoli. A tunable algorithm to update a reference image. *Signal Processing: Image Communication*, 16(4):353–365, 2000.
- [2] A. Cavallaro and E. Touradj. Change detection based on color edges. *IEEE ISCAS*, 2(10):141–144, 2001.
- [3] A. Elgammal, R. Duraiswami, D. Harwood, and L. Davis. Background and foreground modeling using nonparametric kernel density estimation for visual surveillance. *Proc. of the IEEE*, 90(7), 2002.
- [4] I. Haritaoglu, D. Harwood, and L. Davis. W4 who? when? where? what? a real time system for detecting and tracking people. In *Int. Conf. on Automatic Face and Gesture Recognition*, 1998.
- [5] S. Jabri, Z. Duric, H. Wechsler, and A. Rosenfeld. Detection and location of people in video images using adaptive fusion of color and edge information. In *Int. Conf. on Pattern Recognition*, Barcelona, Spain, 2000.

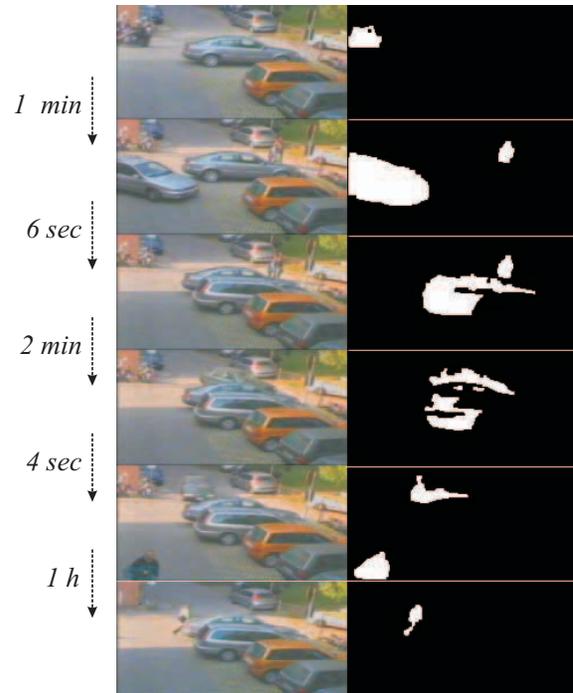


Figure 3. A sequence with a parking vehicle.

- [6] O. Javed, K. Shafique, and M. Sha. A hierarchical approach to robust background subtraction using color and gradient information. In *Proc. of the Workshop on Motion and Video Computing, IEEE*, 2002.
- [7] L. Li and M. Leung. Integrating intensity and texture differences for robust change detection. *IEEE Trans. on Image Processing*, 11(2):105–112, 2002.
- [8] C. Stauffer and W. Crimson. Adaptive background mixture models for real-time tracking. *Int. Conf. CVPR*, pages 246–252, 1999.
- [9] K. Toyama, J. Krumm, B. Brumitt, and B. Meyers. Wallflower: Principles and practice of background maintenance. In *ICCV'99*, volume 1, pages 255–261, 1999.
- [10] C. Wren and et al. Pfunder: Real-time tracking of the human body. *IEEE PAMI*, 19(7), 1997.