Robust Detection of Single-frame Defects in Archived Film

Stefanie Wechtitsch Hannes Fassold Peter Schallauer
JOANNEUM RESEARCH, DIGITAL - Institute for Information and Communication Technologies, Graz, Austria
{*stefanie.wechtisch, hannes.fassold, peter.schallauer*}*@joanneum.at*

Abstract

The main issue in current algorithms for the detection of single-frame defects like dust, dirt and blotches in archived film is the significant number of false alarms due to motion compensation errors and film grain. This typically leads to disturbing artifacts occurring in the subsequent defect removal process. We propose a novel algorithm for the detection of single-frame defects which addresses this issue. A continuous response map is defined which indicates the probability of a pixel being part of a single-frame defect. Furthermore, we introduce the novel co-support which is applied to the response map for both noise suppression and spatial completion of potential defects. Finally, regions which are likely affected by motion compensation errors are identified (e.g., by analyzing the motion field) and the single-frame-defect response map is damped in these regions. Experimental results show that the proposed method outperforms state-of-the-art algorithms in terms of accuracy and robustness to motion estimation issues.

1. Introduction

Single-frame defects (SFDs) like dust, dirt, hairs and blotches are probably the most frequently encountered defects in old film stock and therefore their reliable detection (and subsequent removal) is a key functionality in every automated film restoration tool. The methods for SFD detection can be classified roughly into purely spatial and spatiotemporal methods. Spatial methods use only the information from the current image for detection and typically employ some sort of median or morphological filtering. Spatiotemporal methods (e.g., SDIa [3], SDIp [3], DTMF [5], ROD, sROD [4]) use also the temporally neighboring images and exploit the fact that a SFD occurs only in the current image. They typically inspect the pixel-wise intensity differences between the current image and its neighbor images in order to detect an 'intensity spike'. Most spatiotemporal methods use some sort of motion compensation of the neighbor images in order to increase the robustness of the algorithm to camera and object motion. The spatiotemporal methods, especially those with motion-compensation, tend to be more robust than spatial methods. Furthermore, model-based methods based on Markov random fields have been proposed [1], which have the disadvantage of being computationally expensive. A comprehensive survey of state-of-the-art SFD detection methods and objective evaluation of their performance (detection rate, false alarm rate) can be found in [4]. As can be seen from the ROC curves in this survey, a limitation of current methods is that they show a considerable amount of false detections. The reason for this can often be attributed to errors during the motion compensation process due to occlusions, uncoverings, motion blur etc. This is problematic as the subsequent SFD removal process will add artifacts to the restored image. Another issue with many methods is that film grain triggers many false detections, especially when the methods are parametrized to detect also low-contrast SFDs. In this work, we focus on the reduction of these false detections by a combination of several specifically tailored methods. We do not address the subsequent removal of the defects, which is typically done by insertion of motion compensated neighboring content. In section 2 the proposed detection method for SFD is presented in detail, section 3 describes experiments and the evaluation of the proposed detector and the conclusion is given in section 4.

2. Proposed algorithm

The proposed SFD detection approach represents a probabilistic approach which combines several continuous measures. The advantage of such an approach is that it better preserves information when compared with many state-of-the-art methods, which employ an early binary pixel-wise decision. We first calculate a so-

called primary SFD response map for each-pixel based on a modification of an already existing binary SFD detector. Afterwards we apply the novel co-support operator on the SFD response map, a nonlinear operator which is specifically designed to suppress noise and to complete potential SFDs spatially. In the last step, several damping measures are calculated, where each measure indicates areas where so-called pathological motion [2] (e.g. complex motion, occlusion, uncovering) has occurred. These damping measures are inferred from the motion field as well as from the motion-compensated neighbor images. Response values of pixels that are affected by those measures are damped down by multiplying the SFD response with each of the inferred damping measure. From the resulting final SFD response map we calculate a binary dust mask by a simple thresholding operation. For the necessary motion estimation during motion-compensation of the temporal neighbor images we employ a GPU implementation of the TV-L1 optical flow algorithm [6]. In the following, I denotes the center image, $C_{n\_}$ and C_{n_+} are the motion compensated previous and next images in the sequence and τ denotes the threshold which is used in the final phase of the algorithm.

2.1 Primary SFD response

Firstly, we calculate a primary SFD response map where for each pixel $z = (x, y)^T$ its response value $r(z)$ gives an indication of how confident we are that this pixel belongs to a SFD. After investigating several algorithms from [4], we chose the SDIa algorithm [3] as it is very fast, has reasonable detection capabilities and can easily be modified to yield continuous values instead of binary ones. For each pixel of the center image I the primary response $r(z)$ based on the SDIa algorithm is computed as

$$r(z) = min(|I(z) - C_{n\_}(z)|, |I(z) - C_{n_+}(z)|). \quad (1)$$

It is easy to prove that the thresholded SFD response map $r(z)$ is identical to the result of the SDIa detector from [4]. The SFD response is visualized in Figure 1b.

2.2 Co-support operator

The co-support operator is motivated by the observation that SFDs typically form coherent regions. Furthermore, the occurrence of one-pixel defects is rather rare and might indicate a false detection triggered due to film grain in the material. The co-support operator, a novel non-linear operator, was designed keeping these properties in mind. In the following we describe the operator more in detail. We first introduce the general

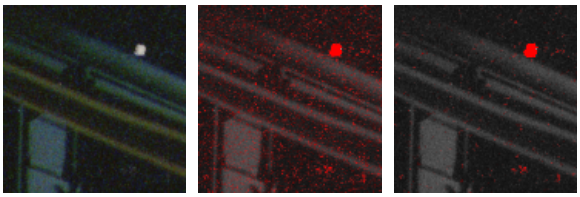

(a) (b) (c)

Figure 1: (a) Original image, (b) Primary SFD response, (c) SFD response after applying the co-support operator. The response maps are visualized in shades of red.

co-support operator, which could have applications beyond the scope of this paper. In particular this operator is applicable for all kind of applications, in which outliers have to be excluded while patterns within the local neighborhood are desired to be preserved. Then we present the specialized co-support operator which is actually applied on the primary SFD response map.

Let v_c be a scalar (or vector) value, representing some feature (e.g. the SFD response, color, etc.) for a pixel. Furthermore, let v_n be the feature value for a specific neighbor in the 8-neighborhood N_8 of the pixel. Then, we define the general co-support operator as

$$cosupp_{gen}(z) = g\{f(v_c, v_n)|v_n \in N_8(v_c)\}. \quad (2)$$

where $f(\cdot, \cdot)$ is some arbitrary function which maps into $\mathbb{R}$ and the operator g is either max_k or min_k (returning the k-highest or k-smallest value of the set).

It's not straightforward to design an appropriate function f for a specific task. Here, we want the specialized co-support operator to act in the following way on a pixel. The pixel's SFD response should be increased if one of its neighbor pixels has a high SFD response. Vice versa, if the pixel does not have a single neighbor pixel with a significant SFD response, then the SFD response of the pixel should be lowered. At first, the *average* function, defined as $\phi(r_c, r_n)$, where r_c and r_n is the SFD response for the center pixel and its neighbor, seems reasonable to be chosen for function f. However, the average function results in an overall growing of defects at their borders which is not desired. To limit the growing effect, we propose a convex combination of *min* and *average*. Ideally, the function $f(r_c, r_n)$ should correspond with function $min(r_c, r_n)$ if r_c is rather low, and it should be $\approx \phi(r_c, r_n)$ if r_c is significant. This is achieved by defining f as

$$f(r_c, r_n) = (1 - \frac{s}{\tau}) \cdot min(r_c, r_n) + \frac{s}{\tau} \cdot \phi(r_c, r_n) \quad (3)$$

where s is calculated as $s = min(r_c, \tau)$. In Figure 1c the result of the specialized co-support operator with f set

according to (3), g set to max_k and k set to 1 is demonstrated.

2.3 Damping measures

Most false alarms typically occur due to motion estimation errors, e.g. in occlusion areas or close to moving edges. Therefore, we calculate a set of damping functions $d(z)$ which indicate these regions (by having high values there) and map into the range $[0, 1]$. The damping of the SFD response map $r(z)$ is then applied by multiplying it with $1 - d(z)$.

The first damping function we propose is derived from the pixel-wise difference of the motion-compensated neighbor images. If the motion is accurately estimated, their difference is supposed to be zero. A similar approach, albeit done in a binary way, is used in the DTMF method [5]. The corresponding continuous damping measure for each pixel is obtained by

$$d_w(z) = min\left(\frac{max(|C_{n_-}(z) - C_{n_+}(z)| - \epsilon_w,\ 0)}{w_{max} - \epsilon_w}, 1\right) \tag{4}$$

where w_{max} is the maximum difference which is mapped to a damping factor of 1 and ϵ_w defines a tolerance value.

Pathological motion often violates the local smoothness assumption of a motion field (u, v). Therefore we calculate a damping measure based on the vector divergence of the motion field by

$$d_{div}(z) = min\left(\frac{max(|\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}| - \epsilon_{div},\ 0)}{div_{max} - \epsilon_{div}}, 1\right). \tag{5}$$

The third damping measure is based on back-matching of the motion vectors. For this measure, we have to calculate the motion field between the center image and its neighbor in both directions, obtaining a forward mapping $m_f(z)$ and backward mapping $m_b(z)$. Let $w(z) = ||z - m_b(m_f(z))||$ be the magnitude of the difference vector when mapping a point z with the forward and then with the backward motion field. For a reliably estimated motion in both directions, $w(z)$ should be zero. Typically, in pathological motion areas (occlusion, uncovering, etc.) $w(z)$ will be significant. Therefore we calculate the back-matching damping measure by

$$d_b(z) = min\left(\frac{max(w(z) - \epsilon_b,\ 0)}{b_{max} - \epsilon_b}, 1\right). \tag{6}$$

The last damping measure focuses on moving edges and textured regions, since it was reported in related work, that false alarms are likely to occur in those areas as well. Therefore, we combine two measures, the magnitude of motion vectors and an image gradient measure. The image gradient magnitude is computed with the central difference quotient. An binary edge image $b_g(z)$ is obtained from it by thresholding using a certain value τ_g. Let $v(z)$ be the magnitude of the motion vector at spatial position z and let v_{max} be the maximum velocity an object is supposed to move within two consecutive frames. By combining $b_g(z)$ and the measure for reliable motion vectors due to its magnitude, a damping measure for moving edges can be obtained by

$$d_m(z) = b_g(z) \cdot min\left(\frac{max(v(z) - \epsilon_m,\ 0)}{v_{max} - \epsilon_m}, 1\right). \tag{7}$$

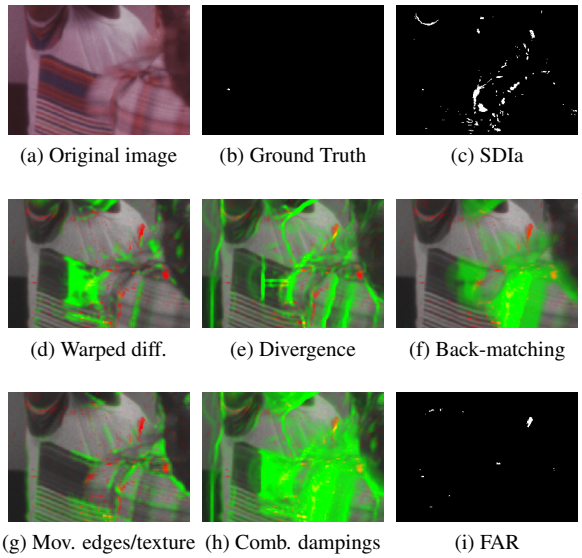

(a) Original image (b) Ground Truth (c) SDIa

(d) Warped diff. (e) Divergence (f) Back-matching

(g) Mov. edges/texture (h) Comb. dampings (i) FAR

Figure 2: Damping measures and the resulting dust masks, where areas in red indicate the SFD response and those in green show the damping.

In Figure 2 the different damping measures are visualized in the green channel. The damping function approach obviously does not increase the detection rate, but represents an effective mechanism for reducing the amount of false alarms.

3. Experimental results and evaluations

In this section a comparative evaluation of the proposed algorithm, denoted as FAR (false alarm reduction), is carried out. The well known ROC curve approach and infrared scans as objective Ground Truth have been employed for evaluating a detector's performance in the same way as in [4].

Three image sequences were used for the evaluation. The *Art* sequence is characterized by moderate to fast translational motion and a high amount of small SFDs. The *Dance* sequence represents a challenging sequence

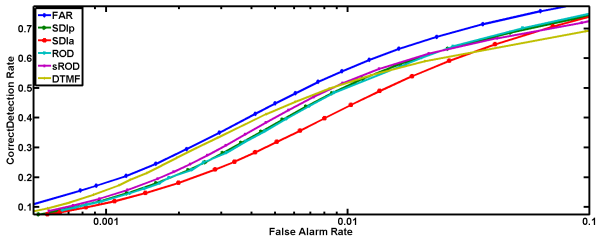

(a) *Dance*

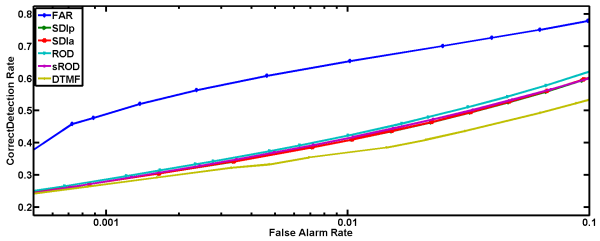

(b) *Jazzman*

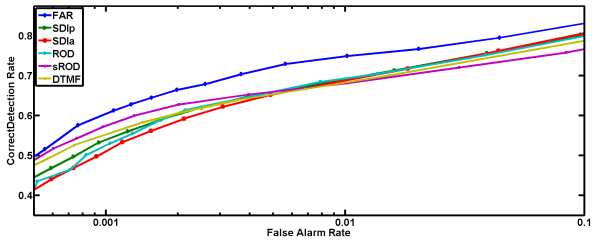

(c) *Art*

Figure 3: ROC analysis of different methods and the proposed FAR method for three test sequences.

with slow camera motion, fast and complex object motion, textured background and several significant SFDs. The last sequence, *Jazzman*, contains only marginal object motion, but a high level of fine noise. The evaluation was done pixel-wise in the same way as in [4]. The results are shown in Figure 3, where the proposed algorithm FAR is compared against the methods SDIa, SDIp, ROD, sROD and DTMF (see [4]). Due to the combination of the various damping measures, it consistently achieves a significantly lower false alarm rate than the other methods. In Figure 2 the combined effect of the various damping functions on the suppression of false alarms due to pathological motion can be seen, when comparing the result of the SDIa detector (Figure 2c) with the result of the proposed FAR detector (Figure 2i). Furthermore, in Figure 3b the benefit of the co-support operator for noisy image sequences is clearly visible.

4. Conclusion

In this work, we have proposed a novel detector for the detection of SFDs (dust, dirt and blotches) which focuses specifically on providing a low false alarm rate. The proposed probabilistic approach combines several continuous measures and preserves information better than methods employing an early binary decision. A primary SFD response map is inferred by transforming the SDIa method so that it provides continuous confidence values. The novel co-support operator is an effective tool to increase the robustness against noise and film grain and furthermore has a positive effect on the spatial completion of potential SFDs. Finally, we apply a combination of several specifically tailored damping measures which mark areas that are likely to produce false alarms. These damping measures are based on the analysis of the motion field and the neighbor images and provide a strong damping of the SFD response map in pathological motion areas (occlusion and uncovering areas, moving edges areas etc.). An objective evaluation shows that our proposed FAR detector achieves a significantly lower false detection rate when compared with state-of-the-art SFD detection algorithms.

5. Acknowledgment

This work was funded by the Austrian Research Promotion Agency (FFG) under the Bridge project "Dust-NG" (No. 827580). Thanks to INA[1] for providing the infrared scans as test data, which have been made available within the project "PrestoSpace"[2] (FP6-507336). We want to thank also Werner Bailer for providing valuable comments and proof-reading the paper.

References

[1] M. N. Chong and D. Krishnan. An edge-preserving MRF model for the detection of missing data in image sequences. *Signal Processing Letters*, 5(4):81–83, 1998.

[2] D. Corrigan, N. Harte, and A. Kokaram. Pathological Motion Detection for Robust Missing Data Treatment in Degraded Archived Media. *ICIP*, pages 621–624, October 2006.

[3] A. C. Kokaram. *Motion Picture Restoration*. Springer, London, 1998.

[4] J. Ren and T. Vlachos. Detection of dirt impairments from archived film sequences: survey and evaluations. *SPIE Optical Engineering*, 49(6), June 2010.

[5] P. Schallauer, A. Pinz, and W. Haas. Automatic Restoration Algorithms for 35 mm Film. *Videre, Journal of Computer Vision Research*, 1(3):59–85, 1999.

[6] M. Werlberger, W. Trobin, T. Pock, A. Wedel, D. Cremers, and H. Bischof. Anisotropic Huber-L1 Optical Flow. *BMVC*, 2009.

[1] Institut National de L'Audiovisuel (http://www.ina.fr)

[2] http://prestospace.org