

TEN YEARS OF DIGITAL VISUAL RESTORATION SYSTEMS

Anil Kokaram

Electronic and Electrical Engineering Department,
Trinity College Dublin, College Green, Dublin 2
anil.kokaram@tcd.ie

ABSTRACT

With the growth in marketability of all forms of visual media comes the need to exploit archive material more effectively, and a requirement to guarantee picture quality regardless of the originating medium. Some kind of touch-up or restoration is rapidly becoming a “must have” in the visual production cycle. This paper reviews the development and impact of Digital Restoration Systems in the broadcast television and film industry over the last ten years. It presents some unifying ideas as well as the relevance and challenges of defect treatment in such harsh picture environments.

Index Terms— Video Restoration, Post-Production, Motion Estimation, Bayesian Inference, Example Based Image Processing, Inpainting, Object removal

1. INTRODUCTION

Research into the improvement of degraded signals is as old as the study of signal processing itself. In the main, researchers in image processing working in this area tend to associate the word “restoration” specifically with problems of noise reduction and deblurring. In a similar manner, Image Sequence restoration has become synonymous with multiframe superresolution techniques and authors such as Galatsanos, Katsagellos, Sezan, Tekalp, Stevenson are well known for early work. In a sense much of this work had relevance to the surveillance industry and principal applications included satellite resolution enhancement and vehicle surveillance.

The rise of Digital Visual Transmission in all its forms from the mid-1990’s caused broadcasters to be concerned about filling available channels. Once Digital TV launched with such huge success in Europe in 1999 led by Sky, and the DVD replaced tape by 2002, holders of visual content realised that their archives were suddenly much more valuable. As broadcasters reached toward archives to fill channels two things became apparent; i) dirty sequences cost more bandwidth than clean ones and ii) digital TV pictures are supposed to be a lot better looking than those supplied by the archives. Manual

‘washing’ of film was a possibility but less flexible than digital restoration of degraded material, and hence an industrial requirement for digital restoration was born. The key point to notice here though is that the problem of deblurring was never an issue, instead the key issues were film grain noise reduction, dirt and sparkle removal (DustBusting as it is now known), and stabilisation of brightness and position.

In the late 1980’s, a BBC Research and Development team led by Richard Storey[1] was the first to consider that degraded broadcast television material could be digitally restored. They concentrated on noise and Dirt removal without motion compensation. In the early 1990’s the first academic work in this area was undertaken [2] that considered motion estimation an integral part of the process. Work in automated digital restoration for degraded motion pictures has therefore been driven by a real need foreseen by industry long before it was a topic in university research laboratories. In the last five years the increased activity in this area has coincided with the notion of the ‘Digital Intermediate’ workflow for film post-production. This essentially means that all film material (modern or archived) is now scanned into digital form before any editing begins, and only in the final cut is the digital film record written to analogue media. In the ten years since 1996 automated digital restoration systems have been increasingly deployed, and university research in this area has increased. The area is rich with problems beyond deblurring and noise reduction.

A taxonomy of defects that have been treated in archive restoration research has been presented in [1]. This paper highlights some salient points that have prevented fully automatic restoration from becoming a reality and hence presents the challenges that remain to the community. To complete the picture an attempt is made to chart the progress of commercially available systems and point out their key innovations.

2. MISSING DATA: DIRT

Missing Data manifests in many forms in visual media. It is common in digital transmission and *error concealment* by picture or coefficient interpolation has long been studied in the MPEG community. In film originated media, missing data manifests as bright or dark patches that are caused by

Thanks to input from Peter Shallauer, Walter Plaschzug, Bernard Besserer. This work was funded by the EU Project PRESTOSPACE.




Fig. 1. Two consecutive images from a sequence showing severe motion blur due to fast motion. Motion vectors are superimposed on the second frame (right). This pathological motion completely confuses motion estimators leading to erroneous detection of large blotches. Observe in particular how the motion blurred regions are behind sharp stationary objects.

dirt sticking to the film in scanning or the emulsion being abraded. The design of algorithms for automated removal of small patches of missing data has been quite successful. Most algorithms are based on the DFD between the current and next, current and past frames. Large DFDs in both directions generally indicate defects. The idea can be articulated under a Bayesian framework [1]. Joyeux et al [3] introduce morphological operators to good effect for small defects.

Large defects remain difficult to automatically remove. Good reconstruction has been obtained by motion interpolation [1], but the detection of large defects is often confused with pathological motion in sequences. Thus fast moving hands, clothing or vehicles are always difficult for a Dirt Detector to ignore. What is missing in the literature is improved detection in the presence of pathological motion combined with a treatment of non-binary occlusion. Figure 1 shows the severity of this problem.

3. MISSING DATA: LINE SCRATCHES

Another very common form of degradation in film is Line Scratches. These are caused during developing or due to material being stuck in the film gate and abrading the material or smearing material over many consecutive frames. Early work by Kokaram [4] on detection and removal of this defect concentrated on spatial interpolation and detection. Detection was limited to straight scratches and relied on the horizontally impulsive nature of the artefact combined with the longitudinal correlation. Besserer, Joyeux et al [5] were the first to exploit the temporal coherence of the line scratch over several frames. They presented an excellent tracking algorithm that connected horizontally impulsive line elements that persisted across several frames at a time. Figure 2 shows the representation used.

However, removing line scratches is notoriously difficult. While convincing spatial interpolation can be achieved in a single frame, over several frames any error in reconstruction




Fig. 2. Left: The image sequence representation useful for detecting line scratches. Centre: Frame with scratch. Right: reconstruction of the line scratch in the region of local motion. Errors in reconstruction are much more visible in a sequence. Work on temporally coherent interpolation is important.

is clearly seen since it is correlated with the same position in many frames. Example based texture synthesis, famously introduced by Efros et al [6] can achieve very good spatial reconstructions, but temporally the result is poor if simply repeated on multiple frames. Figure 2 shows this problem in the region of local motion. An extension to the ideas of Efros to 3D synthesis for line scratch removal was introduced by Bornard et al [7] and was quite successful. It relied on using local motion around the line scratch to fetch useful information outside that region in the next and previous frames. They used global motion only and hence there were difficulties with objects showing local motion. The problem is how to reconstruct the underlying missing data including reconstructing the motion convincingly over several frames.

Recent work by Irani et al [8], Kokaram et al [9] and Sapiro et al [10] has considered the problem of object removal in image sequences. While Kokaram takes a purely motion based approach, Irani and Sapiro both use a combination of more sophisticated methods relying on substituting 3D cubes of data from other parts of the sequence or 3D inpainting. These ideas could be applied to the removal of line scratches. However the nature of the defect is more severe than the cases considered for object removal thus far, in that motion around the region may not be periodic or may be highly varying over short distances. Line Scratch removal technology in industry therefore still requires manual review of manipulated pictures.

4. NOISE, GRAIN AND RESOLUTION

It is generally expected by researchers working on noise reduction algorithms, that the requirement for noise reduction is to completely remove any noise in the image sequence. This is not so. Flat pictures even with all the detail preserved, look bad. Film grain is essential for the film look. While most workers in the area concentrate on preserving detail but removing noise, what is required by those in the post-

production industry is in fact noise reduction that reduces the level of grain but leaves some behind. The practice today is to remove the grain as much as possible, then add grain back into the picture of the type that is required. Compounding the problem is that film grain is a multiplicative process, and the amount of degradation/grain varies with brightness. There is more grain at high brightness levels than at low brightness levels. There is very little published work targetted at film grain noise reduction in particular. One might expect that most of the existing noise reduction techniques could be adapted with these new requirements, but reducing grain levels while leaving the correlation structure of the grain untouched, is a difficult issue [11].

After ten years of research and industrial development in film restoration, sharpening and enhancement are about to become much more important for broadcast applications. This is because of the advent of consumer high definition displays and DVD players. However, it is well known that high levels of noise mitigate against good sharpening results e.g. super-resolution. It remains to be seen how these problems can be resolved in degraded film material since the noise levels especially of film grain can be very high, and possess a non-white correlation structure.

5. SYSTEM ISSUES AND WORKFLOW

A high throughput for restoring films is desirable. Automated techniques are not as yet good enough in general to be robust in every sense and manual quality control is still important. Hence in every restoration job, there must be a diagnostic step to decide when to apply particular artefact removal processes and to decide on the best tool settings. Shot cut detection becomes a difficult task when large patches of a frame can go missing or when brightness fluctuation causes massive changes in brightness of the frame. Preliminary work on diagnosis of defects, a kind of content assessment task for degraded material was addressed in [12]. But the work is far from developed and issues like *detection* of shake and flicker remain open. Representations such as in Figure 3 can prove useful here.

6. AN EVOLVING INDUSTRY

In a sense, the industry for digital visual manipulation is divided into two categories. The television broadcasters are mainly interested in real time processing, while post-production houses (both film and video) who edit and apply effects to movies, are more interested that the picture is as good as possible. That means users in post-production tend to prefer general purpose hardware running various software tools for editing and compositing and hence are usually not concerned with real time processing. Interestingly enough, it is television broadcast hardware producers that had started in the late 1980's to produce restoration systems. The BBC



Fig. 3. *On the left is a set of horizontally projected slices of 200 images in an archive sequence, one image of which is shown on the right. The presence of flicker and shake can be diagnosed from the varying brightness and the jagged edges of features of the slices as they evolve in time.*

used their hardware (Debra) for dirt and noise reduction in-house in the late 1980's and it was real time¹. Digital Vision www.digitalvision.se launched a PC-based system for color correction and dirt/noise reduction in 1988 that exploited dedicated hardware, again allowing real time operation. Around 1997, Snell and Wilcox launched *Archangel*, dedicated hardware specifically targetting motion compensated, real time noise and dirt removal. Significantly, *Archangel* was the first hardware based restoration system to arise from an EU Research Project: AURORA (1994-1998). Terranex Systems teranexlive.dimentians.com is a more recent arrival (~ 2002) using a massively parallel array of processors on a chip to create single hardware units that allow real time noise/dirt/scratch concealment. The algorithms implemented in hardware tend to be deterministic in nature; clever use of motion compensated filtering combined with simple decision making over multiframe. S&W were the first to introduce a real time hardware line scratch removal system which was very successful.

It was in 1997 that the first software based restoration systems emerged for film post-production. Four appeared almost simultaneously. Lowry Digital was the first to use a massive network of Apple Macs to achieve fast throughput and high quality film restoration. They used their own software designed systems and essentially were a high quality post-production unit for restoration. DUST, established in France took a similar approach. MTIFilm www.mtifilm.com was the first dedicated software system to appear for restoration marketed as a restoration system to post-production houses and film studios. They therefore hold the accolade of being the first to design a professional restoration interface in software that showed a timeline, before/after and so on. DaVinci/Digital Revival, emerging from a collaboration with Cambridge University in 1996 followed shortly after offering software for

¹Built out of discrete logic devices it was quite large indeed

restoration on networks of linux machines. HS-ART released their Diamant² software restoration system about the same time. Diamant was also a result of an EU collaboration and driven by the needs of archivists, it is perhaps the only self-contained software system for editing and automated restoration available today. The Diamant product introduced an interesting innovation to show a representation of the entire movie using horizontal projections of each image stacked horizontally (see Fig. 3). This proved surprisingly useful for restoration since flicker and shake in particular can easily be spotted with this representation. It is worthwhile to note that film scanner manufacturers like Philips, Sony, Imagica and Thomson all incorporate some level of software based Dirt and noise reduction in their systems today.

Since 2003 however, the software restoration space has become more interesting. The increasing speed of PCs and the large repositories of video found in *communities* like Google-Video and YouTube, implies that the need for video manipulation has become more mainstream. Software Plug-ins that enable restoration for consumer and professional software platforms like After Effects, Flame, Shake and FinalCutPro, are now available from Adobe Systems, Autodesk, The Foundry, RedGiantSoftware and GreenParrotPictures. The PixelFarm have recently also joined the bandwagon of post-production software manufacturers that have seen the growing niche of restoration systems attractive. The Foundry and RedGiantSoftware together with GreenParrotPictures were the first to use algorithms derived from the use of Markov Random Field priors in DustBusting and Motion Interpolation.

The growth in the industry is clear, and with the imminent release of BluRay High Definition DVD disks, there is no reason to suspect that the demand for restoration will change in the near future.

7. FINAL COMMENTS

This paper has placed much of the work in restoration thus far in context. It highlights some open problems that remain and shows the vibrancy of the industrial applications. A word of caution however: never underestimate the skill of the post-production artist. Those working on image editing daily are actually very good at spotting and correcting defects quickly. While they will never rival real time systems, their visual systems are still much better than current signal processing algorithms. In some environments therefore, where an output of one frame every 30 seconds is sufficient, they are almost impossible to beat.

8. REFERENCES

- [1] A. C. Kokaram, "On missing data treatment for degraded video and film archives: a survey and a new

bayesian approach," *IEEE Transactions on Image Processing*, vol. 13, pp. 397–415, March 2004.

- [2] A. C. Kokaram, *Motion Picture Restoration.*, Ph.D. thesis, Cambridge University, England, May 1993.
- [3] L. Joyeux, S. Boukir, B. Besserer, and O. Buisson, "Reconstruction of degraded image sequences. application to film restoration," *Image and Vision Computing*, , no. 19, pp. 503–516, 2001.
- [4] A. Kokaram, "Detection and removal of line scratches in degraded motion picture sequences.," in *Signal Processing VIII*, September 1996, vol. I, pp. 5–8.
- [5] Bernard Besserer and Cedric Thiré, "Detection and tracking scheme for line scratch removal in an image sequence.," in *ECCV (3)*, 2004, pp. 264–275.
- [6] Alexei A. Efros and Thomas K. Leung, "Texture synthesis by non-parametric sampling," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, September 1999, vol. 2, pp. 1033–1038.
- [7] Raphael Bornard, Emmanuelle Lecan, Louis Laborelli, and Jean-Hugues Chenot, "Missing data correction in still images and image sequences," in *ACM Multimedia*, December 2002.
- [8] Y. Wexler, E. Shechtman, and M. Irani, "Space-time video completion," in *2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'04)*, 2004, vol. 1, pp. 120–127.
- [9] A.C. Kokaram, B. Collis, and S. Robinson, "Automated rig removal with bayesian motion interpolation," *Proceedings of the IEE Journal on Vision, Image and Signal Processing*, vol. 152, pp. 407–414, Aug 2005.
- [10] K.A. Patwardhan, G. Sapiro, and M. Bertalmio, "Video inpainting of occluding and occluded objects," in *IEEE International Conference on Image Processing (ICIP)*, 2005, vol. 2, pp. 69–72.
- [11] S.I. Sadhar and A. N. Rajagopalan, "Image estimation in film-grain noise," *IEEE Signal Processing Letters*, vol. 12, pp. 238–241, March 2005.
- [12] A. Kokaram, Raphael Bornard, Andrei Rares, Denis Sidorov, Jean-Hugues Chenot, Louis Laborelli, and Jan Biemond, "Robust and automatic digital restoration systems: Coping with reality," *Journal of the SMPTE (Society of Motion Picture and Television Engineers)*, vol. 112, pp. 225–231, July/August 2003.

²<http://www.hs-art.com/diamant/index.html>