Blotch Detection for Digital Archives Restoration based on the Fusion of Spatial and Temporal Detectors

3 authors, including:

Louis Laborelli
Institut National de l'Audiovisuel, Bry-sur-Marne, France
18 PUBLICATIONS 227 CITATIONS

Isabelle Bloch
Télécom ParisTech
566 PUBLICATIONS 9,871 CITATIONS

Some of the authors of this publication are also working on these related projects:

Problems at the interface of artificial vs natural cognition View project

HUMN project initiative: review and validation, quality control and prospects for development of automated micronucleus assays using image cytometry systems View project

Blotch Detection for Digital Archives Restoration based on the Fusion of Spatial and Temporal Detectors

Sorin Tilie
DRE, INA
4 av. de l'Europe,
Bry sur Marne, France.
stilie@ina.fr

Louis Laborelli
DRE, INA
4 av. de l'Europe,
Bry sur Marne, France.
llaborelli@ina.fr

Isabelle Bloch
GET-ENST, TSI
CNRS UMR 5141
46 rue Barrault,
75013 Paris, France.
Isabelle.Bloch@enst.fr

Abstract - *This paper proposes a method based on the Dempster-Shafer evidence theory for the detection of blotches in digitized archive film sequences. The detection scheme relies on the fusion of two uncorrelated fast, no motion compensated, spatio-temporal blotch detectors. The imprecision and uncertainty of both detectors are modeled using Dempster-Shafer evidence theory, which improves the decision, by taking into account the ignorance and the conflict between detectors. We found that this combination scheme improves the global performance, and compares favorably to the motion compensated, complex and time consuming blotch detection methods, for real archive film sequences.*

Keywords: Film restoration, spatio-temporal filtering, Dempster-Shafer fusion.

1 Introduction

The last century has provided a large amount of audiovisual documents, which are endangered, because of the smooth, but irreversible degradation of their support. The main challenge for the traditional stakeholders is to achieve the migration to digital formats and the long-term preservation of these digitized holdings [1].

Digital restoration is a key step, because the increasing need in image quality required by the new digital broadcast formats (HDDVD, HDTV and HD cinema). However, digital restoration seems to be the bottle neck of the preservation process, because of the high costs and low throughput, involving expensive hardware and skilled operators. Digital restoration has to deal with several (sometimes contradictory) constraints: improving the restoration quality in order to deal with high-resolution images (HD, 2k or 4k), lowering the costs, which are closely related to the automation of the process (operator costs), to the costs and the evolvability of the hardware, and finally to the speed of the restoration process. These constraints should be fulfilled by the introduction of software solutions, working on standard low-cost PCs, and by the research of new, fast, adaptive, and high-level algorithms.

Film is the oldest and the most fragile moving picture media. Film impairments are related to the storage conditions (moisture, vinegar syndrome, dye fading), to improper handling (scratches, dust, dirt), and to poorly maintained equipment (scratches, unsteadiness).

In this paper we focus on the detection of the most frequent defects, which are dirt and sparkle. Dirt and sparkles are impulsive (single frame) defects. Dirt can be seen as opaque or semi-transparent clusters with random size, shape and position, caused by dust and dirt stuck on the film, while sparkle are white clusters, caused by the local abrasions of film gelatin. Figure 1 illustrates some examples of opaque and semi-transparent dirt extracted from 16mm archive film.

Figure 1: Examples of archive film blotches : (left) examples of semi-transparent blotches on static background; (right) examples of opaque blotches on textured moving background.

In this paper, we propose a method based on the Dempster-Shafer (DS) fusion framework for optimizing the combination of spatial and temporal information. DS framework models both imprecision and uncertainty of each detector, and combines them in order to take the best decision given this partial information. The fusion also provides a risk index, allowing a variable degree of treatment.

The remainder of this paper is organized as follows. Section 2 gives an overview of the previous blotch detection methods, while Section 3 describes two blotch detectors that we combine in Section 4. Section 5 shows experimental results, and conclusions are presented in Section 6.

2 Previous works

Dirt and sparkle detection often relies on the assumption of smoothness of motion in the sequence, and on temporal and/or spatial inconsistency of defects. Usually, the image is filtered in the temporal and/or spatial domain, and the defects are detected as the thresholded difference between the original frame and the filtered one.

Temporal detection methods have been used early in the problem of dirt concealment [2]. Basically, these methods are based on the computation of the DFD, which is the difference between two consecutive motion compensated (MC) frames. A. Kokaram introduced the spike detection index (SDIa) [3] which computes the minimum between backward and forward DFD, and SDIp that reduces the number of false alarms, by introducing an additional constraint that requires the signs of both DFDs to be identical before a blotch can be detected. In a similar way, P. Schallauer [4] introduced the constraint that the absolute differences of previous and next MC frames have to be below a threshold.

In spite of good global performances, these methods have high computational costs, and produce false alarms when motion estimation fails, i.e. for so-called "pathological motion" [5] (occlusions, uncovering, intermittent motion, erratic motion, motion blur, large displacements, transparency).

Spatial filtering methods assume spatial inconsistency of defects. Median filters have been extensively used because of their ability to eliminate outliers while preserving edges. A. Nieminen proposed [6] a multilevel median filter (MMF), which was implemented as the median of the outputs of several median filters with different topologies. R. Hardie and C. Boncelet [7] introduced the Lower-Upper-Middle (LUM) filter, which output is assigned to the upper or to the lower local median values computed on an upper and a lower intervals defined by two parameters. O. Buisson [8] used a top hat morphological filter, because of its ability to detect specific patterns, such as dust and hair.

However, these spatial filters provide false alarms on sharp and textured regions, i.e. when image spatial patterns look like defects patterns, and fail to detect blotches exceeding the filter size.

Spatio-temporal methods extend spatial filtering to the temporal domain, often using motion compensated (MC) frames. G. Arce [9] introduced a multi-stage order statistic filter (MOS) as an extension of the min/max MMF filter to three (non MC) frames. In the same manner, B. Alp [10] presented the ML3D filter, which uses median operations and provides, according to [11], better impulse noise rejection than those proposed by Arce. A. Kokaram [11] improved ML3D filter (ML3Dex) by using three MC frames. M. Nadenau and S. Mitra [12] presented the rank order detector (ROD), which compares the rank ordered differences between the current pixel and six neighbors (from the previous and next image) against three thresholds. P. Van Roosmalen [13] simplified ROD (SROD) by using a single threshold for the maximum distance between the current pixel and the minimum or maximum of the neighbors. Gangal [14] improved ROD performances by using more robust motion estimation and extending ROD to five MC frames. O. Buisson [8] presented a hybrid detector, based on the combination of SDIa with a spatial morphological filter, while E. Decenciere [15] extended the class of morphological area filters to the spatio-temporal domain.

An original adaptive filter was recently presented by M. Hamid [16]. A soft morphological filter (SMF) works on three non MC images, and its size and shape parameters are learned using a genetic algorithm, supervised by both artificially corrupted and uncorrupted sequences. This filter gives less false alarms than LUM and ML3Dex for fast moving objects, but the learning step is very slow, and the filter parameters should be adapted to each new sequence.

Probabilistic methods have also been proposed, using a Bayesian framework and Markovian models for the detection of temporal discontinuities on MC frames [17, 18, 19]. These methods perform well in real situations, but have a high computational cost, which become intractable when neighborhood order exceeds first or second order [17]. R. Bornard [17] extended the MRF model to five MC frames, and introduced an interpretation step (based on spatio-temporal redundancy of false alarms), that reduced the number of false alarms within pathological motion areas.

Spatio-temporal methods achieve better performance than spatial or temporal methods alone. However, the computational load is higher, and often false alarms due to "pathological motion" persist, which is the consequence of a non optimal combination schema.

A better modeling of the combination of these detectors should improve the quality of detection and decrease the computational time. Temporal detection should be used to leave the ambiguity between dirt clusters and objects of a similar spatial structure, while spatial detection should be used to confirm temporal detections, or to replace these ones when motion estimation fails.

3 Blotch Detectors

Two blotch detectors have been chosen, which provide both redundant and complementary detections. It can be assumed that these detectors are independent (in the sense of belief function fusion) since one is working in the temporal domain and the second one is working in the spatial domain.

3.1 Temporal detector

In order to avoid the high computational cost of motion estimation, we have chosen a simple, but fast temporal detector, working on three non motion-compensated frames. We use the principle of the Simplified Rank Order Detector (SROD) [13], which flags a pixel in the current frame as "blotched", if its value is an outlier of the distribution of grey levels of the neighboring pixels,
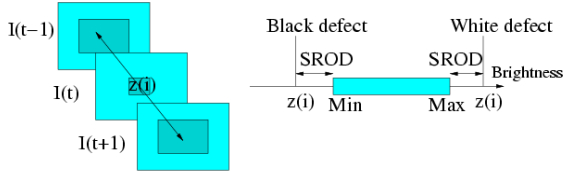
taken from the previous and the next frames (Figure 2).

Figure 2: Principle of SROD.

The "outlierness" of a pixel in the current image is then computed as the difference between the minimum of the distribution and the pixel value for a black outlier, or as the difference between the pixel value and the maximum of the distribution, for a white outlier:

$$SROD(i) =$$
$$\begin{cases} \min(p_k) - z(i) & if \ \min(p_k) - z(i) > 0, \\ z(i) - \max(p_k) & if \ z(i) - \max(p_k) > 0, \\ 0 & sinon. \end{cases} \quad (1)$$

with $k = \{1, ..., N\}$, where N is the number of neighbors, $z(i)$ the value of the current pixel and p_k the value of a neighbor k.

This detector has one free parameter S, which is the size of the windows in the previous and next images (so $N = 2 \times S^2$). S should be larger than the displacements between two consecutive images, but small enough to avoid the loss of sensibility around edges. Efficient computation of local min/max is achieved by first computing min/max on lines, and thereafter on columns.

This detector is able to deal with small motions, but because of the lack of motion compensation fast moving objects can provide false alarms (cf. Figure 3).
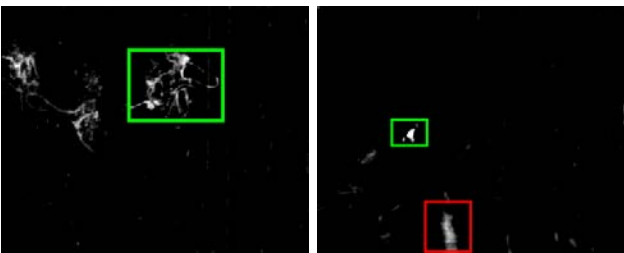

Figure 3: Examples of temporal detections of blotches. Green boxes show examples of good detections, while red boxes show examples of false alarms.

3.2 Spatial detector

In order to decrease the number of false alarms provided by the previous detector, we introduce a second detector, acting in the spatial domain. As blotches are often local extrema with sharp edges in the images, we assume that they can be detected by a morphological filter, using a structuring element with a shape close to real blotch shapes.

Firstly, a top hat morphological filter using non flat large square structuring element with a tunable profile [8] has been implemented. It gives good results for the detection of sharp dust, but its profile is not well adapted to the detection of large and semi-transparent blotches, which are frequent on 16mm film. As structural morphological filters performances are limited by the great variety of blotch shapes and sizes, we have relaxed the constraints on the blotch shape to a single constraint on the blotch area.

Area constraint was introduced using morphological area operators in [20]. The area opening (closing) cuts (fills) a peak (gap or valley) until the area of the cutted (filled) peak (valley) exceeds a given area value. Image structures which do not satisfy the increasing criterion of area are left unchanged. These morphological opening and closing can be efficiently implemented using Tarjan's union find algorithm [21], with a lower computational cost than that of the corresponding structural morphological filter.

The only free parameter of this algorithm is the area parameter λ that should be larger than the blotch size.

This detector provides false alarms for image patterns which are close to the blotch size (Figure 4). However, as it relies on spatial information, its false alarms are different from the false alarms provided by the temporal detector, and should be detected as a "conflict" between detectors at the fusion step.
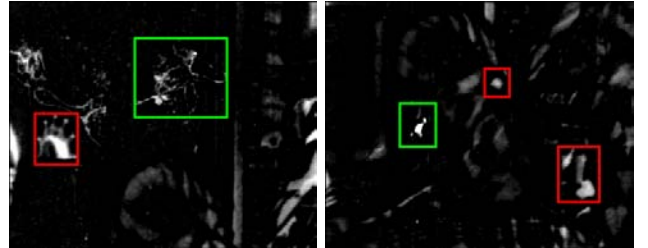

Figure 4: Examples of spatial detection of blotches. Green boxes show examples of good detections, while red boxes show examples of false alarms.

4 Data fusion

In order to take advantage of the redundancy and complementarity of the detectors, we combine their outputs in the framework of the evidence theory, which easily handles the concepts of uncertainty, ignorance and imprecision of data, and is therefore well adapted to our aim.

4.1 Belief structures

We assume that each pixel can be explained by two mutually exclusive and exhaustive hypotheses: "defect" denoted D or "no-defect" denoted $\overline{D}$. The frame of discernment is denoted by Θ:

$$\Theta = \left\{D, \overline{D}\right\}. \quad (2)$$

The set of subsets of Θ is then:

$$2^{\Theta} = \left\{\emptyset, \{D\}, \{\overline{D}\}, \Theta\right\}. \quad (3)$$

Disjunctions or compound hypotheses allow the representation of the ignorance of a source, i.e. if the source cannot distinguish between hypotheses A and B, the corresponding belief is assigned to the union of these hypotheses (i.e. to the compound hypothesis $A \cup B$). A special case is the total ignorance of a source (in the case of missing data), which can be modeled by assigning a non null belief to Θ.

In our two-class application, the only compound hypothesis is $D \cup \overline{D} = \Theta$, and represents the detector inability to classify a pixel as D or $\overline{D}$, near the boundary between D and $\overline{D}$ classes.

The DS evidence theory allows representing both imprecision and uncertainty through two functions: belief (Bel) and plausibility (Pls), which are derived from a mass function (m):

$$\left\{ \begin{array}{l} Bel(A) = \sum_{B \subseteq A, B \neq \emptyset} m(B) \\ Pls(A) = \sum_{B \cap A \neq \emptyset} m(B) \end{array} \right. \qquad (4)$$

For any hypothesis A of 2^Θ, $m(A) \in [0,1]$ and should fulfill the following normalization constraint:

$$\left\{ \begin{array}{l} m(\emptyset) = 0, \\ \Sigma_{A \in 2^\Theta} m(A) = 1 \end{array} \right. \qquad (5)$$

Here, each detector provides numerical output values in an interval denoted as I, and we define each $m(A)$ as a function from I into $[0,1]$ (i.e., $m(A)$ is not simply a number, but a function).

The definition and computation of the mass function is not straightforward, because no general method exists. Mass functions associated with a detector are usually deduced from the empirical conditional probability distributions of the detectors outputs, with respect to the normalization constraint provided by Equation (5).

However, additional constraints on the mass function shape can be derived from the specificity of our application: $m(D)$ should be an increasing function on I because the confidence of hypothesis D increases with the detector output value, while $m(\overline{D})$ should be a decreasing function on I, because the confidence of hypothesis $\overline{D}$ decreases as the detector output value increases. The mass assigned to ignorance ($m(\Theta)$) should be maximum near the boundary between D and $\overline{D}$ classes, accounting for the detector inability to distinguish between D and $\overline{D}$.

For the sake of simplicity, $m(D)$ and $m(\overline{D})$ are modeled by piecewise linear functions, each function being determined by two parameters.

In order to respect the normalization constraint, (Equation (5)), $m(\Theta)$ is computed as follows:

$$m(\Theta) = 1 - m(D) - m(\overline{D}) \qquad (6)$$

Mass functions parameters have been chosen in a supervised manner, using the empirical conditional probability distributions (of the detector output values conditionnally to the classes) with respect to the previous constraints. Figure 5 shows the mass functions which have been chosen for both detectors.
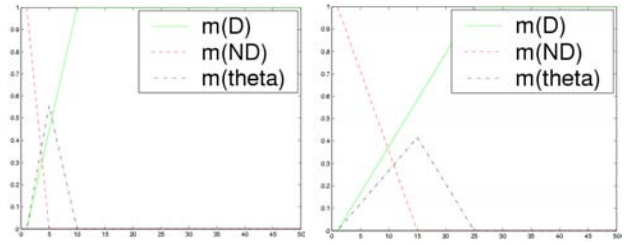

Figure 5: Mass fonctions corresponding to the first (left) and second (right) detector, function of the detectors output values.

4.2 DS combination

The combination of the two detectors has been done by applying Dempster's combination rule [22, 23] on the previously estimated masses:

$$m_{12}(A) = \sum_{B \cap C = A} m_1(B) \times m_2(C) \qquad (7)$$

This computation is performed point-wise over I.

The intersection table is illustrated in Table 1. This table also shows the principle of the DS fusion: some of the mass is transfered from the combined hypotheses to simple hypotheses.

Table 1: Intersections between the two sets of hypotheses.

	D	$\overline{D}$	Θ
D	D	$\emptyset$	D
$\overline{D}$	$\emptyset$	$\overline{D}$	$\overline{D}$
Θ	D	$\overline{D}$	Θ

The empty set mass $m_{12}(\emptyset)$, which measures the conflict between detectors, should deserve a special attention. The conflicting mass can be reassigned to other hypotheses, and the employed method is highly dependent on the problem modeling.

Several solutions have been proposed for the conflict Management [24]. If we consider that information sources are perfectly reliable, we can use Dempster's rule of combination (Equation (7) normalized by $(1 - m(\emptyset))$), or Smets' rule (Equation (7)) if we are not sure that the frame of discernment is exhaustive (open world assumption). However, if we consider that information sources are not reliable, we must apply discounting if possible, or use a disjunctive rule of combination [24, 23].

In our application, we consider that the frame of discernment was correctly modeled, but that sometimes detectors can fail (closed world assumption). For example, the first detector can fail for large motions, while the second detector fails for texture patterns which are similar to the blotch patterns.

Thus, the masses should be discounted by using the reliability of each detector. However the reliability of the temporal and spatial detectors is not easy to estimate as one global factor.

Another solution is to use a disjunctive rule of combination, such as Dubois' rule [25], where the conflicting mass of two subsets B and C is assigned to their union (compound hypothesis) $B \cup C$:

$$m_D(B \cup C) = m_{12}(B \cup C) + \sum_{B \cap C = \emptyset} m_1(B) \times m_2(C) \quad (8)$$

We can notice that in our model with two singleton hypotheses, Dubois's rule simplifies to Yager's rule [26], which assigns the conflicting mass to the whole set Θ (total ignorance):

$$m_Y(\Theta) = m_{12}(\Theta) + m_{12}(\emptyset) \quad (9)$$

Yager's rule of combination merges the total ignorance and the conflict, which can raise some semantics problems in some applications, but this allows defining a single index of ignorance, as long as we consider that the conflict is due to the fact that we ignore which detector is unreliable.

4.3 Decision rule

The decision is based on the belief and plausibility functions (Equation (4)), which are the minimum and the maximum of uncertainty of a given hypothesis.

The length of the belief interval ($[Bel(A), Pls(A)]$) can been interpreted as the imprecision about the uncertainty value, and in our two singleton hypotheses case it is equal to $m(\Theta)$.

The decision rule we have chosen selects the singleton hypothesis which corresponds to the maximum of belief, but only if this maximum exceeds a given threshold Th; otherwise, the decision is Θ (total ignorance):

$$\begin{cases} D & if \ Bel(D) > Bel(\overline{D}) \quad and \ m(\Theta) \geq Sd, \\ \overline{D} & if \ Bel(\overline{D}) > Bel(D) \quad and \ m(\Theta) \geq Sd, \\ \Theta & else \end{cases}$$
$$(10)$$

Figure 6 shows two examples of decision, illustrating the segmentation of the image in "Defect", "No-Defect", and "Ignorance" classes, along with the related amount of ignorance (Figure 7), which in our case measures the "risk" that we assume by taking a decision.
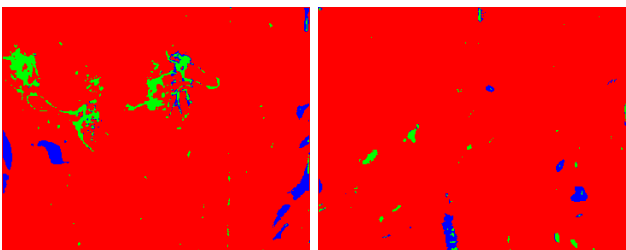

Figure 6: Decisions taken after the fusion step: D is shown in green, $\overline{D}$ in red, and Θ in blue.

5 Experimental results

The performance assessment was performed using the sequences "Dance" and "Art", which have been
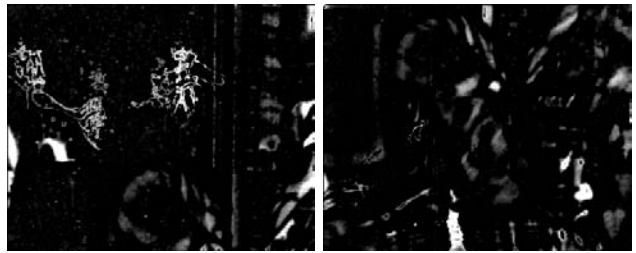

Figure 7: Images of ignorance after the fusion step ($m(\Theta)$).

scanned from archive 16mm color film. These sequences have 136 and 83 images, respectively, and were scanned at the broadcast resolution (720x576). These sequences are challenging, because motions are large and complex, and because the areas of moving objects and of blotches are quite similar. These sequences have also been heavily corrupted by several types of defects (blotches, sparkles, scratches) of various sizes and shapes.

In order to perform quantitative evaluation, the blotch ground truth is required. A traditional approach consists in doing manual segmentation of blotches, or in generating artificial blotches, of random shape, location, and grey levels but with a constant grey level value over the blotch area. However, these models are often unrealistic, as real blotches can be semi-transparent, non-uniform, and have large area (about 1000 pixels).

A good approximation of defects ground truth was achieved using a method of (color) film scanning in infrared light [27]. Infrared images provide an accurate location and transparency of film "physical" defects such as blotches, scratches and gelatine abrasion.

The binary defects ground truth has then been obtained by thresholding infrared images (Figure 8). The threshold value was manually set, in order to find (binary) ground truth patterns as close as possible to the human perception of these defects. In this paper, we chose a high sensitivity threshold, in order to emphasize the semi-transparent blotches.

However, this method has some limitations. Defects which have been revealed in infrared, can be "hidden" by the image content (dark areas), and thus become "invisible" for the defect detection algorithms.
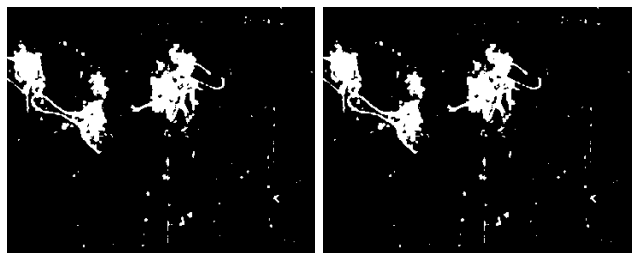

Figure 8: Binary ground truth extracted from infrared images.

The fusion scheme described above has been evaluated using this ground truth. The parameter S of the first detector was set to 15 in order to achieve robust-

ness to small movements, while the parameter of the second detector was set to $\lambda = 800$ pixels, to be larger that the blotch size. Mass functions illustrated in Figure 5 have been used, and the decision threshold Th was set to 0.1.

Figure 9 illustrates the correct detections, the false alarms, and the missed detections of our detector.
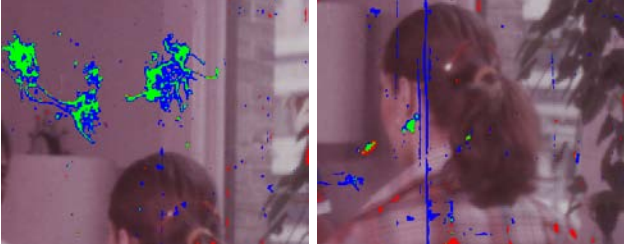

Figure 9: Comparison with ground truth images. Correct detections have been shown in green, false alarms in red, and detections missed in blue.

We can see that semi-transparent blotches have been generally correctly detected, even in areas with pathological motion. However, some false alarms occur on the fast moving textures, showing the limitation of our approach: if small moving patterns are detected by the two detectors, two false alarms are merged, giving a false alarm.

Quantitative performance assessment was achieved using classical performance indexes: recall, which is the percentage of ground truth correctly detected, precision, which represents the percentage of correct detections, and false alarm rate, that is the percentage of false alarms among the no-blotched pixels.

If we denote by DE the binary mask of detection, GT the binary mask of ground truth and $\overline{GT}$ its complement in the image, recall, precision and false alarm rates are defined as follows:

$$Recall = \frac{|DE \cap GT|}{|GT|} \times 100,$$
$$Precision = \frac{|DE \cap GT|}{|DE|} \times 100, \qquad (11)$$
$$FArate = \frac{|DE \cap \overline{GT}|}{|\overline{GT}|} \times 100.$$

The quantitative performance of our blotch detection method is reported in Tables 2 and 3.

Table 2: Comparative performance assessment (sequence "Dance", 136 images).

	Recall	Precision	FA rate
Fusion	22.4%	12.0%	0.56%
SDIa [3]	6.1%	11.2%	0.40%
ROD [12]	8.1%	4.8%	1.31%
Morris [19]	6.2%	5.1%	0.94%
Bornard [17]	3.4%	53.0%	0.024%

However, these results are biased by the presence of nonimpulsive defects (scratches) which cannot be detected by the blotch detectors, and by the defects revealed in the infrared images that are "hidden" by dark areas in the visible image.

Table 3: Comparative performance assessment (sequence "Art", 83 images).

	Recall	Precision	FA rate
Fusion	14.3%	12.4%	0.083%
SDIa [3]	19.1%	5.3%	0.28%
ROD [12]	20.3%	1.8%	0.91%
Morris [19]	7.6%	5.3%	0.11%
Bornard [17]	6.2%	56.1%	0.0037%

Thus a comparative evaluation has been performed, using the same sequence of images and the same evaluation method for several blotch detectors. Four blotch detectors have been used, using 3 or 5 motion compensated images, estimated with a subpixel accuracy phase correlation method. As these methods have several parameters, we have chosen to show only the set of parameters which give the best results.

First, two simple detectors have been chosen for the tests. SDIa [3] computes the thresholded minimum of the differences between the current frame and the previous or next motion compensated frames, and has been used with the threshold $T = 15$. Second detector is the ROD [12], that compares the first three rank ordered differences (between the current pixel and its six spatio-temporal neighbours) against three thresholds. In practice, only the first threshold is significant and was set to $T_1 = 15$. The second and third threshold variation doesn't change the good detection and the false alarms rate, so they have been fixed to $T_2 = 39$ and $T_3 = 55$.

More complex, but accurate MRF detectors have also been tested. Morris detector [19] has been tested with the temporal smoothness bias α set to 0.005, spatial smoothness bias β_1 set to 0.65 and absence of discontinuities bias β_2 set to 0.4.

As R. Bornard's detector [17] extends Morris's detector, the same α, β_1 and β_2 parameters have been used. Additional multiple discontinuities bias β_3 has been set to -0.07, while the surrounding size parameter, which is used in the false alarm detection and elimination post processing, was been set to 38 pixels. However this post processing increases the precision of blotch detection in fast moving sequences, but also decreases the recall, because of the elimination of detections that are close to complex motion areas.

Comparative performance evaluation has been reported in Tables 2 and 3. Our method compares favorably in terms of balance between the recall, precision and false alarms rate to other motion compensated blotch detection methods, and even to the complex and time consuming Markovian methods.

6 Conclusion

In this paper, we propose a fusion scheme for the detection of blotches in digitized archive film material. This method uses the evidence theory framework for com-

bining two blotch detectors, taking advantage of their redundancy, complementarity and incompleteness.

Performances in terms of correct detections and false alarms were improved, as the decision has been performed after the combination step, taking the conflict between detectors into account.

The conflict has been reported on the ignorance mass, thus allowing the introduction of a single risk index. Variable degrees of treatment can be achived by thresholding this risk index in the decision step.

Further improvements will concern the automatic computation of the mass functions, the introduction of discounting factors in order to take into account the detectors reliability, and the introduction of spatial information by spatial fusion.

Finally, this fusion schema can easily be extended to new detectors, improving detection reliability.

7 Acknowledgement

References

[1] http://www.prestospace.org/.

[2] R. Storey. Electronic detection and concealment of film dirt. *SMPTE Journal*, pages 642–647, June 1985.

[3] A. C. Kokaram and P. J. Rayner. A system for the removal of impulsive noise in image sequences. In *SPIE Visual Communications and Image Processing*, pages 322–331, Boston, MA, 1993.

[4] P. Schallauer, A. Pinz, and W. Hass. Automatic restoration algorithms for 35mm film. *Journal of Computer Vision Research*, 1(13):59–85, January 1999.

[5] P. M. B. Van Roosmalen. Technical report for ina. Technical report, INA, 1999.

[6] A. Nieminen, P. Heinonen, and Y. Neuvo. A new class of detail-preserving filters for image processing. *IEEE Transactions on Machine Intelligence*, 9(1):74–90, January 1987.

[7] R. Hardie and C. Boncelet. Lum filters: a class of rank-order based filters for smoothing and sharpening. *IEEE Transactions on Signal Processing*, 41(3):1061–1076, March 1993.

[8] O. Buisson, B. Besserer, S. Boukir, and F. Helt. Deterioration detection for digital film restoration. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, volume 1, pages 78–84, June 1997.

[9] G. Arce. Multistage order statistic filters for image sequence processing. *IEEE Transactions on Signal Processing*, 39(5), November 1991.

[10] B. Alp, P. Haavisto, T. Jarske, K. Oistaemoe, and Y. Neuvo. Median-based algorithms for image sequence processing. In *Proceedings of the SPIE Visual Communications and Image Processing*, pages 122–133, October 1990.

[11] A. C. Kokaram. *Motion picture restoration*. PhD thesis, Cambridge University, 1993.

[12] M. J. Nadenau and s. K. Mitra. Blotch and scratch detection in image sequences based on rank ordered differences. In *Fifth International Workshop on Time-vaying Image Processing and Moving Object Recognition*, pages 1–7, Florence, Italy, September 1996.

[13] P. M. B. Van Roosmalen. *Restoration of archived film and video*. PhD thesis, Deft University of Technology, 1999.

[14] A. Gangal, T. Kayikcioglu, and B. Dizdarogle. An improved motion-compensated restoration method for damaged color motion films. *Signal Processing: Image Communication*, 19(4):353–368, April 2004.

[15] E. Decencière and J. Serra. Detection of local defects in old motion pictures. In *VII National Symposium on Pattern Recognition and Image Analysis*, pages 145–150, Barcelona, Spain, April 1997.

[16] M. S. Hamid, N. R. Harvey, and S. Marshall. Genetic algorithm optimization of multidimensional grayscale soft morphological filters with applications in film archive restoration. *IEEE Transactions on Circuits and Systems for Video Technology*, 13(5):406–416, 2003.

[17] R. Bornard. *Probabilistic approaches for the digital restoration of television archives*. PhD thesis, Ecole Centrale de Paris, 2002.

[18] A. Kokaram. *Motion picture restoration*. Springer-Verlag, 1998.

[19] R. D. Morris. *Image sequence restoration using Gibbs distributions*. PhD thesis, Cambridge University, 1995.

[20] L. Vincent. Grayscale area openings and closings, their efficient implementation and applications. In *Mathematical Morphology and Its Applications to Signal Processing*, pages 22–27, Barcelona, Spain, May 1993.

[21] M. H. F Wilkinson and J. B. Roerdink. Fast morphological attribute operations using tarjan's union-find algorithm. In *International Symposium on Mathematical Morphology*, pages 311–320, Palo Alto, CA, June 2000.

[22] G. Shafer. *A mathematical theory of evidence*. Princeton University Press, 1976.

[23] P. Smets. The combination of evidence in the transferable belief model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(5):447–458, May 1990.

[24] E. Lefevre, O. Collot, and P. Vannoorenberghe. Belief function combination and conflict management. In *Information Fusion*, pages 149–162, 2002.

[25] D. Dubois and H. Prade. Representation and combination of uncertainty with belief functions and possibility measures. *Computational Intelligence*, 4:244–264, November 1998.

[26] R. R. Yager. On the dempster-shafer framework and new combination rules. *Information Sciences*, 41:93–138, 1987.

[27] I. Childs and J. R. Sanders. New capabilities for a line-array ccd telecine. *SMPTE Journal*, pages 1294–1301, 1983.