

An Automatic Optical Flow Based Method for the Detection and Restoration of Non-repetitive Damaged Zones in Image Sequences*

Roman Dudek, Carmelo Cuenca, and Francisca Quintana

Departamento de Informática y Sistemas,
Campus Universitario de Tafira, Las Palmas de Gran Canaria, Spain
roman@idecnet.com, {ccuenca, fquintana}@dis.ulpgc.es

Abstract. In this paper, we describe an automatic method for detecting and repairing non-repetitive damages in image sequences, caused by dust, fibers or local defects of the film emulsion. The method is a three frame window scheme based on the calculation of the optical flow (OF) relating adjacent frames and the first and the last frames of the sequence. The OF validity is checked in order to detect non-repetitive damage, and is later repaired using filtering and smooth blending of the damaged zones. The method works correctly for the set of tested image sequences providing perfect visual repairs of the damaged zones.

1 Introduction

There is a vast archive of celluloid film motion pictures documenting the modern history. With prolonged periods of storage in not always ideal conditions, they suffer from serious conservation problems. Neither the celluloid nor the chemical layer on it is stable over large periods of time. Most archive film motion pictures suffer from damages and need to be digitized and restored in order to preserve the invaluable cultural heritage they contain. Specially the older archive films can contain large amounts of film damage, like spots, dust, scratches, celluloid shrinking or unstable frame position. While for the news and documentary purposes digital video recording has gradually replaced the use of film negative, almost all existing feature film productions have been shot using the traditional film technology, with only some of the latest features shot digitally. The usage of traditional celluloid film technology inherently implies the presence of some drops, mostly caused by film emulsion defects and dust.

Among the different types of defects and damages there are the following:

- Processing liquids spots: This is a non-repetitive defect caused by the incorrect drying of the processing liquids, which can create spots on individual images of the sequence, like the spot shown in Figure 1(a).

* This work has been partially supported by The Spanish Ministry of Science and Innovation under contract TIN2007-60625 and by FEDER funds.

- Film emulsion drops: The film emulsion itself can contain defects. They usually appear as random dots or spots on individual frames. These dots are non-repetitive. We can see an example in Figure 1(b). If the defect is "hard" the emulsion can appear completely peeled off, as in Figure 2(b).
- Dust on film: It is normal that some dust particles get stuck on the film. The dust particles are fragments of fibers, hairs, plants, etc. Their amount can vary depending on the care taken during the film processing. A careful manipulation of the film can decrease these defects, but it will never make them disappear. The film dust will be seen as random fibers or dots on individual frames. They are not repetitive along multiple consecutive frames. In Figure 1(c) an example of a fiber stuck on the film can be seen.
- Mold: The photosensitive emulsion is hold on the celluloid by a common, animal based gelatine. The gelatine is prone to be attacked by mold in presence of humidity, thus creating non-repetitive spots on individual frames of the sequence.
- Dust on camera or film duplicating optics: If the dust is stuck in some part of the camera optics, in the path of the incoming light, then there will be a case like the above (dust on film), but in this case the dust will then be impressed on the same position along a sequence of frames.
- Film scratches: The film is mechanically threaded through the camera. If a hard dust particle is stuck in a point where the film slides along, a longitudinal scratch will be created along the film. These scratches will be visible as vertical lines in the film, provided that in typical cameras the film advances vertically. The scratches are usually visible in the same position along many consecutive frames of the sequence.

Some of the above defects can be physically cleansed in the laboratory, using special liquids and gasses. However, most archive materials are actually not the original film negatives. In the standard laboratory process, the negatives are optically transferred to distribution positives. The original negatives were frequently discarded or lost, so any dust present on the negative, or in the equipment, was permanently impressed on the positive, and in the copies made after these positives, thus propagating the effect.

Many works have been published regarding film restoring and dealing with various aspects of film repair [1] [2] [3] [4]. In the field of repairing localized defects we can distinguish methods aimed at repairing repetitive or non-repetitive damage. Non-repetitive damage appears only on individual frames, and generally can be repaired using information from adjacent undamaged frames. Repetitive damage is a damage which appear along a long sequence of frames, like a film scratch [5] or dust in optics. The information lost to such damage can not be easily patched using adjacent frames and need to be "guessed up" using much more advanced, complex assumptions about the scene depicted, like assuming that a scene consists basically from a static background with objects moving over, or that a walking person repeats over the same movements in each step of its walk [6]. Such advanced assumptions restrain the list of scenes that can be treated, however they can be the only available solution in many cases.



Fig. 1. Examples of film damages

Our work concentrates on non-repetitive damage, like dust on film, emulsion defects, liquid drops and mold damage. We will concentrate on creating a simple to use, yet reliable automatic method to detect the damaged zones and integrate the repair as seamlessly as possible into the original footage.

The rest of the paper is structured as follows. Section 2 explains the automatic method based on the optical flow that detect and repairs non-repetitive damage. Section 3 presents the results obtained for different image sequences. Finally, section 4 depicts some conclusions.

2 Our Proposal

Our work deals with the film restoration of certain types of film damage using information from the adjacent frames in the film image sequence. Let suppose that in the present frame of the sequence there is a localized defect, like a spot, a fiber or a dot. Our solution is copying a similar but undamaged part of another image belonging to the same image sequence. We propose bringing the corresponding parts from the adjacent images in the image sequence to reconstruct the damaged part of the current image. Unfortunately, direct copying of pixels from the adjacent frame would only work on static image sequences. Using of the optical flow (OF) based movement interpolation extends the method usability to sequences with movement.

2.1 Optical Flow Based Image Interpolation

In our optical flow implementation, based on comparison studies by [7] [8] [9], we have chosen a variation of the Horn and Shunck [10] algorithm, extended by a multi-pyramidal scheme to account for large movements. We also added a modification into the calculation of the luminance gradient over time, using an image warping technique [11], thus improving the method's results and achieving really good results on real world image sequences.

Let suppose that we are repairing the image I_t , from the sequence of frames $I_1 \dots I_n$. First, we will calculate the optical flow field $\mathbf{v} = (u, v)^T$ between the previous image I_{t-1} and the following image I_{t+1} of the sequence. We compute

$$\nabla I \cdot \mathbf{v} + I(u_x + v_y) = -\frac{\partial I}{\partial t} \quad (1)$$

Now, for each pixel of the I_t image that needs to be reconstructed, we will calculate its value bringing it from the following image:

$$I'_t(x, y) = I_{t+1}\left(x + \frac{u}{2}, y + \frac{v}{2}\right) \quad (2)$$

Notice that the optical field \mathbf{v} describes movements along the interval of two frames. To interpolate the path along the interval of one frame (from I_t to I_{t+1}), we divide the vector size by 2. The coordinates of the pixel accessed in the I_{t+1} will usually be non integer values, so We use linear interpolation to access these pixel positions.

Although it could be thought that it would be a better solution to calculate the optical field between I_t and I_{t+1} , as it would lead to shorter movements to detect, with a better OF quality, this is not the case. Consider that the damaged spots are mostly of different luminance (black or white). In that case the presence of the spot in only one of the pair of the images used for the OF calculation will cause such OF field to be invalid just in the spot of the damage, exactly where we are actually going to use it. To avoid the interference of the damage in the calculation of the OF, we will use the previous and next frames of the sequence, but not the frame itself. If the damage is non-repetitive it is unlikely that the corresponding parts of the adjacent frames will be damaged too.

2.2 Optical Flow Based Process for Automatic Identification of the Damaged Zones

There is an obvious need for an automatic method that identify the damaged zones. While an operator could identify the damaged zones, such process is quite laborious, as the operator would need to review frame by frame the whole sequence, and mark the damaged zones using a pen-on-tablet repair brush interface. Even in case of treating current, correctly developed negative in a film production, which can contain less than one defect per frame for the operator to mark, a skilled operator can review and fix at most a few thousands of frames in a working day (1 minute of film = 24x60 frames = 1440 frames), making human assistance based process a reasonable proposition only for high budget productions. However, purely manual repair of damaged older film could require man-years of work.

For that reason, an automatic method for detecting and repairing damage is needed. However, as the ultimate criterion for film repairing is the visibility of the defects by a human, the operator could review the results and intervene only when it seems necessary, when noticed that the automatic identification did not work properly.

In order to identify the damaged spots we are concentrating on, we will use its non-repetitive nature. A non-repetitive damaged spot is present in the current frame, but it is not present in the previous or in the following frames of the sequence. To detect them, we will use the fact that the optical flow between a pair of images will be invalid or erroneous near an object that is present in only one of the images. This happens because the constant luminance assumption is broken by an object appearing in only one of the frames. We base the damage spot detection on the measure of the OF validity. We assume that where the OF is correct, the OF vectors will map the individual pixels of the source image onto pixels with similar luminance. We estimate the local validity of the OF field using this fact, defining an error measure for the OF error estimation as the square of the difference between the expected and the actual luminance of each pixel mapped by the OF,

$$E_{t-1,t+1}(x, y) = (I_{t+1}(x + u, y + v) - I_{t-1}(x, y))^2 \quad (3)$$

However, damaged spots are only one of the possible phenomena creating OF errors. Many reasons for error in OF are valid phenomena in the scene, such as occlusions, transparencies or moving shadows. These cases would be false positives and therefore we do not have to repair them. To distinguish actual damage from false positives, we use that the damage is only present on an individual frame, not in the previous or in the following frames. Therefore, the OF between the previous and the following frames, but not using the damaged frame, will contain no error caused by the damaged spot. The reason is that an occlusion or moving shadow in progress, or in general any progressive change along the image sequence, will cause certain errors in the calculated OF. In general, these errors increase with the separation in time of the image pair we are using for the OF calculation. The larger the advance of the error-causing phenomena (a progressive occlusion, for example), the larger the OF error. In general, the OF error between a pair of frames separated by a long distance in the image sequence (a distance of two frames in our case), will be larger than the OF error between consecutive frames (a distance of one frame). Expressed as an equation,

$$E_{t,t+1}(x, y) < E_{t-1,t+1}(x, y) \quad (4)$$

Now, consider the case of a non-repetitive damaged spot that is only present in the I_t frame, but not in I_{t-1} or I_{t+1} frames. The opposite of the above equation will be true: the OF error skipping the current frame with damage will be smaller than the OF error of the “shorter” OF relating the previous and the current frames. Using this property, we define the damage measure $D(x, y)$ for detecting zones of spot damage,

$$D(x, y) = E_{t,t+1}(x, y) - E_{t-1,t+1}(x, y) \quad (5)$$

The damage measure D is a scalar value defined for each pixel of the current frame I_t .

Finally, based on the damage measure D , we can create a composition mask to combine the undamaged parts of the original image with reconstructed parts

replacing the damaged zones of the current image. This is the mask the human would manually create in the simpler, human assisted method. The simplest way would be applying a threshold value to distinguish minor noise disturbances from actual film damage. However, the D measure will usually contain a large volume of noise as a consequence of the film grain noise [12] in the processed film sequence. To reduce this noise we used a simple Gaussian filtering, thus creating a filtered measure $D'(x, y)$,

$$D'(x, y) = (g \circ D)(x, y) \quad (6)$$

Using a Gaussian filtering will reduce noise, but it will reduce the resolution of the measure mask. However, the resolution is not critical, as the OF is locally smooth and the detected errors tend to create clusters in the damaged zone.

Finally, we create the actual mask. To further improve the repair integration instead of a hard threshold we will use a linear clamped function to convert the damage measure $D'(x, y)$ into a composition mask,

$$p(x, y) = clamp((D' - c) * s) \quad (7)$$

where c and s are user defined constants and $clamp()$ is a function limiting p to the range $[0, 1]$. Then, the final result is obtained by using a blending equation instead of using a threshold,

$$I_{result}(x, y) = p(x, y) * I'_t(x, y) + (1 - p(x, y)) * I_t(x, y) \quad (8)$$

Adequate values for c and s are defined mostly by the image resolution in pixels and the level of noise in the processed film. As a consequence, usually the same values can be set only once and applied over many film sequences coming from the same origin (from the same day of shooting, for example).

Notice that the process we propose is symmetrical, that is, the OF could be calculated in the opposite direction, (I_{t+1} to I_{t-1}), and then bring the corresponding pixels from the previous image. There seems to be no important advantage in doing so. One possible improvement to the method could be calculating the OF in both directions and then average the results to reduce the average method error. We tested this possible improvement but doing so produced a unwished side effect. In the reconstructed zones, the film grain noise inherently present in all film images, was reduced with the averaging. While noise reduction could be welcome, if applied only on some parts of the image it becomes visible and an unwished artifact on its own.

It is important to note that the proposed method will only work for non-repetitive damage, like emulsion defects, fibers and dust stuck on the film strip. The repetitive damage like film scratches is much more difficult to treat, as the part of the image covered by the defect can be completely lost along the whole image sequence, and would need to be guessed from scratch. Only some special cases, like an scene in constant, regular movement (pan over horizon) could be solved, as the damage is static relatively to the frame, although it is not static relatively to the depicted scene. If the number of frames affected by the defect is

greater than one, but it does not include the whole of frames in the sequence, a more complex method could be designed, which included more than only three sequence frames.

3 Results

This section presents the results of our method on an example of three consecutive images, shown in Figure 2(a), 2(b) and 2(c) extracted from a longer film sequence. We are only showing a cutout of the complete frames so the results can be seen in more detail. The middle image, Figure 2(b), contains a large damage in form of white spots, probably caused by a pealed emulsion on the positive film print. The previous frame (Figure 2(a)) and the following frame (Figure 2(c)) contain no visible damage in the same zone, so if the damage is




Fig. 2. A sequence of three consecutive frames. The intermediate frame contains white spots damage.



(a) A mask image as it would be created by a human operator if the process were manual. This mask image marks the damaged zone in the frame. (b) The mask image created by our automated detection process, without Gaussian filtering. (c) The mask image created by our automated detection process, with Gaussian filtering.

Fig. 3. Examples of mask images



(a) Simple repair example: Bringing the same zones from third frame of the sequence, using the mask created by our method. (b) Optical Flow corrected repair example: Bringing the corresponding image zones from the third frame of the sequence and compensating the motion in the scene using the OF calculated between the first and third frames.

Fig. 4. Simple and OF corrected restorations

non-repetitive and the overall damage is not too extensive, it is unlikely they will contain damage exactly in the same zone.

First, in Figure 3(a) we present the mask image that would be created by a human operator, if the process was to be done manually. The human operator marked the damaged zone of the image we are repairing (Figure 2(b)). This mask can be created using a pen-on-tablet interface. In that case, only a few frames a minute can be treated if the damage is not extensive. Note that the outline of the mask is soft, allowing for smooth blending of the repaired zones with the original frame, thus making the repair less visible.

In Figure 3(b) we present the mask created by our automated detection process by simple thresholding of the damage measure D . No filtering is used on the damage measure, D , in this case. We can see that the mask outlines are quite hard and irregular, making the borders of the repair potentially visible. In order to avoid these borders, Figure 3(c) shows the mask obtained after Gaussian filtering the damage measure D . In D' the noise present in our initial D damage measure has been removed before extracting the mask. As a convenient side effect, the Gaussian filtering also widens the zone to repair, thus producing certain safety margin around the damaged spot, and it softens the spots outlines achieving better visual integration of the repaired zones.

We use the mask shown in Figure 3(c) to compose parts of the image following the damaged frame in the sequence (Figure 2(c)) over the image we are repairing (Figure 2(b)).

In Figure 4(a) we directly bring the same zones from the following image of the sequence using the mask created with our method. The damage is less obvious



(a) A processing liquid spot. (b) A dark spot. (c) A fiber.

Fig. 5. Examples of film corrected damages for images in Figure 1



(a) The automatic mask for liquid spot. (b) The automatic mask for dark spot. (c) The automatic mask for the fiber spot.

Fig. 6. Examples of automatic image masks for images in Figure 1

now. However, a close examination reveals that the substituted zones do not fit well and the repaired spots look distorted. Finally, in Figure 4(b), we bring the corresponding image zones from the following image of the sequence also using the mask created with our method, while compensating the motion in the scene using the OF calculated between the previous and the following frames. We can see in that Figure a visually perfect repair of the damaged zones. The repair will be perfect with the condition that the following frame is not exactly damaged on the same spot.

We have also tested the method with different kinds of damages in image sequences, such as a processing liquid spot, a dark spot and a fiber spot. In Figure 5 we can see the same three frames as in the Figure 1 but in this case the damage has been properly corrected using optical flow. The images in Figure 6 correspond to the three automatically computed image masks. In all frames, white zones identify the damages in the frames and black zones identify parts of the images which are considered by the application of the method without spots.

4 Conclusions

There is a vast historical film archive that it is worthwhile preserving and restoring as it is considered to be a cultural heritage. A large part of this historical footage, and to a lesser extent more recent or even current film documents, are affected by multiple kinds of film damage, ranging from old film shrinking to dust in camera. We concentrated our work on repairing irregular damage, like non-repetitive spots or drops in the images. We proposed and tested an automated method identifying this kind of damage, distinguishing it from other phenomena in the actual depicted scene. Our method is a three frame window scheme based on the calculation of the optical flow relating the adjacent frames and the first and last frames, skipping the central frame, comparing the OF validity for detecting non-repetitive damage. We use both filtering and smooth repair blending to achieve seamless repair of the detected damaged zones. We also use the same optical field relating the first and third frame of our window to create an interpolated frame to use as source for “patching” the damaged second frame of the window.

The automatic identification works correctly for most of the tested scenes. However some sequences containing special phenomena, difficult to distinguish from film damage, like scenes of snow, fire or a fly present in only one picture, can not be adjusted to prevent appearance of either false positives or false negatives. Some human assistance would be needed in order to identify correctly the damage. In these cases, the automatically created mask can be used as base to be corrected by the human operator, speeding up at least partially the process. Nevertheless, even in such difficult cases, once the operator points the defects, the system can still repair these automatically, with good visual results.

References

1. Van Roosmalen, P.: Restoration of Archived Film and Video, Delft University of Technology, Ph.D. Thesis (1999)
2. Joyeux, L., Boukir, S., Besserer, B., Buisson, O.: Reconstruction of Degraded Image Sequences. Application to Film Restoration. *Image and Vision Computing* 19(8), 503–516 (2001)
3. Harvey, N., Marshall, S.: Application of Non-Linear Image Processing: Digital Video Restoration. In: *Int. Conf. on Image Processing*, vol. 1, pp. 731–734 (1997)
4. Boulanger, J., Kervrann, C., Bouthemy, P.: Space-Time Adaptation for Patch-Based Image Sequence Restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29(6), 1096–1102 (2007)
5. Joyeux, L., Buisson, O., Besserer, B., Boukir, S.: Detection and Removal of Line Scratches in Motion Picture Films. In: *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, vol. 1, pp. 548–553 (1999)
6. Patwardhan, K., Sapiro, G., Bertalmío, M.: Video Inpainting Under Constrained Camera Motion. *IEEE Transactions on Image Processing* 16(2), 545–553 (2007)
7. Álvarez, L., Weickert, J., Sánchez, J.: Reliable Estimation of Dense Optical Flow Fields with Large Displacements. *Int. Journal of Computer Vision* 39(1), 41–56 (2000)

8. Beauchemin, S., Barron, J.: The Computation of Optical Flow. *ACM Computing Surveys* 27(3), 433–467 (1995)
9. McCane, B., Novins, K., Crannitch, D., Galvin, B.: On Benchmarking Optical Flow. *Computer Vision and Image Understanding* 84(1), 126–143 (2001)
10. Horn, B., Schunk, B.: Determining Optical Flow, AI Memo 572. Massachusetts Institute of Technology (1980)
11. Papenberg, N., Bruhn, A., Brox, T., Didas, S., Weickert, J.: Highly Accurate Optic Flow Computation with Theoretically Justified Warping. *Int. Journal of Computer Vision* 67(2), 141–158 (2006)
12. Minelly, S., Curley, A., Giaccone, P., Jones, G.: Reducing Chromatic Grain Noise in Film Sequences. In: *IEEE Colloquium on Non-Linear Signal and Image Processing* (Ref. No. 1998/284), vol. 5, pp. 1–5 (1998)