Fédération Internationale
des Archives du Film
International Federation
of Film Archives
Federación Internacional
de Archivos Fílmicos

Choosing a Film Scanner

How does an archive scan the images on film, store them, and present them in a cinema or stream them on the web while maintaining high quality and preserving the original viewing experience? The whole chain of digitisation and related activities is a complex one, but in this article we will look at aspects of film scanning and how to choose a scanner. Even if an archive does not plan to purchase its own scanner, it is worth knowing what type of scanner a contractor is using.

Screen resolution – what can the audience see?

One question that arises is how much resolution is required from a viewer's perspective? How much detail can a cinema patron actually see? If one looks at the requirements of a THX-certified[1] cinema, the first row must not have a viewing angle greater than 90° and the last row must not have a viewing angle narrower than 36°, when looking at a 2.39:1 Cinemascope image. "Normal vision" is defined as being able to resolve a detail covering 1/60 of a degree, but even a visual acuity of 1/90 of a degree is common. From this you can apply a calculation that leads to the number of pixels (horizontal and vertical) on a screen that a human can actually discern.

From a seat located half-way between the first and last rows, a person with a human visual acuity of 1°/75 can discern approximately 4K horizontal pixels across a Cinemascope image and a little bit less than 2K pixels over the vertical. Seated in the front row, one can see around 9K (although many people would find this uncomfortably close):

Maximum discernable pixels in a 2.39:1 Cinemascope image for human visual acuity of 1°/75

	Horizontal	Vertical
Back row of cinema (36°)	2792	1168
Middle row of cinema (52°)	**4192**	**1754**
Front row of cinema (90°)	8594	3596

[1] A THX-certified cinema complies with a set of standards defined by the THX company.

The same person in the same rows viewing a 1.85:1 image of the same screen height would be able to discern the same size of pixel, but there would be fewer pixels across the narrower image:

Maximum discernable pixels in a 1.85:1 widescreen image for human visual acuity of 1°/75

	Horizontal	Vertical
Back row of cinema	2162	1168
Middle row of cinema	**3245**	**1754**
Front row of cinema	6653	3596

And for a 1.37:1 image of the same height in the same cinema:

Maximum discernable pixels in a 1.37:1 academy image for human visual acuity of 1°/75

	Horizontal	Vertical
Back row of cinema	1601	1168
Middle row of cinema	**2391**	**1754**
Front row of cinema	4926	3596

From this it can be seen that even in the middle row of such a cinema, it should be possible for someone with good vision to see the individual pixels in a 2K screening, and such a person should be able to resolve the pixels in a 4K screening when sitting in the front row. From this, it can be deduced that the capabilities of the human visual system are not the sole factor in dictating the resolution of a cinema screening. Perhaps we should not be surprised, since viewers expect to see the image structure – the grain – in a conventional projected film image. The random appearance of grain, of course, is quite different from the fixed array of pixels in a digital image.

Aspect Ratios – what happens when creating a DCP?

Another issue to consider are the different aspect ratios throughout the chain. Let's start by scanning a 1.37:1 academy aspect ratio image in 4K (4096 x 2990, or 11.7 Megapixels). If you want to project this image wrapped in a DCP package, you will need to scale it down to fit into the 4K / 1.89:1 container, which is 4096 x 2160. The result is an image of approximately 2959 x 2160, or 6.1 Megapixels in size. With digital cinema technology today, this is the highest number of pixels you can project from an academy format film image onto a screen. If you instead create a 2K DCP, which has a container 2048 x 1080 pixels in size, a 1.37:1 image will be 1440 x 1080, or 1.6 Megapixels, which only contains 13% of the spatial information of the original scanned image. This makes a compelling case for producing 4K DCPs from archive films.

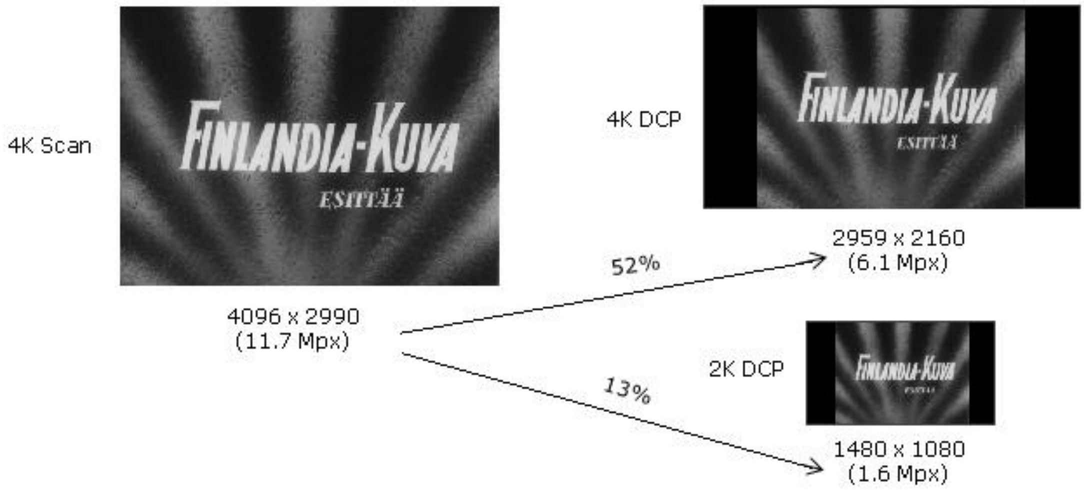

Resolution – how many pixels are needed?

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 a grain on a film is a very different kind of picture element than a digital pixel sample. The ITU (International Telecommunication Union) conducted research[2] 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[3]), was that the maximum detail which can be recorded on and retrieved from a 35mm answer print[4] 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 better, 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.

It is, however, not easy to relate the findings of this study to the resolution required for scanning. Perfect digital sampling of an analogue image requires a sampling rate higher than the finest detail on the original, and it could 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 grain structure beyond a certain point only leads to more accurate digitisation of the grain, and not of the image it forms. Broadly though, one can conclude from the study that *at least* 4K resolution is required to digitise a modern original negative for restoration or preservation purposes, particularly if manipulations such as stabilising and de-warping, which can lead to an effective loss of resolution, are carried out. On the other hand, 2K resolution is likely to be good enough for making

[2] ITU (2001), 35mm Cinema Film Resolution Test Report, ITU-R Doc. 6-9/3
[3] 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.
[4] The study defines 'answer print' as one made directly from the original negative

digital access versions of 35mm films. In fact many people find that capturing a film at HD video resolution produces entirely satisfactory 2K DCPs given that the 16:9 aspect ratio of HD video is very close in pixel dimensions (1920 x 1080) to the 2K DCP container. It should be remembered that 2K scanning does save a lot of time and resources (bandwidth, storage, and processing time) compared to 4K scanning.

Speed – how fast does the scanner need to run?

Scanning speed is a key consideration during the digitisation process if the aim is to digitise the archive, rather than just carry out restorations on selected films. Every step—physical film preparation, film scanning, colour correction, sound scanning, sound syncing, file packaging and compression, and file transfer—takes time. Ideally, the scanning equipment should be powerful and fast enough to not impose its own restrictions for the overall processing schedule. Some older film scanners which were designed primarily for feature film production operate at speeds of slower than one second per frame and are not suited to a archive digitisation workflow. In a high-volume scanning scenario, there simply is no place for a scanner taking multiple days to scan a feature film. Some newer scanners aimed at the archive market are capable of high quality scanning at speeds up to real-time, although such speeds require very fast systems for capturing and storing the resulting data.

Bit Depth and Light Levels

The bit depth determines, in effect, how accurately the colour and brightness of each pixel is measured. 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 and 16 bit systems have 4096 and 65536 possible levels respectively. The more bits used, the larger the files of course.

10 bit scanning may be sufficient if little or no correction to colour balance or light levels is required to the resulting scan, but if there is a wide variation in exposure, there may not be sufficient information in the image to achieve an acceptable result when it is graded after scanning. The illustration below shows an exaggerated example of the effect of insufficient bits:

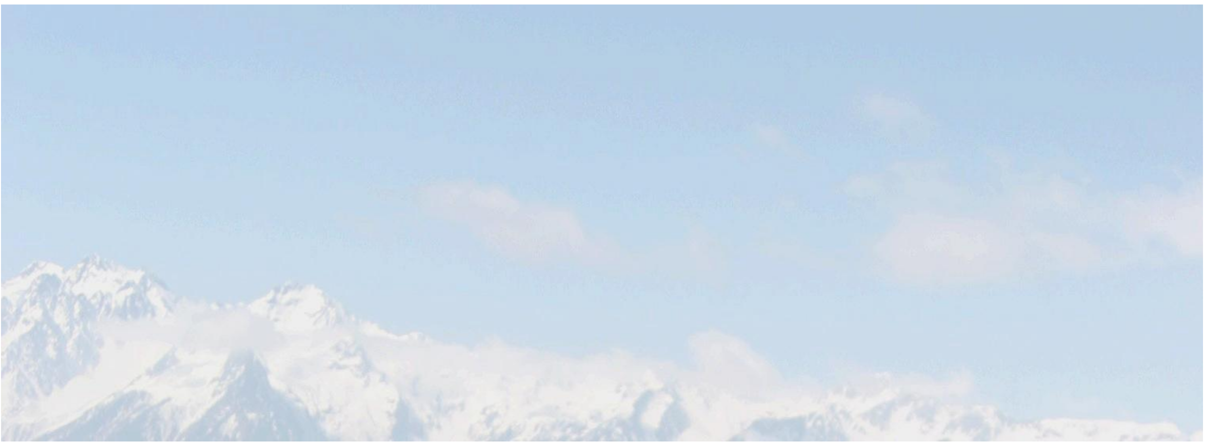

This image is very light, and has no dark tones

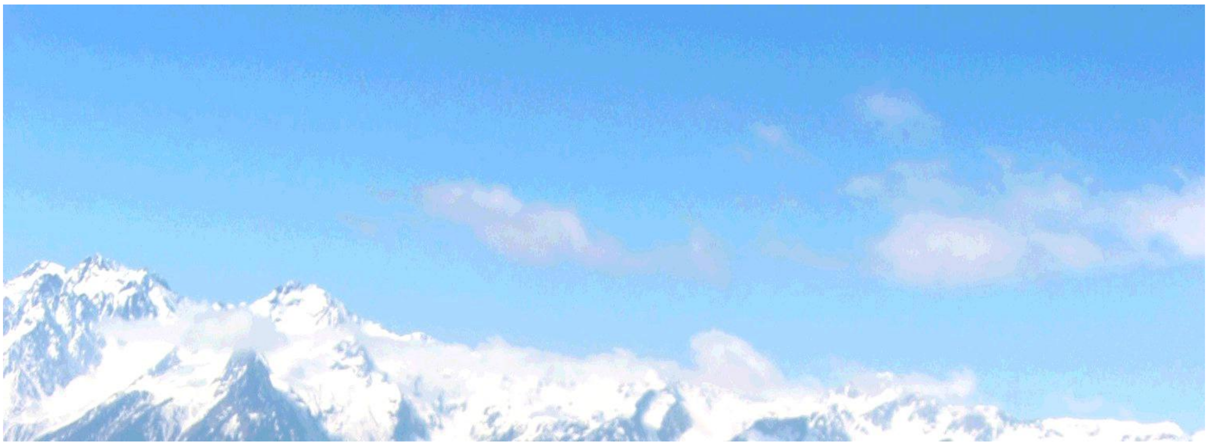

The same image, with brightness and contrast adjusted, now shows obvious
'contouring' or 'banding' between different brightness levels.

Some scanners allow the exposure of each scene to be set in advance, or may even allow exposure to be adjusted 'on the fly' in the same way that a telecine is operated; other scanners may only offer the option to set up a 'best light' exposure for the entire reel. If the resulting scan requires subsequent scene by scene correction to any significant extent, then scanning at higher than 10 bit depth is likely to be required. Higher bit depth is also recommended if restoration work is intended.

Some scanners perform poorly where the original is very dense (as is not uncommon for nitrate camera negatives, for instance), resulting in insufficient information in the shadows (or highlights where the negative image is reversed). Where an archive plans to scan such material, it is recommended that the scanner's performance with this type of original is tested.

Steadiness – registration pins needed?

Some scanners designed principally for digital film production use a pin registered intermittent mechanism to keep the film perfectly aligned during scanning. Such scanners tend to be slow, and may be very intolerant of damaged and shrunk originals. On the other hand, many scanners designed primarily for archive work use a continuous transport coupled with an optical method of stabilising the resulting image (or more accurately, stabilising the perforations of the scanned film). These scanners can produce very steady results even with quite badly damaged films.

Scanners often have a degree of difficulty keeping the image perfectly steady where there are physical splices in the film, with the result that the image 'bounces' at splices.

Frame size – is it really 4K or 2K?

Few scanners actually capture an image precisely 4096 or 2048 pixels across for every type of film frame presented. Film frames can vary considerably from the standards, while the sensor in a scanner is fixed in size. While it is possible to design an optical system which adjusts the size of image hitting the sensor, in most scanners the optics are fixed and the size of the captured frame is adjusted by selecting the range of pixels saved. Some older '2K' scanners, many of which are still in use, do not have a full 2048 pixel width sensor, and create '2K' scans by artificially increasing the resolution by data processing. Other scanners may only offer 2048 pixels width for a full frame 35mm film, so that the smaller 35mm Academy frame image may only be around 1820 pixels across. Ideally a scanner should be capable of overscanning the image so that it can be sized to fit the film frame after scanning. This is particularly valuable if the framing position varies through the reel of film.

Sound – as important as the image

Until recently, capturing film soundtracks was considered to be an entirely separate process. Scanners with intermittent transport running no more than a few frames per second could not at the same time capture sound intended to run at 24 or 25 fps. Some scanners now offer sound capture as an option. Some are able to do this because they can scan the picture at normal running speed (24/25 fps) and therefore can read the sound at the same time through a separate audio head. Others offer the option of running the film through the scanner a second time at sound speed, and some capture an image of the soundtrack while the picture is being scanned, and then process it in some way to produce an audio file. Scanners offering the last option may not be able to produce the same audio quality as those using conventional sound readers.

Direct capture of optical negative soundtracks, of course, requires specialised technology which emulates the image spread effect of conventional optical printing. Such technology is not currently available on any film scanner.

Scratches – pre- or post-scanning?

It is generally accepted that it is better to lose as many of the scratches and blemishes on a film as possible at the scanning stage, rather than by applying digital restoration software to the scanned images. Any digital restoration software working on the scanned image must necessarily attempt to distinguish the wanted image from the unwanted blemishes which have become an integral part of the image, and is therefore prone to error.

The diffuse illumination used by most scanners will help to minimise surface abrasions on the film base, and many scanners also offer other options for removing scratches at the scanning stage. One technique involves capturing an infra-red image of the film, which is effectively a map of the abrasions that can then be subtracted from the visible image. Note that this only works with colour film, as the black and white silver image is opaque to infra-red. Wet-gate scanning is also a possibility, but not many scanners offer this, and concerns have been raised about its effect on the resolution of the image.

Testing a scanner – how to do it

Unlike traditional photochemical technology, and unlike the world of professional stills imaging, film scanner manufacturers do not generally provide much information regarding the response of their scanners to exposure levels. Nor does the digital film world have standardised systems for controlling and managing colours and levels. It is therefore difficult to apply a standard test to a scanner and its supporting software, and in any case, there is a lack of standardised test material.

Before choosing what scanner should be used on an archive’s film, test footage should be scanned and assessed. Experience shows that asking the scanner manufacturers themselves to scan a test reel is not a very useful approach. It is better to find someone with working experience in the particular type of scanner to carry out the test scans, and ideally to attend the test scanning session. Test material might include both well- and poorly-exposed footage, dense images (particularly negative), damaged footage, scratched footage, footage with imperfect splices, both good and faded colour, footage with changes in frame position (for pre-sound era material), footage with sound, etc..

Both the results of the test and the effort required to produce them should be assessed. Both of these will necessarily be subjective assessments to a large extent, but in assessing the operation of the scanner, the time taken to set up the film on the scanner, and the effort required to deal with such things as large changes in exposure should be considered along with the actual scanning speed. Assessing the results will largely be a matter of opening the files in a suitable application and using tools such as histograms to examine the tonal range.

FIAF Technical Commission 2012