IMPROVED BLOTCH DETECTION BY POSTPROCESSING

J. Biemond, P.M.B. van Roosmalen, and R.L. Lagendijk

Information and Communication Theory Group
Faculty of Information Technology and Systems
Delft University of Technology, The Netherlands

ABSTRACT

Blotches are common artifacts in old film sequences that manifest themselves as disturbing bright or dark spots. Existing methods for detecting blotches can achieve high detection rates. High detection rates are only useful if the corresponding number of false alarms is not too high, visible artifacts in the corrected sequence result otherwise. We show that the performance of blotch detectors can be improved significantly by taking statistical influence of noise on the detection mechanism into account. Further improvements are achieved first by using a double-stage detection strategy and second by a constrained dilation technique.

1. INTRODUCTION

Blotches present a common type of artifact in old film sequences that manifests itself as disturbing bright or dark spots caused by dirt and by the loss of the gelatin covering the film due to ageing effects and bad film quality. Characteristics of blotches are that they seldom appear at the same spatial location in consecutive frames, they tend to be smooth (little texture), and they usually have intensity values that are very different from the original contents they cover. Films corrupted by blotches are often restored using a two-step approach. In the first step blotches are detected and detection masks are generated that indicate for each pixel whether or not it is part of a blotch. In the second step, corrupted pixels are corrected by means of spatio-temporal interpolation [1-3].

Blotch detectors are either object based or pixel based. Pixel based detectors determine for each pixel whether or not it is part of a blotch independently from whether or not its neighboring pixels are considered to be part of a blotch. Object based detectors exploit the spatial coherence within blotches via, e.g., markov random fields. So far, pixels based detectors have shown to achieve similar detection results as object based detectors at a fraction of the computational cost [1].

In every detection problem there is a trade off between the probability of correct detection and the probability of false alarm. Obviously, a blotch detector will not be set so that it generates too many false alarms because the interpolator in the correction stage is fallible. New visible artifacts that are more disturbing than the blotches themselves are introduced into the corrected sequence otherwise. On the other hand, when setting a detector to a lower detection rate, many blotches will be detected only partially or not at all. Two major causes for false alarms are ever present noise and inaccurate motion estimation.

Section 2 of this paper describes the blotch detector we will be using. Section 3 presents three postprocessing operations that can be applied on the candidate blotches output by a blotch detector to improve the quality of the detection masks. The key is that we use a pixel based detector and that in the postprocessing we consider blotches as objects. This allows us to exploit the spatial coherency within blotches while maintaining low complexity and low computational effort. The first postprocessing operation detects and removes possible false alarms by taking into account the probability that the detector wrongly detects a blotch (an object) due to noise. The second postprocessing operation applies a double-stage detection mechanism that finds missing pieces of blotches that are detected only partially otherwise. The final postprocessing operation consists of a constrained dilation operator that fills small *holes* in and on edges of the candidate blotches. Section 4 describes the results and concludes this paper.

2. THE SIMPLIFIED *RANKED ORDER DIFFERENCE (ROD)* DETECTOR

Blotches are characterized by the fact that they seldom appear at the same location in a pair of consecutive frames and that they have intensity values different from the original image contents. Therefore blotches can be detected by detecting temporal discontinuities in image intensity. The additional use of motion compensation significantly reduces the number of false alarms. The *ROD* detector [2] is based on these principles. This detector has three free parameters, T_1, T_2, and T_3, to control the trade off between the number of correct detections and false alarms. By letting $T_2, T_3 \rightarrow \infty$, *ROD* can be simplified to what we call *simplified ROD* (*S-ROD).*

Let $I_n(\mathbf{z})$ denote the intensity of a pixel at a spatial location $\mathbf{z}^T = (x, y)$ in frame n. Let $r_{n,i}(\mathbf{z})$ form a set of six reference pixels obtained from spatially co-sited pixels and their vertical neighbors in motion compensated previous and next frames (see Fig. 1). *S-ROD* is then defined by:

$$d_n(\mathbf{z}) = \begin{cases} min(r_{n,i}(\mathbf{z})) - I_n(\mathbf{z}) & \text{if } min(r_{n,i}(\mathbf{z})) - I_n(\mathbf{z}) > 0 \\ I_n(\mathbf{z}) - max(r_{n,i}(\mathbf{z})) & \text{if } I_n(\mathbf{z}) - max(r_{n,i}(\mathbf{z})) > 0, \\ 0 & otherwise \end{cases} \quad (1)$$

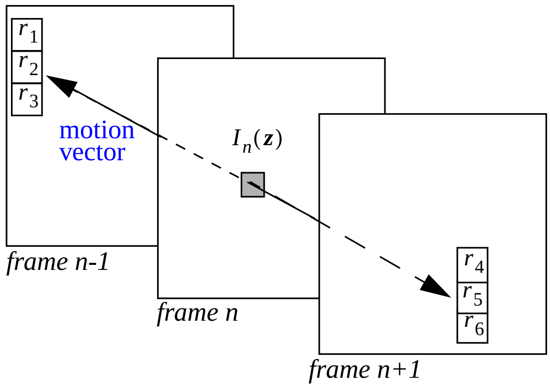

Figure 1. Selection of reference pixels $r_{n,i}(\mathbf{z})$ from previous and next frames using motion compensation.

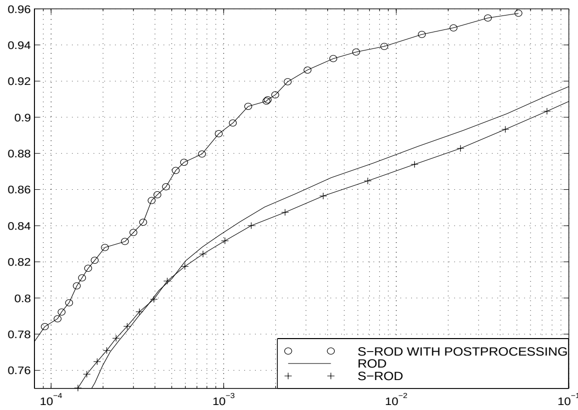

Figure 2. ROC-curves resulting from ROD, S-ROD, and S-ROD with postprocessing applied to the test sequence.

and a blotch is detected when:

$$d_n(\mathbf{z}) > T_1 \qquad \text{with } T_1 \geq 0 . \tag{2}$$

S-ROD basically looks at a range of pixel intensities obtained from motion compensated frames and compares this range to the pixel intensity under investigation. A blotch is detected if the intensity of the current pixel lies far enough outside that range. What is considered *"far enough"* is determined by T_1. If T_1 is small many blotches will be detected correctly but many false alarms will occur. As T_1 becomes larger, fewer blotches are detected and the number of false alarms drops.

Figure 2 shows the *receiver operator characteristic* (ROC) curves obtained from the *Western* test sequence using *ROD*, *S-ROD* and *S-ROD* with the postprocessing proposed in the next section. The *Western* sequence, which was also used in [1,2], is a sequence (64 frames) to which artificial blotches have been added. Each artificial blotch had a fixed gray value which was drawn uniformly between 0 and 255. We observe that the performance of S-*ROD* is slightly below that of *ROD*. We also see that much is gained after applying the postprocessing operations we describe next.

3. IMPROVING THE DETECTION RESULTS BY POST-PROCESSING

We propose a number of postprocessing operations of which the goal is to maximize the ratio of correct detections and false alarms. The key is that the candidate blotches are viewed as objects and not as individual pixels. Section 3.1 defines what we consider to be an object. Section 3.2 presents the first postprocessing technique by deriving how to detect and remove possible false alarms due to noise given a specific detector. Often blotches are detected only partially. Section 3.3 presents our second postprocessing technique that finds more complete blotches in those cases. Section 3.4 introduces a constrained dilation technique that includes small holes on the edges of and in candidate blotches as our third and final postprocessing operation.

3.1 Object definition

We want to manipulate candidate blotches as objects rather than as individual pixels. Because we are particularly interested in blotches it is reasonable to use characteristics of blotches in the object definition. The characteristic we use is the fact that blotches tend to be smooth, i.e. that adjacent pixels have similar intensities. We consider a pair of pixels to be similar if their difference is smaller than twice the standard deviation of the noise. Other characteristics of blotches are taken into account implicitly due to the fact we are only interested in pixels flagged by the blotch detector.

Therefore, adjacent pixels that have similar intensities and that are flagged by the blotch detector are considered to be part of the same candidate blotch. To differentiate between the various candidate blotches a unique label is assigned to each candidate blotch.

3.2 Removing false alarms due to noise

High correct detection rates are achieved by setting the blotch detector to a high degree of sensitivity. However, the detector is then not only sensitive to blotches but also to noise and many false alarms result. We propose computing the probability that the detector gives a specific response, i.e. that a specific set of values $d_n(\mathbf{z})$ results for a candidate blotch, due to noise. If the computed probability exceeds a certain risk R, the candidate blotch is removed from the detection mask.

We demonstrate this approach for the *S-ROD* detector. Given T_1, the probability that *S-ROD* generates a single false alarm due to noise equals:

$$\begin{aligned} P[d_n(\mathbf{z}) > T_1] &= \\ &= P[I_n(\mathbf{z}) - max(r_{n,i}(\mathbf{z})) > T_1, I_n(\mathbf{z}) - max(r_{n,i}(\mathbf{z})) > 0] + \\ &\quad P[min(r_{n,i}(\mathbf{z})) - I_n(\mathbf{z}) > T_1, min(r_{n,i}(\mathbf{z})) - I_n(\mathbf{z}) > 0] \\ &= P[I_n(\mathbf{z}) - max(r_{n,i}(\mathbf{z})) > T_1] + P[min(r_{n,i}(\mathbf{z})) - I_n(\mathbf{z}) > T_1] \end{aligned} \tag{3}$$

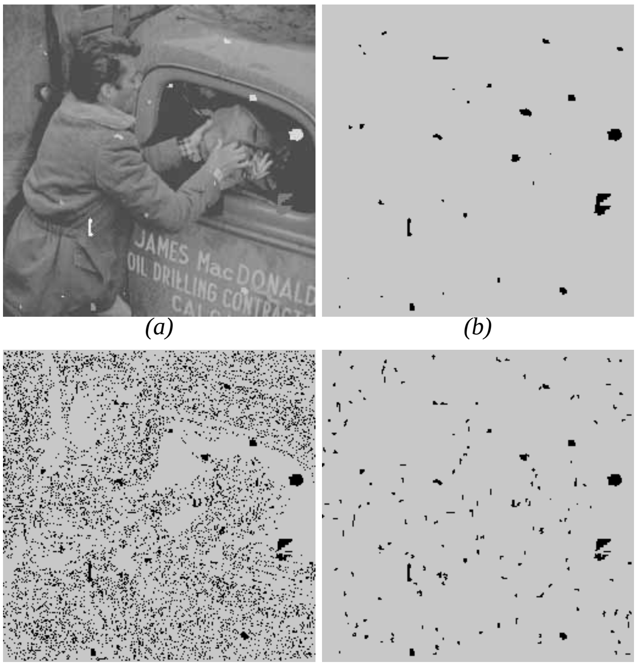

(a) (b) (c) (d)

Figure 3. (a) Blotched frame from test sequence. (b) Mask of artificial blotches. (c) Initial detection mask using S-ROD with $T_1 = 0$. (d) Detection mask after removing possible false alarms due to noise.

Let us assume that the intensities of the reference pixels are equal to the intensity of the pixel under observation in the absence of blotches and noise and let us also assume that the noise is additive, white and gaussian. Then, based on (3), it is easy to compute the probability mass function $P[d_n(\boldsymbol{z}) = X(\boldsymbol{z})]$, i.e. the probability that *S-ROD* gives a specific response $X(\boldsymbol{z})$ for a single pixel at location $\boldsymbol{z}$ due to noise.

After the labelling procedure, a candidate blotch is an object with spatial support S that consists of N pixels each of which with a specific detector output $d_n(\boldsymbol{z})$. Let H_0 denote the hypothesis that this object is purely the result of false alarms due to noise. $P[H_0]$ is then the probability that a collection of N individual pixels are flagged by the S-ROD independently of their location and of their neighbors:

$$P[H_0] = \prod_{\boldsymbol{z} \in S} P[d_n(\boldsymbol{z}) = X(\boldsymbol{z})] \,. \tag{4}$$

We now remove those candidate blotches for which the probability that they are solely the result of noise exceeds a risk R:

$$P[H_0] > R \,. \tag{5}$$

$d_n(\boldsymbol{z})$	probability of false alarm	$d_n(\boldsymbol{z})$	probability of false alarm
1	0.091921	7	0.002224
2	0.060748	8	0.000892
3	0.036622	9	0.000304
4	0.020492	10	0.000108
5	0.010353	11	0.000028
6	0.005168	12	0.000008

Table 1. See text for explanation.

Figure 3 illustrates the result of this approach. It shows frame 8 of the *Western* sequence, the artificial blotch mask and the detection masks before and after postprocessing. The initial detection mask was obtained using *S-ROD* where we let $T_1 = 0$. The noise was assumed to be *i.i.d.* gaussian and the noise variance was estimated to be 9. Table 1 shows the probability of specific detector responses due to noise for single pixels, again assuming that no blotches are present and that the reference pixels differ from the pixel under observation by the noise term only.

In this frame 85.7% of the blotches were detected correctly and 12.9% of the uncorrupted pixels were mistakenly flagged as being part of a blotch before postprocessing. After postprocessing, where we set $R = 10^{-5}$, 85.1% of the blotches were detected correctly and only 1.1% of the clean pixels were mistakenly flagged as being part of a blotch. Clearly this is a large improvement.

3.3 Completing partially detected blotches

The technique for removing possible false alarms due to noise can be applied to any operator setting for T_1. Of course, this method is most effective for low values of T_1, i.e. when the detector is set to a high detection rate. When a blotch detector is set to a low detection rate less gain is to be expected from this strategy. A second strategy for improving the ratio between correct detections and false alarms is described here.

We note that at lower detection rates many blotches are not detected at all and other blotches are detected only partially. Our goal is to make those blotches that are detected only partially more complete. We achieve this by noting from Fig. 2 that as T_1 is lowered the probability of false alarms decreases faster than the probability of correct detections. This means that detections resulting from a blotch detector set to a low detection rate are more likely to be correct and can thus be used to validate the detections from that detector when set to a high detection rate.

This can be implemented by applying hysteresis thresholding [4]. The first stage computes and labels the set of candidate blotches using the operator settings of the blotch detector (in the case of S-ROD this is the operator setting of T_1). Possible false alarms due to noise are removed as described before. The second stage sets the blotch detector to a very high detection rate (i.e., $T_1 = 0$ for S-ROD) and again a set of candidate blotches is computed and labeled. Candidate blotches from the second set can now be validated; they are preserved if corresponding candidate blotch in the first set exist. The other candidate blotches in the second set, which are more likely to have resulted from false alarms, are discarded. Effectively the candidate blotches detected using the operator settings are preserved and are made more complete.

3.4 Constrained dilation for missing details

There is always a probability that a detector fails to detect elements of a blotch. This is illustrated by Fig. 3 where it can be seen that, even though the *S-ROD* detector has been set to its most sensitive setting, not all the blotches have been detected completely. In this final postprocessing step we refine the candidate blotches by removing small *holes* in and on the edges of the candidate blotches.

We propose using a constrained dilation operation for filling in the holes. It applies the following rule: if a pixel's neighbor is flagged as being blotched and its intensity difference with that neighbor is small (e.g., twice the standard deviation of the noise) then that pixel should also be flagged as being part that blotch. The constraint on the differences in intensity reduces the probability that uncorrupted pixels surrounding a blotch are mistakenly flagged as "blotched" because blotches tend to have gray values that are significantly different from their surroundings.

However, It is important not to apply too many iterations of this constrained dilation operation because it is always possible that the contrast between a candidate blotch and its surrounding is low. The result would be that the candidate blotch would grow completely out of its bounds and many false alarms would occur. In practice we found that applying two iterations leads to good results.

4. RESULTS AND CONCLUSION

We already observed in Figure 2 that *S-ROD* in combination with postprocessing is a significant improvement over plain *ROD*. For example, the number false alarms resulting from *S-ROD* with postprocessing is a factor 4.3 lower than that what results from *ROD* at a correct detection rate of 85%.

Figure 4 shows the detection masks that result form S-ROD and S-ROD with postprocessing. Also shown are the artificial mask and the corrected frame. Note the differences in the detection results. Also compare the impaired image in Fig. 3a to the corrected result. The interpolation method as described in [3] was used for interpolating the missing data. These results were obtained by setting the overall false alarm rate to 10^{-4} for the whole test sequence.

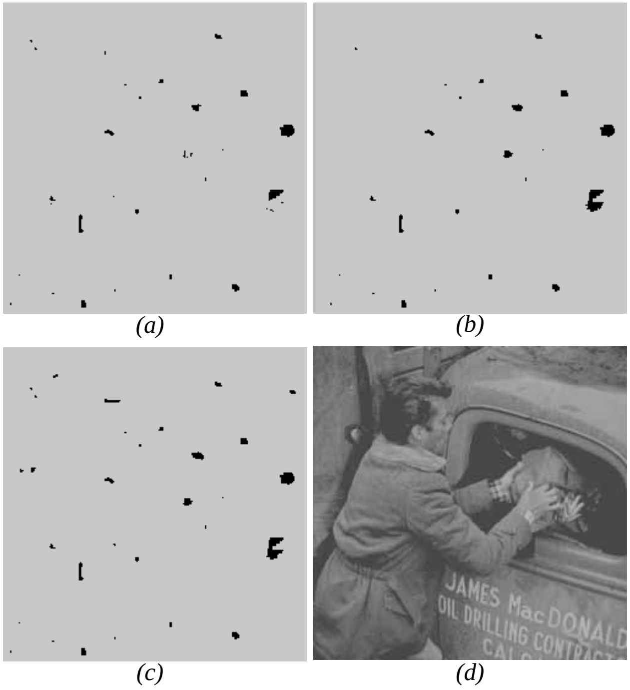

Figure 4. (a) Blotches detected using S-ROD. (b) Blotches detected using S-ROD with postprocessing. (c) Artificial blotch mask. (d) Frame corrected after S-ROD with postprocessing.

In conclusion, the methodology described here can also be applied to other blotch detectors. Alternatively, the rules and constraints posed by the postprocessing could well be defined implicitly in a new detector, e.g. based on markov random fields. This would, however, significantly increase both the complexity and the computational effort.

REFERENCES

1 *A.C. Kokaram, R.D. Morris, W.J. Fitzgerald, and P.J.W. Rayner, "Detection of Missing Data in Image sequences", IEEE Trans. on Image Processing, Vol. 4, No. 11, 1995.*

2 *M.J. Nadenau, and S.K. Mitra, "Blotch and Scratch Detection in Image Sequences based on Rank Ordered Differences", in Time-Varying Image Processing and Moving Object Recognition, V. Cappelini ed., Elsevier, 1997.*

3 *P.M.B. van Roosmalen, A.C. Kokaram, and J. Biemond, "Fast High Quality Interpolation of Missing Data in Image Sequences using a Controlled Pasting Scheme", Submitted to ICASSP99.*

4 *J. Canny, "A Computational Approach to Edge Detection", IEEE PAMI, Vol. 8, No. 6, 1986.*