A New Robust Technique for Stabilizing Brightness Fluctuations in Image Sequences

François Pitié, Rozenn Dahyot, Francis Kelly and Anil Kokaram

Electronic and Electrical Engineering Department
Trinity College Dublin, Ireland
{fpitie,dahyotr,frkelly,akokaram}@tcd.ie

Abstract. Temporal random variation of luminance in images can manifest in film and video due to a wide variety of sources. Typical in archived films, it also affects scenes recorded simultaneously with different cameras (e.g. for film special effect), and scenes affected by illumination problems. Many applications in Computer Vision and Image Processing that try to match images (e.g. for motion estimation, stereo vision, etc.) have to cope with this problem. The success of current techniques for dealing with this is limited by the non-linearity of severe distortion, the presence of motion and missing data (yielding outliers in the estimation process) and the lack of fast implementations in reconfigurable systems. This paper proposes a new process for stabilizing brightness fluctuations that improves the existing models. The article also introduces a new estimation method able to cope with outliers *in the joint distribution* of pairs images. The system implementation is based on the novel use of general purpose PC graphics hardware. The overall system presented here is able to deal with much more severe distortion than previously was the case, and in addition can operate at 7 fps on a 1.6GHz PC with broadcast standard definition images. [1]

1 Introduction

Random fluctuation in the observed brightness of recorded image sequences, also called *flicker*, occur in a variety of situations, all well known in the vision and image processing community. The most commonly consumer observed instance of flicker is in archived film and video (see figure 1 and material [8]). It is caused by the degradation of the medium, varying exposure times of each frame of film, or curious effects of poor standards conversion. In any situation where multiple views are recorded by different cameras, the problem can also be observed due to different camera behaviour or lighting conditions due to camera orientation, even if video cameras have been previously calibrated using radiometric calibration routines. Both outdoor and indoor footage can show flicker due to illumination variation. In addition, flicker can also affect modern film and video media especially when telecine transfer is not properly done.

[1] This work has been founded by Enterprise Ireland, the EU project BRAVA, HEA PRTLI TRIP and the European project Prestospace.

Fig. 1. Example of flicker. See especially the black diagonal on the right frame.

Flicker is a global phenomenon, not spatially uniform, affecting the whole recorded image plane. As such its presence also has a detrimental effect on any application involving image matching since they tend to assume brightness constancy between frames, like in motion estimation [6, 4]. In addition, flicker reduces the redundancy between frames and hence increases the bandwidth of transmitted sequences. This is particularly a problem for Digital Television and the compression of multiview scenes. Dealing with the problem of flicker requires some attempt to model or measure the brightness fluctuation between frames (the estimation process), and then to remove this fluctuation in some way (the correction process). This paper presents a new robust technique to correct flicker in image sequences. By merging models previously proposed in the literature, we propose a new generic non-linear model to link two images affected by flicker. This non-linear transformation is then estimated using two approaches, one being the standard non-robust histogram matching technique [9]. As an alternative, we propose to make the estimation more robust using two methods. First, the non-linear function is directly estimated by Maximum A Posteriori on the joint distribution of the intensities of the two considered frames. Second, because this distribution is a mixture of data linked by the flicker model (inliers) and some mismatched data (outliers due to occlusions or motion), we propose an original inlier enhancement process to reduce the influence of outliers in the estimation. In addition, we present a novel implementation based on general purpose PC graphics hardware. Our results show that we are able to cope with a much wider range of flicker effect than previously was the case. We also present results showing the effect of flicker on MPEG4 compression of different kinds of sequences including multi-view camera sequences, and the improvement in bandwidth usage with our proposed flicker reduction method.

2 Related works

The following paragraphs present several approaches that have been proposed in the literature - most of them independently - for the removal of flicker in videos. Three main models have been proposed for intensity distortion. All try to recover the unknown original intensity image (ground truth) I_n^o given the observed intensity images I_n at time n. The estimation $\hat{\mathrm{I}}_n^o$ of the non-degraded image gives the restored frame I_n^R.

Linear model with spatial dependencies. Decenciere [1] proposed a linear model, involving a gain a_n and an offset b_n, linking the intensities of the observed image I_n to the original flicker free image I_n^o. As illustrated in figure 1, the flicker is not the same across an entire frame. Roosmalen [12, 13] introduced then the spatial dependency (x, y) as follows:

$$\mathrm{I}_n^o(x, y) = a_n(x, y)\ \mathrm{I}_n(x, y) + b_n(x, y) \tag{1}$$

While processing the video, the original flicker free image $\mathrm{I}_n^o(x, y)$ is not yet available at time n and is replaced in equation 1 by the previous estimated one $\hat{\mathrm{I}}_{n-1}^o$. In Roosmalen's approach, spatial variation is introduced on a block basis, thus $a_n,\ b_n$ are piecewise constant. Parameters are estimated by least squares estimation [14]. Each estimate is associated with a confidence measure [12].

Some pairs of blocks can not be matched because of missing data occurring for instance, with blotches, occlusions, etc. To cope with this problem, Roosmalen [12] and Yang [14] suggest to detect occluding areas based on spotting large intensity differences that cannot be explained by flicker alone. Parameter estimation is then performed only on the blocks in which there are no outliers detected. Estimates for the "missing blocks" are then generated by some suitable interpolation algorithm. Unfortunately, this method for detecting outliers fails in the presence of heavy flicker degradation.

Because the restored images are generated by *locking* the brightness to previous frames, errors can accumulate. To avoid this problem, due partly to a bias in the estimation, the final restored intensity image is given by a mixture of $\hat{\mathrm{I}}_n^o$ and the observed image :

$$\mathrm{I}_n^R(x, y) = k\ \hat{\mathrm{I}}_n^o(x, y) + (1 - k)\ \mathrm{I}_n(x, y) \tag{2}$$

where k is the forgetting factor set experimentally at 0.85 [13].

Linear Model with Robust Regression. Ohuchi [10] and also Lai [6] consider the same linear model as in equation (1) but the spatial variation of the gain and the offset is expressed directly using second order polynomials [6, 10]. The parameters of those polynomials are then estimated using robust M-estimation [6, 10, 3, 11] involving a Reweighted Least Squares algorithm. Robustness of the estimator is needed to deal with outliers that frequently occur in old videos.

However, as noticed by Kokaram et al.[5], due to the correlation between the frames, the regression (robust and non-robust) introduces a bias in the estimates that can damage seriously the restoration process in case of heavily degraded sequences. Therefore, the authors [5] have introduced a slightly modified linear model that allows reduction of this bias. Another improvement proposed in [5] is to change the polynomial basis to a cosine basis to express the gain and the offset. Since the success of global motion estimation is linked to flicker correction (and vice versa), some have proposed to couple both estimations [6, 13, 5]. In this paper, images will be registered prior to any flicker estimation.

Non-linear model. However, the brightness distortion can sometimes be non-linear. Naranjo et al. [9] proposed the following non-linear model: $\mathrm{I}_n(x, y) = f_n(\mathrm{I}^o_n(x, y))$, where f_n is any increasing function estimated by comparing the intensity histogram of the current observed image and the average histogram of neighbouring frames [9] (see equation 5). Notably, this model does not account for the spatial variation of the flicker defect.

3 A New Model

In our experience with this problem, spatial dependence and non-linearity are key to modelling a wide range of flicker defects. Since the causes of flicker are usually unknown and various, we prefer to adopt a very weak prior on the distortion function. We propose therefore to extend the non-linear model proposed by Naranjo [9] to integrate spatial variations: $\mathrm{I}^o_n(\mathbf{x}) = f_n\left(\mathrm{I}_n(\mathbf{x}), \mathbf{x}\right)$, where $\mathbf{x} = (x, y)$ is the pixel location, I^o and I the flicker free and the observed frames respectively.

Ideally, an estimate of a transformation should be performed for each pixel $\mathbf{x}$. However, this is computationally expensive and it would require an additive spatial smoothness constraint. Alternatively, we propose to estimate the transformations f^i_n at regularly spaced control points $\mathbf{x}_i$ in the image using interpolating splines of order 3. The splines yield and implicit smoothness constraint and our corrected pixel at $\mathbf{x}$ can be written: $\mathrm{I}^o_n(\mathbf{x}) = \sum_i^N w(\mathbf{x} - \mathbf{x}_i)\ f^i_n\left(\mathrm{I}_n(\mathbf{x})\right)$, where f^i_n is the transformation at control point $\mathbf{x}_i$ and $w(\mathbf{x})$ the interpolating 2D mask. The number of control points on one axis is defined as the *flicker order*.

Temporal Integration. A key deviation from previous efforts in flicker removal is that we model the change in brightness between observed images and we do not model the brightness change between the observed and the hidden clean images. Thus f is estimated between frames n, m and our new model is:

$$\mathrm{I}_m(\mathbf{x}) = \sum_i^N w(\mathbf{x} - \mathbf{x}_i)\ f^i_{n,m}\left(\mathrm{I}_n(\mathbf{x})\right) \tag{3}$$

Therefore our task is to smooth brightness changes between frames and not necessarily reveal the true original image. Brightness variations between images I_m and I_n can be caused by: 1) intentional effects like shadows or gradual brightness changes, due to special editing effects for instance, and 2) the flicker degradation which is unintentional. The first effect is generally low frequency and exhibits slow temporal variation. The second effect is temporally impulsive and it is this signal that has to be removed.

As illustrated in figure 2, we consider all the estimations $f_{n,m}$, between I_n and its neighbouring frames $\{\mathrm{I}_m\}$. These estimated parameters $f_{n,m}$ correspond to the impulsive flicker mixed with the gradual informative variations. To separate the impulsive flicker and keep the informative variation, a temporal robust

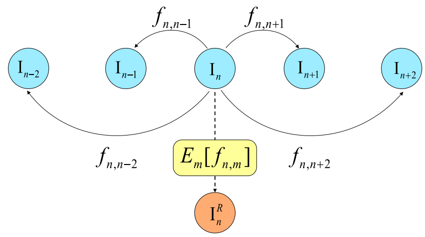

Fig. 2. Compensation of a frame: $\{\mathrm{I}_n\}_n$ original pictures, I_n^R : restoration of I_n, $f_{n,m}$ brightness variation parameters between frames I_n and I_m.

expectation is computed:

$$\hat{f}_n(k) = \mathbb{E}_m\left[f_{n,m}(k)\right] = \arg\min_{f_n(k)} \sum_{m=n-M}^{n+M} e^{-\frac{(m-n)^2}{\sigma_w}} \rho(f_{n,m}(k) - f_n(k)) \qquad (4)$$

where $f_{n,m}(k)$ is the k^{th} component of the lookup table $f_{n,m}$, ρ a robust function [11] and σ_w a temporal scale factor. The temporal window has been fixed experimentally to 15 frames ($M = 7$). The expectation $\hat{f}_n$ is estimated using Reweighted Least Squares and finally applied to I_n to generate the restored image I_n^R.

Speeding Up the Correction Scheme. It is possible to dramatically reduce the complexity by estimating $f_{n,m}$ through the simple combination of successive estimations $f_{n,n+1}$. However, when the flicker is too severe that only a portion of the intensity range is occupied by an image (skewed pdf) the parameter estimates are poor. To conclude, two strategies are available: one involving exhaustive estimations but able to cope with extreme fluctuations in intensity without propagation of errors and one involving a minimal number of estimations but more liable to errors.

4 Estimation of the Brightness Distortion Function $f^i_{n,m}$

In the case of a spatially constant brightness distortion ($\mathrm{I}_m = f_{n,m}(\mathrm{I}_n)$), it has been shown in [9] that the distortion function can be estimated by [2]:

$$f_{n,m}(\mathrm{I}_n) = \mathrm{C}_m^{-1} \circ \mathrm{C}_n\left(\mathrm{I}_n\right) \qquad (5)$$

where C_n and C_m are the cumulative histograms of I_n and I_m. In our case we can assume the distortion locally constant. At each control point $\mathbf{x}^i$, the local pdf is given by [2]:

$$p_n^i(u) = \frac{\sum_{\mathbf{x}|\mathrm{I}_n(\mathbf{x})=u} w(\mathbf{x}-\mathbf{x}_i)}{\sum_{\mathbf{x}} w(\mathbf{x}-\mathbf{x}_i)} = \sum_{\mathbf{x}|\mathrm{I}_n(\mathbf{x})=u} w(\mathbf{x}-\mathbf{x}_i) \qquad (6)$$

[2] where $\sum_{\mathbf{x}} w(\mathbf{x}) = 1$ for splines.

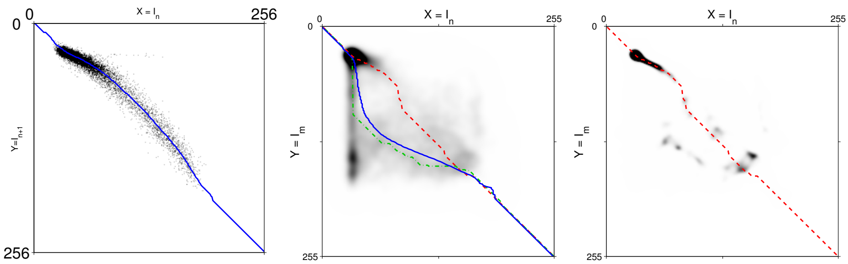

Fig. 3. Examples of joint pdf and their corresponding estimated distortion functions. On (a) an example of non linear distortion and its pdf matching estimation; on (b) an example of blotch noticeable by the vertical line and its consequence on the estimation by pdf matching (solid blue), by using MAP (dotted green) and by using the inlier enhancer and MAP (dashed red); on (b*): the joint pdf of (b) after inlier enhancement.

where $\mathbf{x} \to w(\mathbf{x})$ is the spatial weighting function and $\mathbf{x}_i$ a control point. $\hat{f}^i_{n,m}$ is estimated using equation (5).

This *Pdf matching* scheme (eq. 5) is a non-biased estimator, unlike the regression methods [5]. Another interesting point about considering pdfs separately on each frames is that pdfs are not affected by local motion within the area of consideration, whereas regression methods assume that there is no local motion between both frames. Nevertheless, as shown in the next paragraph, this estimator is very sensitive to large occlusions. We present next an original method that allows us to cope with severe levels of occlusion.

A New Treatment of Outliers. Previous methods for outlier detection [12, 14] are based on brightness differences between frames. The detection is taken by thresholding the mean absolute error. However, in case of strong flicker, the differences might be only due to the flicker and not to the outliers. We propose here a new way to dissociate outliers and inliers based on the observation of the local joint pdf of the image intensities (similarly defined as in eq. 6). To simplify notation we denote U and V the random variables I_n and I_m and we write f instead of $f^i_{n,m}$.

Estimation of f by Maximum A Posteriori. We propose here to estimate f that maximises of the number of couples (u, v) that follow the flicker model (i.e. $f(u) = v$). In other words we want to maximise the likelihood $\prod_u \mathcal{P}(u, v = f(u)|f)$ with the prior knowledge that f is increasing[3]. This can be simply done in a Markovian framework by considering $f = [f(0), \dots, f(u)]_{u \in 0:255}$ as a Markov chain of order 1, whose transition probabilities $\mathcal{P}(f(u+1)|f(u))$ are such as to

[3] For simplicity, we use $\mathcal{P}(u, v)$ instead of $\mathcal{P}_{UV}\,(U \in \mathcal{C}_u, V \in \mathcal{C}_v)$ and $\mathcal{P}(u)$ instead of $\mathcal{P}_U(U \in \mathcal{C}_u)$. The probabilities are in practice approximated by histograms.

ensure $f(u+1) > f(u)$. It is possible as well to add boundary conditions on f so that $f(0) = 0$ and $f(255) = 255$.

$$\hat{f} = \arg\max_f \prod_{u=1}^{255} \mathcal{P}(u, f(u)|f) \cdot \mathcal{P}(f(u)|f(u-1)) \tag{7}$$

The Maximum A Posteriori $\hat{f}$ can be found by using the Viterbi algorithm [7].

Inliers Enhancement. However the MAP estimation can perform poorly in the presence of important occlusions.Figure 3-b shows the joint distribution of intensity between two images I_n and I_{n+1} corresponding to axis u and v respectively. If the images presented no outliers and were only degraded by brightness fluctuation (ie. no occlusions, no local motion), the distribution would be along the plot $v = f(u)$ (i.e. the inlier class $\mathcal{I}$). The case illustrated shows a region of missing data in frame $n+1$ (typically a large region of constant color). This outlier causes a spurious ridge in the distribution manifesting as an horizontal line $u = u_0$. Sadly some of the probabilities on this outlier line are greater than the corresponding inliers one (i.e. $\mathcal{P}(u_0, f(u)) > \mathcal{P}(u, f(u))$) and the MAP estimation will try to follow this outlier line.

The task is to successfully reject the outlier ridge. Simply weighting out parts that are too far away from the identity line $u = v$ is the classical method used in robust regressions [5] but it becomes inefficient when the transformation f makes inlier pairs deviate a lot. To suppress the ridge, it would be intuitively beneficial to explore a weight $\mathcal{R}(u, v)$ on the joint distribution which make inliers row and column wise maximum:

$$\text{if } (u,v) \in \mathcal{I} \text{ then} \begin{cases} \forall u' \neq u, & \mathcal{R}(u,v) > \mathcal{R}(u',v) \\ \forall v' \neq v, & \mathcal{R}(u,v) > \mathcal{R}(u,v') \end{cases} \quad \text{where } f(u) = v \tag{8}$$

We propose to examine the factor :

$$\mathcal{R}(u,v) = \mathcal{P}(u|v)\ \mathcal{P}(v|u) = \frac{\mathcal{P}(u,v)^2}{\mathcal{P}(u)\ \mathcal{P}(v)} \tag{9}$$

Under some hypotheses (cf. Appendix A), it can be shown that the factor $\mathcal{R}$ fulfills conditions (8) and can therefore be used iteratively to enhance the distribution of the inliers in the mixture while reducing the outliers. Initialising $\mathcal{P}^{(0)}(u, v)$ at $\mathcal{P}(u, v)$, we iterate:

$$\mathcal{P}^{(n+1)}(u,v) = K^{(n+1)}\ \mathcal{P}(u,v)\ \mathcal{R}^{(n)}(u,v) \tag{10}$$

where $K^{(n+1)}$ is a normalising constant. It is shown experimentally that applying this procedure for $n = 3$ produces sufficient attenuation of the outlier distribution and then to improve the estimation of f (cf. section 6). Figure 3-b* shows some results obtained with this iterative process. We can see that the occlusion in sequence (b) has been efficiently removed. The estimation of f using MAP on $\mathcal{P}^{(n)}(u, v)$ is then improved.

5 Practical Issues

Tuning the Restoration Process Parameters. The flicker order (i.e. the number of control points per dimension) typically vary from 3—for the vast majority of image sequences—to 6 or even 14 for old movies whose film has chemically changed. For flicker orders greater than 6, occlusions (blotches and local motions) force us to use the robust scheme presented in paragraph 4. Working on color sequences can be done by simply stabilizing the luminance component of the frames.

Flicker Compensation using Graphics Hardware. Computational load can be a key issue in a restoration process involving thousands of frames. To try and alleviate this, we focused on the flicker compensation stage of our algorithm. We found that on average flicker compensation accounted for 80% of the total restoration time for our non-robust restoration process.Modern computer graphics card are becoming much more programmable and contain powerful graphic processing units (GPUs). In fact they can now be considered as useful co-processors to the CPU. Flicker compensation on the GPU is only possible because of the latest advances in graphics hardware architecture especially full support for floating point accuracy. Using fragment programs we can perform the necessary operations to map f_n^i on an image block and multiply it by $w(\mathbf{x}-\mathbf{x}_i)$. We use vertex programs to correctly position the interpolating kernels on the image. Render-to-texture and floating point data are required for the summation of $f_n^i(\mathrm{I}(\mathbf{x}))\ w(\mathbf{x}-\mathbf{x}_i)$.

Performance. Figure 4 shows the results of using the GPU compared to the CPU. These results were obtained on a 1.6GHz Pentium 4 machine running Windows 2000, with an Nvidia GeForce FX5600 graphics card. Using the GPU implementation we reduced the time taken for flicker compensaton from 80% of the total restoration time to 55%. On average the GPU implementation is 3.5 times faster than the CPU implementation. Altogether the full non-robust scheme processes on average 6.8 frames per second at flicker order 3. The full robust scheme takes 1.5 frames per second whereas Roosmalen's process takes around 2 to 3 seconds a frame on similar hardware. For heavily degraded sequences the full process can take up to 20 seconds per frame.

6 Results

Sequences. The first sequence, *Snake*, composed of images from different uncalibrated cameras. This is a modern sequence actually used for post production.The second one, *Paula*, is a 8mm movie from a private repository, captured by pointing a DV camera at the projection screen; the asynchrony of the frame rates is the principle source of flicker. We also processed the *Tunnel* sequence (obtained by telecine) [13] and eventually we processed sequences from *Rory O'More*, a severely degraded movie from 1911.

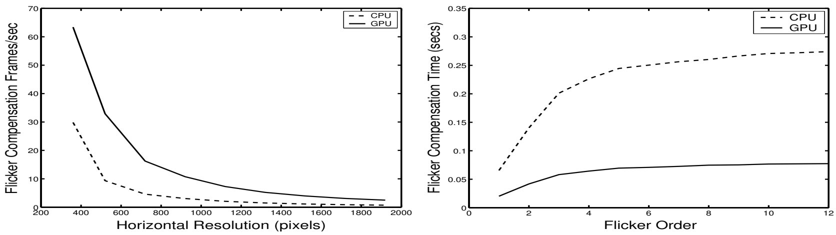

Fig. 4. Flicker Compensation CPU Vs GPU.

Scenarios. Different restoration scenarios have been explored : no restoration at all (ob), naranjo's restoration (na), our method with *pdf matching* (pm), with MAP only(v), with MAP and inlier enhancement (hv) and finally our correction scheme assuming f linear and estimated as specified in [5] (af).

Evaluation by comparing the mean and the variance. Assessing the performance of the systems on real degraded sequences is difficult because of the lack of objective criteria for assessing the quality of the restoration. However, as shown in [13, 10, 9, 5] it is feasible to expect that a good de-flicker process would reduce the fluctuations in the mean and the variance of image intensities from frame to frame. Figure 5-b shows some results for (pm) on the *Paula* and *Snake* sequences. We can see how the filter smoothed the brightness fluctuations. Figure 5-a shows the importance of the non-linear treatment. Whereas (af) cannot remove completely the fluctuations, (pm) stabilizes the brightness. However (na) hardly corrects the flicker.

Evaluation by comparing the MPEG4 compression performances. However the mean and the variance cannot characterize subtle differences between restorations, especially if the dirts and blotches make the mean and variance fluctuate. We propose therefore a novel criteria for assessing the quality of flicker reduced, by comparing the compression ratio given by a MPEG4 encoder (in our case the Microsoft MPEG4 encoder). Results on the *Tunnel* sequence corroborate our previous remarks : (pm) improves the compression performances by 48.6%, (af) by 45.8% and (na) by only 38.4%. This criteria is also very useful to assess the improvements of our inlier enhancement heuristic (hv) as clearly shown in table 1 for damaged sequences from *Rory*.

However, as for the mean and the variance measure, this evaluation is still biased because it favours restoration processes that reduce details level.

7 Conclusion

Images sequences can be affected by brightness fluctuation for many reasons. We presented in this paper a new model able to deal with various kinds of flicker, a new process of correction, and a new efficient way for removal of outliers in the

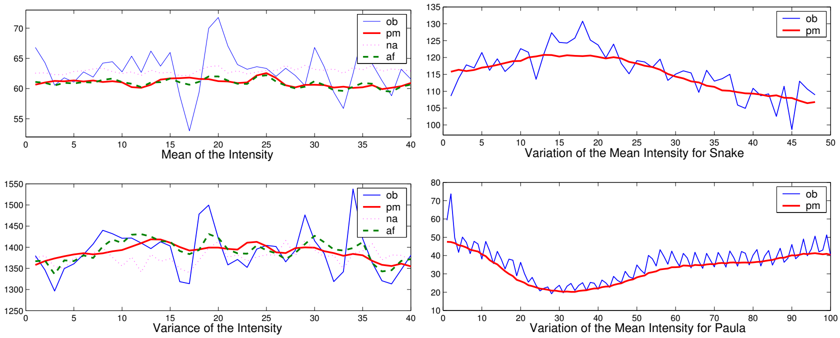

Fig. 5. Comparison for different scenarios of the mean and variance (a) (*Tunnel*). Variations of the mean (b)(top: *Snake*, bottom: *Paula*).

Table 1. Compressions ratios between the original compressed sequences and the restored ones. Better restorations are obtained for higher compressions ratios.

Sequence	pm	v	hv
Rory (shot 11)	17%	9.9%	20.5%
Rory (shot 13)	12.4%	10.7%	17.2%
Rory (shot 16)	15.2%	8.0%	17.7%
Rory (shot 19)	4.3%	8.0%	9.0%
Tunnel	48.6%	-	45.6%
Paula	2%	-	12%

estimation of our brightness distortion function. These algorithms are fast and coupled with the use of cheap graphics hardware, we can reach near real-time performance on standard computers for most real image sequences.

A The Factor $\mathcal{R}(u, v)$

We want to show in a first case that $\mathcal{R}$ can actually hold condition (8) without any precondition on the outliers distribution. The only hypothesis made here is that the *proportion* of inliers for each color in both images is greater than the proportion of outliers: if $(u, v) \in \mathcal{I}$, we have $\mathcal{P}(u|v) > \frac{1}{2}$ and $\mathcal{P}(v|u) > \frac{1}{2}$.

$$\text{thus,} \quad \begin{cases} \text{if}(u,v) \in \mathcal{O} : \mathcal{P}(u|v) < \frac{1}{2} \\ \text{if}(u,v) \in \mathcal{O} : \ \mathcal{P}(v|u) < \frac{1}{2} \end{cases} \Rightarrow \begin{cases} \text{if}(u,v) \in \mathcal{I} : \mathcal{R}(u,v) > \frac{1}{4} \\ \text{if}(u,v) \in \mathcal{O} : \mathcal{R}(u,v) < \frac{1}{4} \end{cases} \tag{11}$$

The condition (8) is clearly fulfilled. But as the assumptions made on the amount of inliers is very strong, we would like to examine another case where a color might be occluded by more than 50% of outliers. As the case illustrated by figure

6-a is a very frequent one, we consider below its case study. In this problem, occlusions, all of color u_0, are only present in image I_n. Therefore the only non null probabilities are situated in the inlier parts $\mathcal{P}(u, f(u)) \geq 0$ and in the outliers parts corresponding to the monocolor occlusions $\mathcal{P}(u_0, .) \geq 0$. Thus if we introduce an inlier pair $(u_1, v_1) \neq (u_0, v_0)$ with $v_1 = f(u_1)$ and $v_0 = f(u_0)$, for condition (8) to hold, we need only:

$$\mathcal{R}(u_0, v_0) > \mathcal{R}(u_0, v_1) \ , \quad \mathcal{R}(u_1, v_1) > \mathcal{R}(u_0, v_1) \tag{12}$$

Demonstration: Let α be the proportion of occluded pixels. If we assume that the colour distribution of the occluded region ('under' the blotches) is the same as the colour distribution in unoccluded regions, we have:

$$\mathcal{P}(u_0, v_1) = \alpha \mathcal{P}(v_1) \Rightarrow \mathcal{P}(u_0|v_1) = \alpha \tag{13}$$

However, as occluded regions are usually spatially coherent, the intensity distributions will be different and we must consider a weaker hypothesis:

$$(H): \quad \alpha \leq M \leq 1 \mid \quad \mathcal{P}(u_0, v_1) < M \cdot \mathcal{P}(v_1) \quad ; \quad \mathcal{P}(u_0|v_1) < M \tag{14}$$

Note that this hypothesis means when $M > 1/2$ that for one intensity v_1 in image I_n, the amount of outliers can be greater than the amount of inliers (ie. $\mathcal{P}(u_0|v_1) > .5$). Now, from (H) we derive directly:

$$\mathcal{P}(u_0, v_0) = \mathcal{P}(v_0) = \frac{\mathcal{P}(v_0)}{\mathcal{P}(v_1)} \mathcal{P}(v_1) > \frac{\mathcal{P}(v_0)}{\mathcal{P}(v_1)} \frac{\mathcal{P}(u_0, v_1)}{M} \tag{15}$$

$$\mathcal{P}(v_1|u_0) < \frac{M}{K} \mathcal{P}(v_0|u_0) \quad \text{with} \quad K = \frac{\mathcal{P}(v_0)}{\mathcal{P}(v_1)} \tag{16}$$

$$\mathcal{P}(v_0|u_0) + \mathcal{P}(v_1|u_0) \leq 1 \quad \Rightarrow \quad \mathcal{P}(u_0|v_0) < \frac{1}{1 + M/K} \tag{17}$$

$$\mathcal{R}(u_0, v_0) = \mathcal{P}(u_0|v_0)\ \mathcal{P}(v_0|u_0) = \mathcal{P}(v_0|u_0) < \frac{1}{1 + M/K} \tag{18}$$

$$\mathcal{R}(u_1, v_0) = \mathcal{P}(u_1|v_0)\ \mathcal{P}(v_0|u_1) < \frac{M^2}{K} \mathcal{P}(v_0|u_0) = \frac{M^2}{K} \mathcal{R}(u_0, v_0) \quad \text{cf. } H, 16 \tag{19}$$

$$\mathcal{R}(u_1, v_1) = \mathcal{P}(u_1|v_1)\ \mathcal{P}(v_1|u_1) = \mathcal{P}(u_1|v_1) > 1 - M > \frac{1 - M}{1 + M/K} \mathcal{R}(u_0, v_0) \tag{20}$$

Finally to fulfill conditions (12), we need $\frac{M^2}{K} < 1$ and $\frac{1-M}{1+M/K} > \frac{M^2}{K}$. Solutions for K and M are shown on figure 6-b. Stated briefly, the more frequent the occlusion's corresponding color (v_0) in the picture $n + 1$, the better the algorithm works. In particular, if the occlusion color appears in the most frequent color $(K > 1)$, the algorithm will work even if the full picture is half occluded $(M > .5)$. In practice of course, the assumptions made here do not necessarily apply exactly. We find that this means that inliers could be slightly attenuated. However, because this process is applied together with the MAP estimation discussed in section 4, flicker is properly suppressed in all examples tested so far as shown in the material[8].

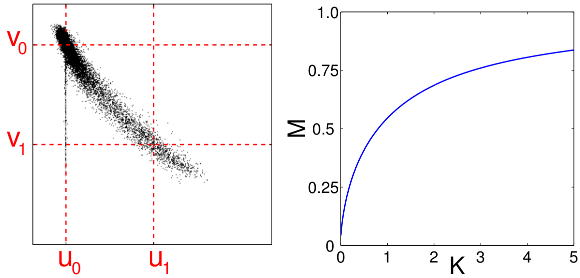

(a) (b)

Fig. 6. On (a), example of single color occlusion. On (b) $K = \frac{\mathcal{P}(v_0)}{\mathcal{P}(v_1)}$ versus M (in case of uniform occlusion, M corresponds to the proportion of occluded pixels).

References

1. E. Decencière. *Restauration automatique de films anciens*. PhD thesis, Ecole Nationale Supérieure des Mines de Paris, December 1997.
2. R. Gonzalez and P. Wintz. *Digital Image Processing*. 2nd edition, 1987.
3. P.J. Huber. *Robust Statistics*. John Wiley and Sons, 1981.
4. H. Jin, P. Favaro, and S. Soatto. Real-Time feature tracking and outlier rejection with changes in illumination. In *ICCV*, pages 684–689, July 2001.
5. A. C. Kokaram, R. Dahyot, F. Pitie, and H. Denman. Simultaneous luminance and position stabilization for film and video. In *Visual Communications and Image Processing*, San Jose, California USA, January 2003.
6. S.-H. Lai and M. Fang. Robust and efficient image alignement. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 167–172, Fort Collins, Colorado, June 1999.
7. B.H. Juang L.R. Rabiner. An introduction to hidden markov models. *IEEE ASSP Mag.*, pages 4–16, 1986.
8. Video Material. URL: `http://papabois.mee.tcd.ie/sigmedia/publications/`.
9. V. Naranjo and A. Albiol. Flicker reduction in old films. In *Proc. of the 2000 International Conference on Image Processing (ICIP-2000)*, September 2000.
10. T. Ohuchi, T. Seto, T. Komatsu, and T. Saito. A robust method of image flicker correction for heavily-corrupted old film sequences. In *Proc. of the 2000 International Conference on Image Processing (ICIP-2000)*, September 2000.
11. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C - The Art of Scientific Computing*. Cambridge University Press, 1995.
12. P.M.B. van Roosmalen. *Restoration of archived film and video*. PhD thesis, Delft University of Technology, October 1999.
13. P.M.B. van Roosmalen, R.L. Lagendijk, and J. Biemond. Correction of intensity flicker in old film sequences. *IEEE Transactions on Circuits and Systems for Video Technology*, 9(7):1013–1019, 1999.
14. X. Yang and N. Chong. Enhanced approach to film flicker removal. *Proceedings of SPIE Applications of Digital Image Processing XXIII*, 4115:39–47, 2000.