Detection and Recovery of Film Dirt for Archive Restoration Applications

Jinchang Ren[1,2], *Theodore Vlachos*[1]

[1] Centre for Vision, Speech and Signal Processing, University of Surrey, Guildford, UK.
[2] Digital Imaging Research Group, School of Informatics, University of Bradford, UK.

j.ren@{surrey,bradford}.ac.uk t.vlachos@surrey.ac.uk

ABSTRACT

A novel spatio-temporal method is proposed for film dirt detection and recovery. Firstly, a more reliable confidence measurement of dirt is extracted for color films. False alarms caused by motion are filtered using consistency checks among several measurements. Then, candidate dirt is detected by filtering and thresholding this confidence measurement. Finally, bi-directional local motion compensation and ML3Dex filtering are taken for the recovery of dirt pixels. Experiments on real data demonstrate the efficiency and effectiveness of our method in terms of both detection and recovery of dirt.

Index Terms: film dirt detection, archive restoration, spatio-temporal filtering, missing data recovery.

1. INTRODUCTION

Archive restoration of degraded films in the digital domain has recently attracted lots of interest [1-5, 10-11], and several high-profile projects have received EU funding, such as AURORA, BRAVA and more recently PrestoSpace. In general, dirt is among the most common impairments in archived films, which occur when dust or other material adheres to the film due to electrostatic effects, or when the film is passed through various transport mechanisms [1-3]. Usually dirt only lasts for one frame, appearing mostly as dark or bright opaque spots of random size, shape and location.

Depending on whether the recovery process is applied exclusively to dirt pixels or not, existing methods can be characterised as one- and two-pass. The former utilize global filtering of frame images to recover dirt thus it has the potential to over-smooth non-dirt pixels as well. The latter usually has a detection process to identify candidate dirt pixels such that the reconstruction algorithm will concentrate on these areas and reduce errors induced during recovery [1, 3].

For dirt detection, we have three further classes, namely spatial, temporal and spatio-temporal. Spatial processing involves only intra-frame information, i.e. a local neighborhood for each missing pixel [7-8]. Using recursive median filtering and several rules, they fail to recover large or fast moving dirt and generate many false alarms [5]. In temporal processing inter-frame information from several neighboring frames is used and in each frame only one corresponding pixel is treated by comparing motion compensated errors against one or more thresholds [1-2]. Spatio-temporal processing is a combination of the above two, such as the rank-order-detector (ROD) [4] and soft morphological filtering [3].

After detection, "interpolation" in spatial and/or temporal domain(s) is usually employed for the recovery of missing data. In Kokaram *et al* [1], ML3Dex filtering, MRF and AR interpolation are used. Again, the performance of these three methods depends on accurate motion compensation. In [9], global motion compensation and coherence-based search are employed for adaptive recovery of missing data using temporal or spatial interpolation. This approach does not perform well in the presence of complex motion and only thin-line structure like scratches can be corrected. In addition, some model-based methods have also been introduced using Markov/ Gibbs random fields or autoregressive modelling in both detection and recovery of missing data [1]. However, these methods may fail if accurate motion compensation cannot be achieved.

In this paper, we propose a spatio-temporal method for detection and recovery of dirt in degraded colour films. Firstly, we present a dirt detection scheme suitable for colour images, which is effective even when motion compensation fails owing to similar luminance but different colours. Secondly, a confidence measurement is extracted, which is helpful in both detection and recovering of dirt. Thirdly, false alarms caused by motion are avoided by consistency checks among several measurements. Finally, local motion estimation and compensation is selectively applied to candidate dirt regions for computational efficiency.

2. DIRT DETECTION AND RECOVERY

Confidence measurement of dirt was firstly introduced by us in [5]. Nevertheless our approach ultimately proved sensitive to false alarms caused by scene motion and moreover could only detect small size dirt. Here we overcome these shortcomings and also avoid overdependency on thresholding. In addition, we apply local motion estimation and compensation for the recovery of dirt pixels.

2.1. Extended confidence measurement of dirt

Let f_{n-} and f_{n+} be two neighboring frames of f_n, and $M(f_n)(i,j) = \max(r_{i,j}^{(n)}, g_{i,j}^{(n)}, b_{i,j}^{(n)})$ denotes the maximum value among the three colour components in f_n. This is based on a reasonable assumption that dirt pixels should appear in grey in colour sequences thus the maximum component is more accurate to represent its energy. Two elementary differences, d_{n-} and d_{n+}, is then defined as

$$\begin{cases} d_{n-} = M(f_n) - M(f_{n-}) \\ d_{n+} = M(f_n) - M(f_{n+}) \end{cases} \quad (1)$$

As dirt is a temporal impulse, consistent differences are expected: A more close d_{n-} and d_{n+} stands for a more likely dirt, and vice versa. If both d_{n-} and d_{n+} are negative or positive, this relates respectively to dark or bright dirt pixels.

In Kokaram [1] and Schallauer *et al* [2], dirt is detected by comparing absolute differences between compensated frames with some given threshold(s). A poor choice of threshold may lead to irrecoverable loss of information. Instead we extract a confidence measurement in a similar way as defined in [5]. We additionally define a combined measurement of d_{n-} and d_{n+} in (2), and we also define $D_n = 0$ if $d_{n-}d_{n+} \le 0$.

$$D_n(\lambda) = \frac{d_{n-}d_{n+}}{\lambda |d_{n-}| + (1-\lambda)|d_{n+}|} \quad (2)$$

where $\lambda \in [0,1]$ is a parameter to weight these two elementary differences. It is apparent that the definition in [5] is a special case of (1) where $\lambda = 1/2$.

With $D_n(\lambda)$, a confidence measurement of dirt can be obtained in the same way as defined in [5] and [10]. To remove noisy measurements, median filtering is applied to this confidence image, and the result is denoted as $Conf_n(\lambda)$, and dirt pixels are further determined by thresholding this confidence measurement.

Figure 1 shows one color image (with main dirt areas marked in white boxes) and different confidence images extracted from this color image and its corresponding grey image with or without median filtering. For comparison, the corresponding objective ground truth (GT) is also given. This GT is obtained by infrared scanning of original films, which is then thresholded to generate a binary mask as shown in Figure 1(d). The results from (b) to (c) are extracted from luminance component (or grey source), as used in [5], and they are less accurate than those from color components in (e) and (f). Moreover, median filtering seems overall useful towards removing false alarms and improving accuracy.

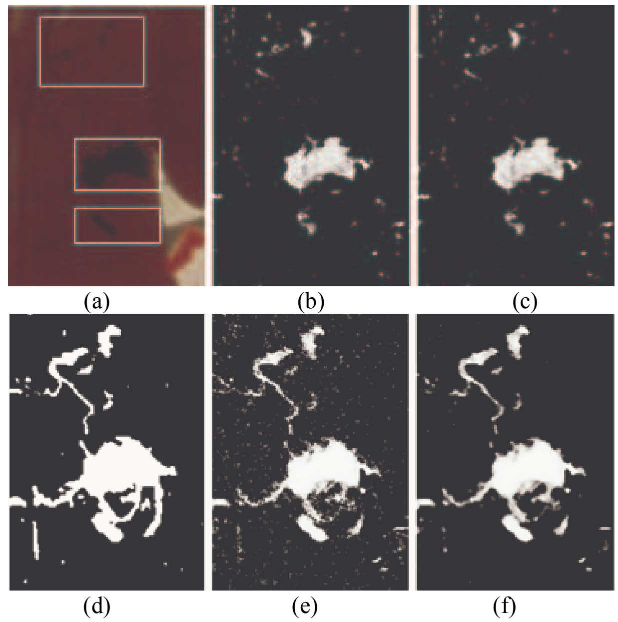

Figure 1: One color image (a) and its ground truth map of dirt (d) and confidence of dirt extracted from grey image (b) and from color image (e). (c) and (f) and results after median filtering of (b) and (e), respectively.

2.2. Avoiding false alarms caused by motion

Since D_n varies with λ an unsuitable value for λ may lead to many false alarms (caused by motion). As a result, $Conf_n$ and detection efficiency also rely on λ.

If dirt pixels lie within a static background, this is an idealized case where we have $d_{n-} = d_{n+}$ and also a local maximum value of D_n. Consequently, in this case no false alarms will occur.

If the moving area has near-constant illumination or color (especially around dirt pixels), which is mostly true owing to piecewise smoothly-varying intensity in natural scenes, then this case is not much different than the idealized one and so false alarms can also be avoided.

Regarding other cases, as dirt pixels are motion-independent, we expect their confidence measurements to be insensitive to the change of λ. Consequently, we choose three values for λ, i.e. 0.1, 0.5 and 0.9, and take pixels of consistently high confidence measurements in the corresponding three confidence images as real candidates, otherwise we consider them as false alarms.

2.3. Recovering dirt pixels

To recover missing data due to dirt, motion estimation and compensation are required. Let $f_{n-1}(i_-, j_-)$ and $f_{n+1}(i_+, j_+)$ be two motion-compensated pixels corresponding to the missing data of $f_n(i,j)$. A simple method was introduced in [2] taking the average of $f_{n-1}(i_-, j_-)$ and $f_{n+1}(i_+, j_+)$ as an estimate of $f_n(i,j)$. In [1], a two-stage median filtering, ML3Dex, was employed, as shown in Figure 2, to recover $f_n(i,j)$:

$$\begin{cases} z_l = median[W_l] \qquad l \in \{1,2,3,4,5\} \\ \bar{f}_n(i,j) = median[z_1, z_2, z_3, z_4, z_5] \end{cases} \tag{3}$$

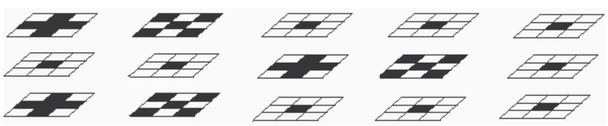

Figure 2: Five sub-filter masks (W_1 to W_5) defined in ML3Dex for missing data recovery in [1].

For each candidate region of dirt $R_{n,k}$, a larger rectangular region $S_{n,k}$ is defined using an enlarged outer boundary of $R_{n,k}$ satisfying (4), where $|.|$ denotes cardinality i.e. number of pixels within a region.

$$R_{n,k} \subset S_{n,k}, |S_{n,k}| \geq 3|R_{n,k}| \tag{4}$$

Meanwhile, for each $S_{n,k}$, two corresponding regions are defined in f_{n-1} and f_{n+1} as $S_{n-,k}$ and $S_{n+,k}$. These occupy the same the spatial location in each of the three consecutive frames. Then, forward and backward local motion estimation and compensation are employed between $S_{n-,k}$ and $S_{n,k}$, $S_{n+,k}$ and $S_{n,k}$, respectively. Afterwards, ML3Dex filtering is utilised. Please note that both motion estimation and ML3Dex filtering are only applied to dirt areas, rather than the whole image, for efficiency and avoiding over-smoothing of image details.

3. RESULTS AND DISCUSSION

3.1. Detection results

We compare the performance of our method against several well-established approaches in the literature. For detection, four other methods are considered including ML3Dex [1], LUM [8], SDIp [1] and ROD [4]. The threshold in the first three methods is set as 10, and the three thresholds used for ROD are 5, 10 and 15. In LUM, a spatio-temporal $3\times3\times3$ window with $N = 27$, and $k = 9$ was used. In SDIp, ROD and ML3Dex, motion estimation and compensation is applied by using the Black-Anandan optical flow algorithm [6], which is also used for local motion compensation in our approach.

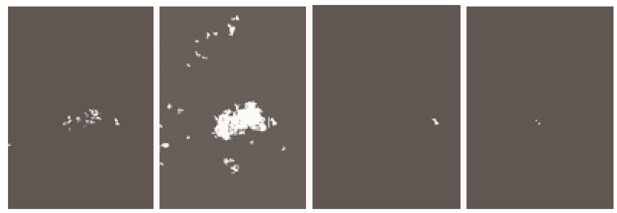

(a) SDIp (b) ROD (c) ML3Dex (d) LUM

Figure 3: Detected dirt for the source image in Figure 1(a) using SDIp, ROD, ML3Dex and LUM.

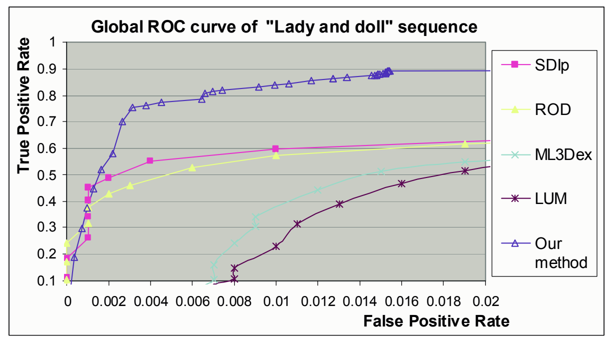

Figure 4: ROC comparison for several methods for test image in Figure 1(a) using SDIp, ROD, ML3Dex, LUM and our method.

Using objective ground truth (GT) maps, a quantitative performance assessment is possible by computing of Receiver Operating Characteristics (ROCs). By counting correctly detected or missed dirt pixels between the dirt mask and the GT map, true positive detection rate, R_{tp} and false positive detection rate, R_{fp}, are computed. Figure 4 demonstrates an average ROC curve for 80 consecutive frames. Greylevel GT data were thresholded at level 85 to generate a binary mask. In Figure 4, it can be clearly seen that our method outperforms all others when false positive rate is more than 0.1%. Although SDIp and ROD occasionally yield better higher true positive rate when a much lower false positive rate (say less than 0.1%) is required, they are less efficient because they need bi-directional motion estimation and motion compensation over the whole image.

3.2. Reconstruction results

In this set of experiments, ML3Dex filtering is applied to the whole image and also to detected dirt masks for missing data recovery. In Figure 5, results on global filtering as well as local filtering on the dirt masks from SDIp and our method are compared. For better visualization purpose, the source image and reconstructed images are contrast enhanced. In the enhanced image, we can easily identify dirt areas and compare the performance from different methods. From Figure 5 we can see that our method generates comparable results to global ML3Dex filtering while local ML3Dex fails to reconstruct all the missing data owing to incomplete masks of dirt detected.

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

Figure 5: Contrast-enhanced section of test image (a) and reconstructed images using ML3Dex applied to SDIp detection (b), whole image (c) and our method (d).

Figure 6 shows another reconstruction example in which dirt is marked within a grey rectangle in Fig 6(a). Although our method fails to detect and also recover several boundary pixels of the dirt area, it has preserved well most of the details and avoids over-smoothing caused by global ML3Dex filtering (see regions between fingers).

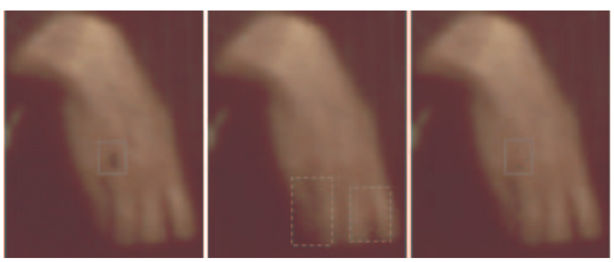

(a) (b) (c)

Figure 6: Contrast-enhanced section of test image (a) and reconstructed images using global ML3Dex filtering (b) and our local filtering (c).

Finally, our method is more efficient as bi-directional motion estimation and compensation is only employed to candidate dirt regions. Dirt pixels normally occupy no more than 1% of the whole frame area, hence we provide a good balance between accuracy and efficiency.

4. CONCLUSIONS

A spatio-temporal method was presented for dirt detection and recovery in dirt-impaired archived film material. We demonstrated that confidence measurement of dirt extracted from color channels in our approach is more accurate and robust relative to greylevel information only. This confidence measurement was shown to be useful for dirt detection. Results from real dirt samples of degraded films show that our algorithm performs well in terms of accuracy, robustness and efficiency and compares favorably with competing approaches in the field.

5. ACKNOWLEDGEMENT

This work was funded by project PrestoSpace, supported by the European Commission (FP6-IST-507336). We would like to thank the research staff at Snell & Wilcox, UK for valuable discussions and also thank INA, Paris, for providing the sequences with objective ground truth.

6. REFERENCES

[1] A.C. Kokaram, "Motion picture restoration", Berlin, Germany: Springer-Verlag, 1998.

[2] P. Schallauer, A. Pinz and W. Haas, "Automatic restoration algorithms for 35mm film", *J. Computer Vision Res.*,1(3), 1999, pp. 59-85.

[3] 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 T-CSVT,* 13(5), 2003, pp. 406-416.

[4] M.J. Nadenau and S.K. Mitra, "Blotch and scratch detection in image sequences based on rank ordered differences", *Proc. of 5th Int. Workshop on Time-Varying Image Processing*, Florence, 1996, pp. 27-35.

[5] J. Ren and T. Vlachos, "Segmentation-assisted dirt detection for the restoration of archived films", *Proc. BMVC*, 2005, pp. 359-368.

[6] M.J. Black and P. Anandan, "The robust estimation of multiple motions: parametric and piecewise-smooth flow fields", *J. CVIU,* 63(1), 1996, pp. 75-104.

[7] G. R. Arce, "Multistage order statistic filters for image sequence processing", *IEEE Trans. Signal Proc.,* 39(5), 1991, pp. 1146-1163.

[8] R. Hardie and C. Boncelet, "LUM filters: a class of rank-order-based filters for smoothing and sharpening", *IEEE Trans. Signal Proc.,* 41(3), 1993, pp. 1061-1076.

[9] R. Bornard, E. Lecan, L. Laborelli and J.-H. Chenot, "Missing data correction in still Images and image sequences", *Proc. ACM Multimedia*, 2002, pp. 355-361.

[10] J. Ren and T. Vlachos, "Efficient detection of temporally impulsive dirt impairments in archived films", *Signal Processing*, 87(3), 2007, pp. 541-551.

[11] A.C. Kokaram, R. Morris, W.J. Fitzgerald, and P.J. W. Rayner, "Detection of missing data in image sequences," *IEEE Trans. Image Proc.*, 40(11), pp. 1496-1508, 1995.