Choosing a Film Scanner

This document considers how film scanners operate and looks at the criteria that are important when choosing a film scanner. Even if an archive does not plan to purchase its own scanner, it is useful to know what type of scanner a contractor is using in order to assess where they might encounter problems. These guidelines should not be taken as recommendations for any one model of scanner; scanners are mentioned by name only as examples of different technologies.

The Theory

Resolution – how many pixels are needed?

The information content of a 35mm film frame and the scanning resolution needed to capture it has been the subject of much debate. It is a complex issue because the image made up from random microscopic clusters of metallic silver or colour dyes on a film is very different to an image constructed from discrete pixels. The ITU (International Telecommunication Union) conducted research[1] into the matter in the early part of the 2000s, and although it is not strictly possible to precisely describe the resolution of an analogue film image in terms of lines or pixels, the conclusion drawn from analyses of modulation transfer functions (MTF[2]), was that the maximum detail which can be recorded on and retrieved from a 35mm answer print[3] is around 2000 lines per picture height at an aspect ratio of 1.85:1, which equates to about 2700 lines per picture height for a full frame 4:3 film image. This is equivalent to an information content of just under 4k across the image. 4K can thus be considered to approach the resolution of a modern 35mm answer print. The study found that the resolution of the original negative was considerably higher, but that of a 'release' print, produced by the printing sequence negative→interpositive→internegative→print, was considerably worse. In subjective tests on release prints, it was found that the highest resolution that could be detected by assessors when the release print was projected was only about 875 lines per picture height.

[1] ITU (2001), 35mm Cinema Film Resolution Test Report, ITU-R Doc. 6-9/3
[2] MTF is an established tool for measuring the effective resolution of a film image. Unlike digital resolution, where there is an absolute cut-off point, the MTF is expressed in terms of the ratio of image contrast to object contrast.
[3] The study defines 'answer print' as one made directly from the original negative

It is, however, not a simple matter to relate the findings of this study to the resolution required for scanning. Digital sampling of an analogue image requires a sampling rate higher than the finest detail on the original (actually twice the highest frequency of the original, as described by the Nyquist theorem) but it can be argued at one extreme that there is no theoretical limit to the resolution required to perfectly render an analogue film image, right down to the microstructure of the film grain. In practice however, increasingly faithful reproduction of the image structure beyond a certain point may only lead to a more accurate representation of the underlying random structure rather than the image it forms.

The faithfulness of the digital scan is heavily dependent of a number of other factors such as the film transport mechanism, the scanner optics, the condition of the film, the skill of the operator, and so on. It is also influenced by the design of the imaging sensor, since a CCD array does not capture 100 percent of the image falling on it due to the finite spaces between the pixel elements (known as the 'fill factor'): the smaller the gaps between each element, the more accurate the rendition of the original. For this reason, a 4K sensor generally produces a superior 2K digital image than a 2K sensor. (The fill factor is typically bigger in CMOS sensors than in CCDs, but there are many other factors governing the choice between these two technologies.)

The conventional view is that with current scanner technology, to obtain an image acceptable for preservation[4], at least an 8K scan is required for a modern original 35mm negative, and at least a 4K scan for a 35mm print made from such a negative. A few film scanners are now (in 2021) offered with more than 4K capability, but managing the resulting amount of uncompressed data presents an extreme challenge to the digital infrastructure. We still remain some way from the ideal position where we are routinely able to capture **all** the significant data on a film with a comfortable margin, unlike audio archives who have been able to do this with audio for some years, commonly sampling sound well above the human audible range at 96 KHz or even 192 KHz.

For film archives a degree of compromise is required, primarily dependent on two factors: the quality of the original, and the purpose of the digitisation.

The Battle of the Ancre (1917), a coarse-grained duplicate made in 1931. ©IWM

[4] 'Preservation' here means a master copy made as a replacement for a vulnerable original.

Regarding the first, archive film varies widely in resolution depending on its age and generation, and 2K scanning may be sufficient even for a typical modern 35mm print, if generated from a duplicate negative. Even if the digitisation is for preservation purposes, 2K scanning may be entirely acceptable for a great deal of film found in archives, such as elderly dupes of silent-era cinema which are likely to be coarse-grained and lacking definition.

However, if the purpose of digitisation is not the creation of a preservation master (where the original film is to be retained in preservation-standard storage), then it may be acceptable to tailor the quality level of the scanning to the intended purpose, such as online access, DVD/BluRay or digital cinema. In order to determine the optimal scanning resolution for digital cinema, it first needs to be understood that the terms '2K' and '4K' are not necessarily equivalent in scanning and in projection. A DCP image is set to a standard 'container' size into which the image must be fitted, whereas the aspect ratio of a scanned image is usually adjusted to fit the original film. The 2K DCP container is 2048 x 1080 pixels, which is close to a 1:1.85 film scanned at 2K (2048 x 1107 pixels). If the film is in academy ratio (1:1.37) however, a 2K scanned image would be 2048 x 1495, so in order to fit this into a DCP container it must be scaled down to 1480 x 1080. Similarly, a 4K scan of the same film (nominally 4096 x 2990) needs to be scaled down to 2959 x 2160 in order to fit into the 4096 x 2160 DCP container for 4K projection (and of course scaled down much further for 2K projection).

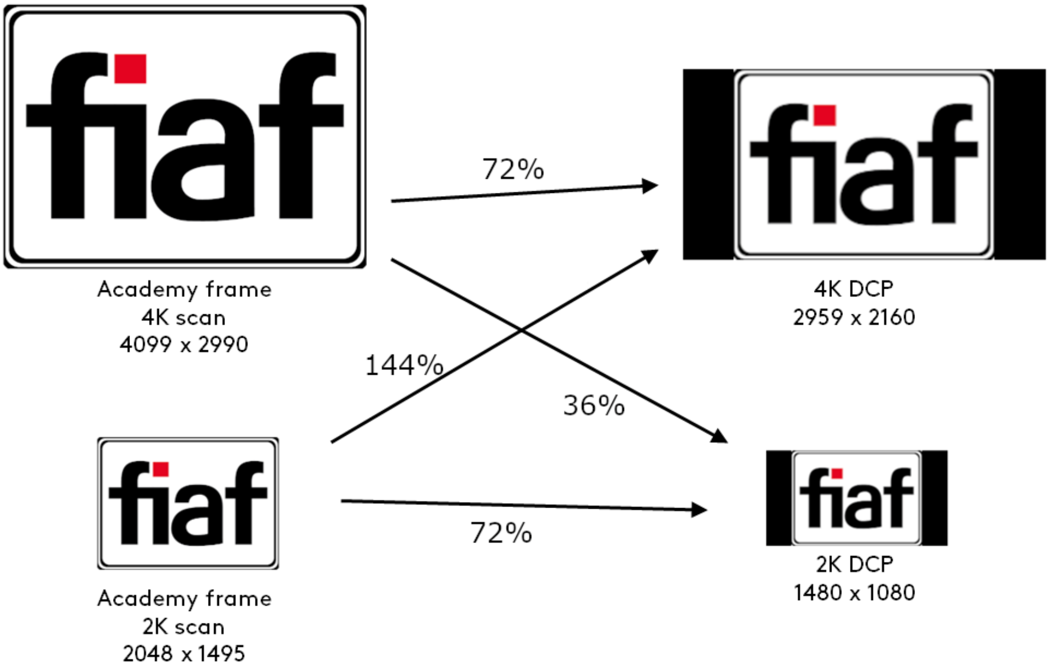

It can be seen that a 2K DCP can severely limit resolution, and therefore for academy ratio films there is an advantage in creating a 4K DCP even from a 2K scan. However the degree to which image resolution is significant for an audience is debatable given the number of other factors which affect human perception.

Bit Depth and Light Levels

Resolution is not the only thing that matters. The bit depth determines how precisely the colour and brightness of each pixel is recorded. An 8 bits-per-colour-channel system means that for each pixel, the brightness of each of the three colours will be set at one of a possible 256 levels between a maximum and a minimum point. A 10 bit system will have 1024 possible levels for each colour, so there are much finer steps between each colour value. 12 bit systems have 4096 possible levels. How the available bits are used depends on the system: linear data encoding has a set of even steps, whereas logarithmic encoding enables more precision in the lower end of the scale and less at the higher.

From the development of film scanning in the 1990s, it has been common to capture film scans at 10 bit (either log or linear), and this may be sufficient if the original has a relatively even scene-to-scene exposure, or if the exposure level is adjusted during scanning for each change in scene density throughout the reel, in which case little or no correction should be required to the resulting scan. (Most scanner software allows the exposure of each scene to be set in advance, or even adjusted 'on the fly'.) Without scene-to-scene correction during the scan, wide variations in exposure in the film may result in clipped whites or crushed blacks, or a compressed dynamic range which has insufficient data (not enough data points) to allow the image to be satisfactorily corrected post-scan.

However, using 10 bit data for film scans is very much a compromise and, as with resolution, archivists increasingly want systems which can record the entire dynamic range of a camera negative in one pass. High Dynamic Range (HDR) technology is now a common option with film scanners, holding out the promise of reaching this ideal in theory, if not always in practice. There are two approaches to this, one is to take multiple exposures of each frame at different light levels which are then combined into a single file, the other is to use a new generation of HDR sensors which can capture a wider range in a single exposure. Typically the resulting files have 16 bits per channel, but beware that just because a file is saved or converted to 16 bit, it doesn't mean the original data was captured at that bit depth – check the manufacturer's specification carefully!

At least one scanner is now compliant with ACES (the Academy Colour Encoding System) which was developed primarily for film production purposes. This system not only encompasses HDR, but also imposes a colour management system on the workflow so that there is consistency throughout the workflow. Scanner manufacturers on the whole seem uninterested in this system, but there is no doubt that a more rigorous approach to maintaining consistent standards in the film digitisation process is long overdue.

Frame size and position

Film frames can deviate considerably from standard sizes and positions, so all scanners allow adjustment of the size and position of the captured image, either by physically adjusting the optics or more commonly by selecting a portion of an overscanned image.

A scanner should have the option of overscanning the image so that the full frame and even the perforations are captured, the latter providing potentially useful data about image steadiness. However, capturing even part of the perforations may cause flare in the image, since the light passing through the perforations will reach the sensor unmodulated by any film density. Overscanning can also be valuable if the frame position varies through the reel of film, allowing the complete frame to be captured without adjustment during the scan so that its position can be corrected post-scan.

Overscanning is also important when creating digital preservation masters where the goal should be to capture as much information as possible. Indeed, for the ideal digital preservation master, it can be argued that the entire width of the film from edge to edge should be captured, as this will include such things as edge numbers, manufacturer's marks, information about perforation damage, and so on. Such a full width/full height scan of a sound film will in addition capture a potentially usable image of the soundtrack (see below), and it will at the very least have value as a preservation record of the format of the original soundtrack.

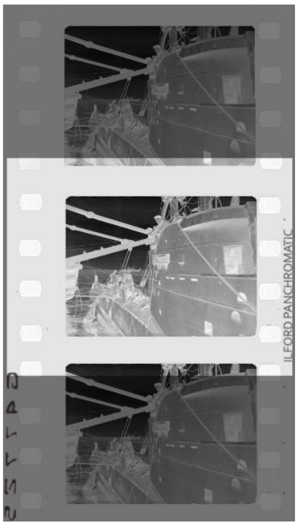

Edge to edge scanning of an Academy frame negative: in this case, approximately 60% of the available horizontal resolution is used for the image itself.

However, capturing the entire width of the film means the resolution of the image itself is greatly reduced, so this should be contemplated only with a high resolution scanner.

Scanner design

Sensors and image capture

The CCD (Charge Coupled Device) or CMOS (complementary metal oxide semiconductor) image sensors in scanners fall into two main groups, area arrays and line arrays.

Area array scanners operate much like digital cameras in that a digital photograph is taken of each frame of film as a discrete operation. Scanners designed principally with high quality production, preservation or restoration in mind, such as the Arriscan and Northlight (and the original Kodak Genesis scanner from the 1990s) are intermittent scanners, that is, the film is advanced one frame at a time and held steady in the gate while the image is captured, in much the same way as a film step printer. Other scanners such as the MWA Vario and Lasergraphics ScanStation run the film continuously across the gate and use a short-pulse stroboscopic light source to freeze the frames as they pass, in effect using flash photography to capture a moving image.

The sensors in line array scanners consist of a single line of pixel elements, and these scanners operate in a fashion analogous to continuous film printers in that the image of each frame is built up from a series of line by line exposures as the film moves over a narrow slit gate. This type of imaging is used by such scanners as the Spirit and the Goldeneye 4. A refinement of this approach is adopted by the DFT Scanity, which has multiple line array sensors, so that each slice of the image is captured by each line sensor in turn, and a composite result created from the output of all the sensors.

Scanner imaging is further divided between those that capture a full resolution image for each red/blue/green colour channel, and those that rely on a Bayer-pattern sensor.

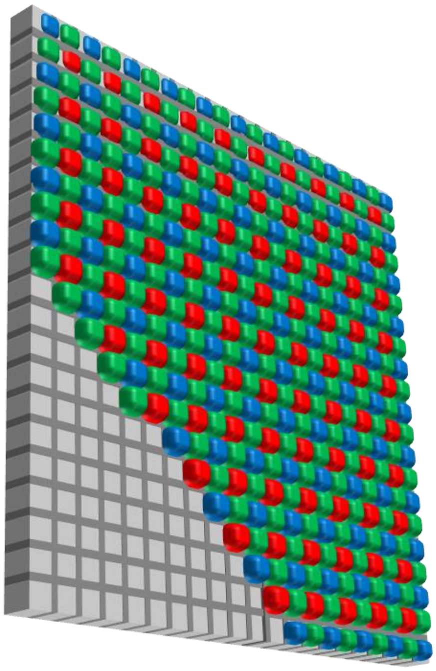

This type of sensor allows three colour channels to be extracted from a single area array by placing colour filters over each element in a fixed mosaic pattern. The resulting image requires 'de-mosaicing' in order to produce separate red/blue/green images, and although the algorithms used to do this are sophisticated, the resulting images are not strictly full resolution in each colour channel. Intermittent scanners with a single area array sensor (such as the Arriscan) operate by capturing three (or more) images of each frame using red, green, and blue light in turn in order to generate a full colour image. Scanners with line arrays generally use white light and then split the image by optical filters into red, green and blue elements, which are directed to separate sensors.

Light sources and scratches

Most scanners now have LED light sources. These are efficient, cool and reliable.

Regardless of whether the scanner uses white light combined with filters or separate red, green, and blue light, the combination of light source, optics, and sensor is unlikely to perfectly match the dyes in colour film stock. This is not a problem when dealing with modern film in good condition, but some scanners are much more effective than others at recovering badly faded colours. Also, where the film has less conventional colours (such as tinted films, two-colour systems, etc) the scanner may not capture a satisfactory result.

Ideally, it would be possible to tune the wavelengths sampled by the scanner to match the dyes in each film, but this is not currently an option with any scanners on the market. Some experts have demonstrated the effectiveness of multispectral scanning, that is, sampling the colour spectrum at many points to allow the selection of the optimum result. This can be extremely effective with 'problem' films, but it does require specialised systems and the ability to handle a huge amount of data.

The arrangement of the illumination and imaging optics has a significant impact on the degree to which scratches and other surface blemishes and marks are captured as part of the image. Most scanners have diffuse light sources because these reduce degree to which such defects are imaged, but it has been shown that there is a surprising amount of variation between different scanners in how effectively unwanted blemishes are disguised.

Some scanners have the option to use infrared imaging of the film to create a map of the dirt and scratches so that these can be subtracted automatically from the film image. This technique relies on the fact that a conventional colour film image is transparent to infra-red light, so when illuminated by infra-red light it appears clear apart from any scratches and marks on its surface. This approach cannot be used for black and white film, since the silver image is just as opaque to infra-red as the blemishes on the film.

Some scanners are available with a wet gate option, and this can be very effective at scratch reduction, although using wet gate does introduce additional complexity to the scanning process. There are claims that the use of wet gate reduces the definition of the image, but this is not well studied.

As a general rule it is preferable to minimise the capture of film defects in the scanning process, rather than attempting to remove them using digital restoration tools after scanning.

Steadiness

Intermittent scanners usually have a pin registered mechanism to keep the film perfectly aligned during scanning, but this may result in them being very intolerant of damaged and shrunk originals. The pin registration may have to be disabled for archive film, which then may necessitate the use of image stabilisation software on the resulting scans.

The way that area array and line array scanners respond to damaged film differs. In the former case, an image of the whole frame is captured at one time, so where perforation damage or tears in the film disrupts the smooth progress of the film, the result is that the image of the frame as a whole shifts in position. This can be rectified relatively easily post-scan. For a line array scanner, any disruption of the progress of the film over the gate will affect the way the image of the whole frame is built up from each line, and this can result in warping or rippling of the image. Such scanners have software designed to limit this effect, but it is not always completely effective.

In either case, it seems all scanners have a degree of difficulty keeping the image perfectly steady where there are physical splices in the film, resulting in the image 'bouncing' (albeit very slightly) at splices. There have been a number of reports that this can be particularly problematic when scanning 16mm A and B roll masters.

Transport

Most scanners are now capable of handling old and fragile film without damaging them, typically using a friction drive roller and a carefully controlled take-up. Continuous motion scanners need some method to determine the precise position of the frames as they pass over the gate: some use optical systems to detect the position of the film's perforations, using either visible light or infrared light, and others have a free-wheeling sprocket roller which is used by the scanner software to calculate the frame position and transport speed. Both approaches are effective, although the optical system is likely to be more tolerant of perforation damage. Scanners which are based on traditional editing tables (Steenbeck, Moviola, etc) which use sprocket rollers to drive the film are best avoided for elderly film: these machines were designed for handling new work prints rather than fragile masters, and there is a high risk of damage to the film.

Speed – how fast does the scanner need to run?

How important is the speed at which the scanner runs will depend on whether an archive is using it for individual high quality projects such as restorations, in which case a slow scanning speed will not impact greatly on the overall workflow, or for a high-volume project to digitise the content of the archive, in which case a fast scanning speed is an advantage.

The first generation of film scanners in the 1990s, such as the Kodak Genesis, operated at speeds of slower than one second per frame, resulting in multiple days to scan a feature film. Modern intermittent scanners may still only manage a few frames per second, and while these may be excellent for preservation or restoration, they are probably not ideal for a large volume digitisation project.

Continuous scanners now operate at speeds up to and beyond real time, although the archive must be aware that such speeds demand very fast systems for capturing and storing the resulting data, and that reaching the highest resolution advertised by the scanner may require a slow scanning speed.

If the scanner uses a conventional optical (photoelectric) sound head, the scanner should be running at normal projection speed in order to capture the sound, otherwise the sound quality is likely to be compromised.

Sound – as important as the image

Too often the soundtrack is considered as an afterthought. The film's soundtrack should be captured at a quality consistent with that of the image, but the way that the soundtrack is captured on a film scanner is not necessarily optimal for the type of track being scanned.

Scanner manufacturers take a variety of approaches to sound capture. Scanners capable of running at 24 or 25 fps may be fitted with a conventional optical sound head which will capture adequate sound from a film print with a good optical track. However this type of sound reader may be less than ideal if the track is scratched or damaged, and it is not designed to capture sound from an optical *negative* soundtrack. Typically sound captured from a negative sound track using a conventional photoelectric sound reader will have a high noise level and will suffer from distortion, particularly of sibilants.

There are a number of other options for soundtrack capture, either integrated with the scanner or as standalone systems. These systems work by capturing an image of the soundtrack and processing it in order to create an audio signal. The advantage of this is that the processing allows a degree of manipulation to the image of the track, such as removing scratches and blemishes, adjusting the position, the width, the light levels, and even emulating the image spread effect that is important when a print is made from the soundtrack.

One increasingly common approach is to assemble a whole soundtrack by stitching together the individual fragments of the track captured alongside each overscanned frame. This can then be processed to produce a sound file. This approach allows sound to be captured on an intermittent scanner even where there is no separate sound head, but it does mean that each frame must be overscanned enough to capture an overlap with the adjacent frame so that the pieces can then be stitched together, which may impact on the resolution (see above.)

The more sophisticated of these systems offer a great deal of control over the process, and with the right expertise can produce very good results.

Testing a scanner – how to do it

No single scanner is a perfect solution to digitising all motion picture film. Because scanners currently operate at a similar quality level to film in terms of resolution and dynamic range, and film originals are highly variable in their type and condition, digitising film remains a compromise. A cheaper scanner may produce perfectly acceptable results in some circumstances, while the most expensive scanner may be unsuitable in others.

Film scanner manufacturers do not generally provide much objective information regarding the performance of their scanners in terms of exposure response curves, resolution, processing etc. In addition there is a lack of standardised test material, compounded by the fact that film originals come in a huge range of types and in all sorts of conditions. It is therefore difficult to apply a standard test to a scanner and its supporting software, so assessing whether a scanner is suitable for your archive is a matter of carefully evaluating the requirement, and then assembling a collection of test pieces to try out on the candidate scanners.

Experience shows that asking the scanner manufacturer for a demonstration with their own test material is not a useful exercise, neither is it ideal to send your own test reel to them because it is important to be present when the test material is scanned so that the performance and operation can be evaluated. A scanner may produce an excellent result from an original in poor condition, but if this takes the operator the whole day to achieve, it is not likely to be suitable for this type of film in practice.

Test material should include both well- and poorly-exposed film, dense images (particularly negative), damaged film, scratched film, dirty film, shrunken film, film with imperfect splices, both good and faded colour, film with changes in frame position (for pre-sound era material), film with unorthodox frame ratios, film with combined sound, separate soundtracks, and test charts. If you have a nitrate collection to digitise, you might consider including nitrate material if shipment and test site allow. In short, make sure that you test every type of film that you are likely to want to put on the scanner.

Whether testing scanners with HDR or not, attention should be given to dealing with film that has a lot of scene-to-scene exposure variation: 'best light' scans of such material are likely to result in clipped whites or crushed blacks, so the amount of time and expertise needed to make exposure adjustments through the reel need to be assessed, and the results closely examined.

When assessing the resulting scans, pay attention to how blemishes and marks on the original are captured (there can be a big difference between scanners), how steady the image is, especially at splices, how good the resolution is, and so on. The eventual decision will be guided by the quality of the scans, by the suitability of the scanner for the type of film to be digitised, the ease of operation, and by the intended purpose of the digitisation.

FIAF 2021