

EBU - Tech 3335 : **Methods of measuring the imaging performance of television cameras for the purposes of characterisation and setting**

Alan Roberts, May 23 2012

SUPPLEMENT 002 : Assessment of a Nikon D800 DSLR

The Nikon D4 is a medium-sized, full-frame, digital single lens reflex camera. As is becoming common, it also offers HDTV recording. Although there are few of the usual controls that go with video cameras, it was thought worthwhile to check its performance as an HDTV camera. The results were mixed.

Tests were done on a camera body (serial number 601303) with two lenses (Nikkor AF-S 20-70 1:2.4G, serial number 546609, and AF-S 24-120/4G VR ED, serial number 62001283). The test procedures were as described and recommended by the EBU, in Tech 3335 (<http://tech.ebu.ch/docs/tech/tech3335.pdf>).

The camera is medium-sized (146x123x81mm) and fairly heavy (1kg) for a DSLR. Power is from a lithium-ion battery, rated at 7V/1.9Ah, although an adaptor is available to run it from eight AA-sized batteries. A mains adaptor is available to replace the battery for extensive shooting. Recording is on to 2 card slots, 1 Compact Flash, the other SD.

Recording video is at 1920x1080, 1280x720, or 640x424, at a variety of frame rates (including 25Hz). Recorded files are in QuickTime (MOV) with H.264 compression (AVC, MPEG4). There is no facility for external time-code or genlock. Real-time video output is via HDMI, which appears to be of the same quality as video recording. Therefore it is possible to make better quality recordings externally than can be done via the in-built system. There is a small built-in microphone (adequate for note-taking and guide-sound recording) and an external mic socket (3.5mm). Connection to a computer via USB allows rapid file transfer, although it is probably quicker to swap cards for downloading. Hot-swapping caused no problems. The camera also has an optional wireless connection system.

External controls are by push-buttons and rotary switches. Most of the controls are to do with stills-shooting, and have no effect on video-shooting. Some controls can be assigned to preferred functions. The menu structure, although quite comprehensive, has few items specific to video shooting. Therefore, in this document I have not listed the full menu contents.

Colorimetric and Resolution requirements of cameras

Alan Roberts (May 14~22 2012)

ADDENDUM 72 : Assessment of a Nikon D800 DSLR

The D800 is a full-frame Digital Single Lens Reflex camera, with HDTV recording facilities at 1920x1080 (progressive-only at 29.97, 25, 23.976Hz), 1280x720 (progressive-only at 59.94, 50, 29.97 and 25Hz) and 640x424 (progressive-only at 29.97 and 25Hz). Monitoring and control of the camera are very different from a conventional video camera, making it rather unsuitable for many video purposes. The small LCD panel display was adequate for framing shots, and for getting focus provided the depth of field is short. Otherwise, it is best to use an external monitor fed via HDMI (Setup tools>HDMI).

The camera's single CMOS sensor is 35mm full-size, 35.9x24mm, and has approximately 36.8 million photo-sites of which 36.3 million are 'effective', i.e. used to make pictures (the remainder being blanked off to provide thermal and black level data). The specification lists the choices of resolutions for stills-shooting, the maximum of which is 7,360x4,912, which is 36,152,320 and close enough to 36.3 million. Allowing for, say 52 blanked photo-sites along each edge, the overall size could be about 7,412x4,964, a total of 36,793,168 which is close enough to 36.8 million. Since the image size did not change when changing from stills to video or between video standards, the video feed must be generated from most of the sensor, probably 7,360x4,140 using a 3.833:1 (23:6, not an easy relationship) down-scaling filter for 1080-line and 5.75:1 (23:4, also not easy) for 720-line. It seems likely that simpler relationships were chosen, using slightly less than the full sensor dimensions, however, in theory, this should be a fairly simple conversion because the image is highly over-sampled.

All testing was done by recording onto Compact Flash in the camera, since this is the mode in which the camera is easiest to use and will probably be used. Monitoring was done on a 42" plasma display via HDMI.

1 Colour performance

The camera menus provide control over the system primaries (Menu button > Shooting Menu > Color Space) offering sRGB and Adobe RGB. Since the sRGB primaries are those of ITU.709, the world-wide HDTV standard, only this setting makes sense for shooting video. Effectively, this setting is a choice of matrix rather than of primaries.

There are also Picture Controls (press the button with a key symbol) offering a choice of 6 settings. 'Standard' was used for all the tests. However, by then pressing the 'OK' button, the stored settings can be modified (Sharpening, Contrast, Brightness, Saturation, Hue). Most of these settings were left at the default, factory value, only Sharpening was explored.

Using sRGB and Standard Picture Control, the colour performance was perfectly acceptable on a television display.

2 Resolution (1920x1080)

The camera was exposed to a circular zone plate test chart, containing patterns to test luma, R G and B, and chroma channels. Only one quadrant of the luma pattern is shown here (Fig.1).

This showed some low-level horizontal coloured aliasing, apparently centred on about 2,260, and higher-level vertical aliases centred about 1,270. Some aliasing is to be expected in any camera having sensors not pixel-mapped to the image format, but this asymmetry is interesting for it shows that different processing has been used for horizontal and vertical down-scaling, presumably to save on the power

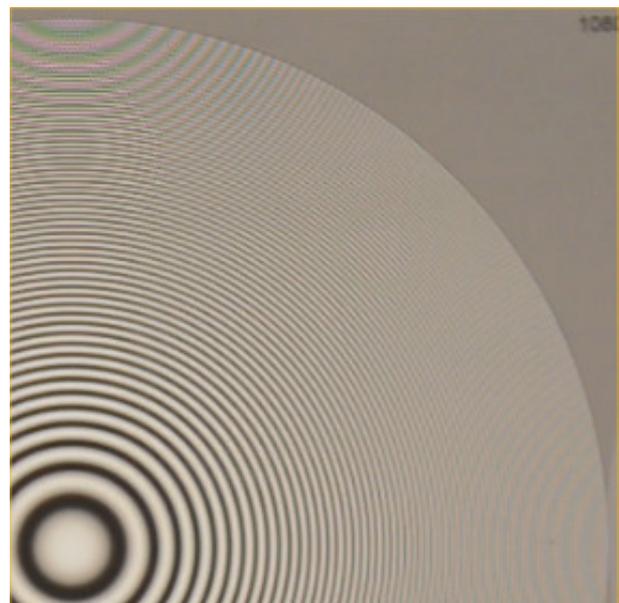


Figure 1 Resolution 1080p, default detail

consumption required for better conversions. The pattern centred on 816 vertically and the similarly-placed horizontal one are probably due to harmonic distortion caused by a mismatch between the camera gamma curve and the assumed transfer characteristic of the standard television display.

While it would be interesting to chase these numbers to find their origins, it hardly helps in the circumstances.

This less-than-ideal down-scaling, and the asymmetry, is to be expected if the size of the down-scaling filter is not very big. While the level of horizontal aliasing is fairly low, the level of the verticals is worrying, since it will be highly visible on moving images, and can be expected to cause problems in MPEG-type compression. The clean resolution does not substantially exceed 1,355x762, and the presence of the coloured aliasing implies that the resolutions of R and B are not the same as G.

The coloured aliases are the result of unequal RGB resolution in the Bayer-patterned sensor. The resolution chart clearly shows this (Fig.2). Note that the red pattern (and therefore also the blue pattern) has aliases centred on about 2,260:0 (i.e 2,260 horizontally and 0 vertically) and about 0:1,270 (0 horizontally and 1,270 vertically), but otherwise shows square resolution up to about 1,355x762, while the green pattern has aliases centred on diagonal frequencies about 2,260:1,270. This is exactly as is expected for a Bayer-patterned sensor.

The resolution chart also contains two smaller patterns, which explore the frequency range up to 3,840x2,160, again, only one quadrant is shown (Fig.2 centre top). Again, there are aliases visible centred on about 2,198 horizontally and 1,236 vertically, establishing an approximately 2.24:1 relationship with the sensor dimensions. It is not at all clear how this comes about, but it is definite evidence that the down-scaling filter is not big enough, and is not rejecting the high-frequencies which the high-resolution lenses can deliver to the high-resolution sensor.

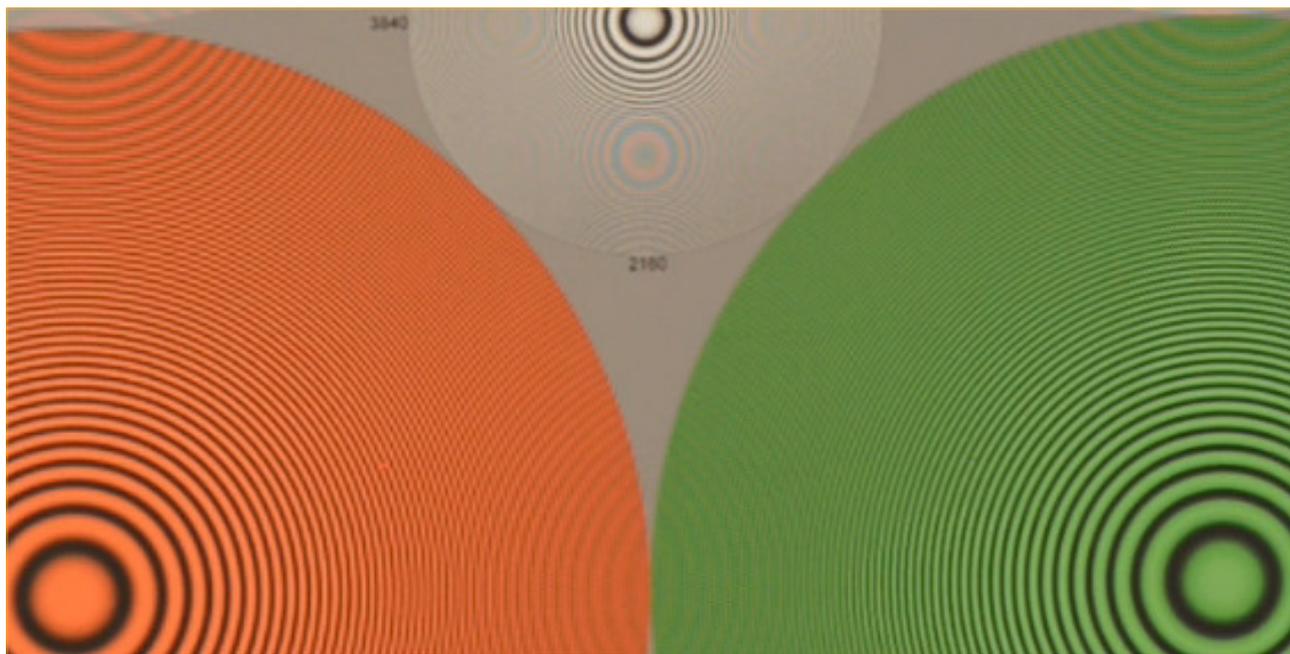


Figure 2 Resolution 1080p, default detail

The delivered resolution is not full 1920x1080, and aliases are clearly visible, which is not good for serious HDTV shooting. A proper video camera of the same price could normally be expected to perform rather better.

Spatial aliasing is caused by less-than-ideal scaling or interpolation, and in still images may not be a big problem, because they don't move. However, aliases in moving pictures are much more of a problem because, when the image moves, the aliased frequency content moves in the opposite direction to the image motion, causing a rippling effect on edges. Since motion-sensitive compressors such as MPEG2 and MPEG4 depend on the cleanliness of edges to measure motion, these aliases can cause the compressor to

allot undue bit-rate to motion and/or result in excessive compression artefacts. Pictures with aliasing at the levels seen here are probably not acceptable in HDTV.

3 Resolution 1280x720

This time, the story much less satisfactory (Fig.3).

The aliases centred on 2,260 and 1,270 are still visible, resulting from down-scaling to 1920x1080, and there is a very strongly coloured vertical alias centred on about 739. This results from the further down-scaling to 720p, and is inevitably more problematic for serious programme-making.

The resolution is clean only up to about 1,030x580. Given the low achieved resolution and the level of multiple aliases, this mode is not recommended for HDTV shooting.

The vertical alias at about 739 is caused by aliasing in the colouring channels. Fig.4 shows the red and green patterns, and the alias at 739. The double-frequency pattern (top-centre) also shows an alias centred on about 1478, again caused by aliasing in the colouring channels. Undoubtedly, both these patterns are caused by simplifications in the down-scaling filters for 720p. Note that these aliases do not flicker, but they are very strong and highly noticeable.



Figure 3 Resolution 720p, default detail

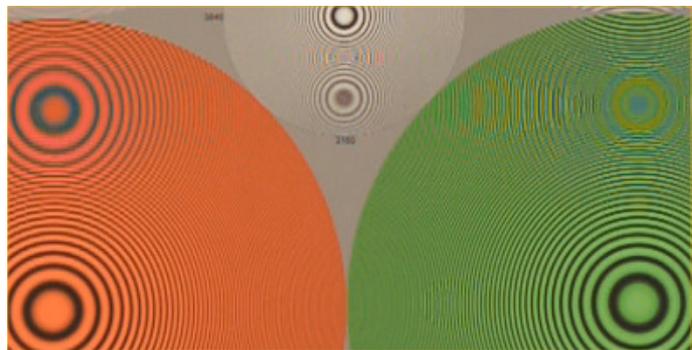


Figure 4 Resolution 720p, R and G

4 Gamma curve

Normally, the distribution of noise levels matches the slope of the gamma curve, so it would have been good to be able to use a standard curve, but there were no such options in this camera. Therefore, several measurements were made of the patches of the grey scale on a Colorchecker chart (which has known reflectances) at exposures from F/2.8 to F/22, all at ISO 500 with shutter set to 1/50 second. Some slight manual adjustments in the exposure calculations were needed to get the curves to overlap nicely. Fig.5 shows the result, plotted on linear and logarithmic axes.

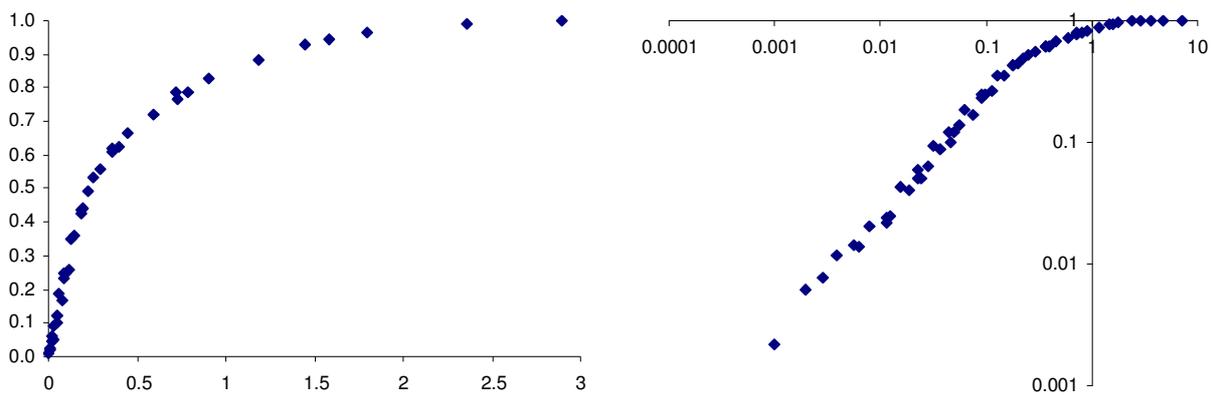


Figure 5 Gamma curve (a) linear axes

(b) logarithmic axes

The 'normal' exposure was arbitrarily chosen such that 100% exposure caused about 85% signal level. In the log plot, it is clear that the logarithmic nature of the curve continues at very low exposure levels. Also there is considerable compression of highlights above about 0.25 exposure, but the change of curvature is gentle. This is typical of proprietary log curves used in stills cameras, and in cameras intended to give a filmic look to video shooting.

5 Noise measurements

As a by-product of the gamma curve measurement, it was possible (although tedious) to get many measurements of noise level, at many exposure levels.

Fig.6a shows the results for a constant signal level of about 50%. The horizontal axis is marked in dB of gain, rather than ISO settings, with 0dB taken to be ISO 800, thus this plot runs from ISO 100 to 6,400. Normally, the noise level could be expected to worsen by 3dB for each change of 6dB in gain (i.e. for each 'speed' setting). In this plot, the average slope is more like 2dB per 6dB of gain, indicating that there may be some form of noise reduction going on at high gain (high ISO) settings.

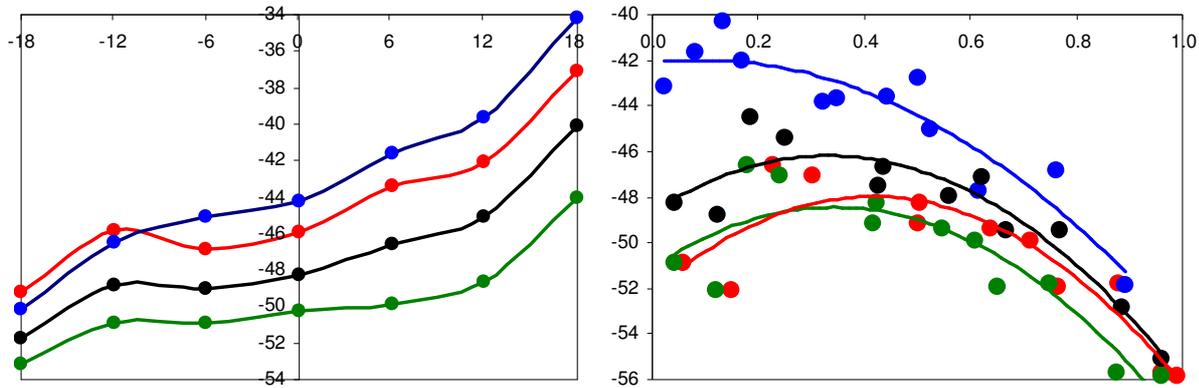


Figure 6 Noise 1080p (a) ISO 500

(b) at about 50% signal level

Fig.6b shows the distribution of noise with signal level at ISO 500. The points are not joined into lines because they come from many disparate measurements, and there is an uncertainty of about ± 1 dB in the measurements, caused by taking too few pixel values in the measurements (about 10,000 each time). The trend lines should normally follow the slope of the transfer characteristic, which means the noise should rise steadily as the signal level falls. Again, this plot shows the noise falling off towards black level, implying some form of noise reduction.

One possibility is that there is a non-linearity in the analogue circuits on the sensor, which would be likely to limit the frequency response as the gain rises near black (since the electronics has a fixed 'gain-bandwidth product'). This seems to be a standard trick in some camera designs, which has the neat effect of making the pictures look a little less noisy than they would normally be, since we tend to judge camera noise by the levels near black.

Nevertheless, noise levels are commendably low, considering the size of the photo-sites on the sensor.

Since the photo-site spacing is approximately $35.9/7,360=4.88\mu\text{m}$, the maximum area for each photo-site is about $23.8\mu\text{m}^2$, the camera should be about as sensitive, and have similar noise levels, as a conventional 3-sensor $\frac{2}{3}$ " camera, which is normally expected to deliver better than 50dB PSNR at its zero-gain setting. Such cameras are normally rated at between ISO 640 and 800. The results appear to confirm this supposition, approximately.

6 Exposure range

This is normally calculated as the ratio of the exposure which just causes white clipping to the exposure level below which no details can be seen. From the gamma-curve plot, the over-exposure range is about 3 times, 1.58 stops, and the gamma curve appears not to break off into a linear portion near black, it more closely resemble a film-type logarithmic curve. Therefore, the only limit on low exposure is the noise level. To confirm this, I took multiple exposures at ISO 500 and 1/100 shutter, and looked at the noise levels.

Exposures from F/4 to F/22 were all almost acceptable. At F/4, the white patch of the Colorchecker chart was clipped, but only just. At F/5.6, all patches were clear of clipping. Therefore the clipping level appears to be at about F/4.5. At F/22 and 1/200 shutter, the black patch hardly distinguishable from its surround (at about 0.6% signal level), and its noise level noise level is still about -50dB, so clearly noise is not the limit in this camera. Therefore, I raised the gain of this exposure by 12dB in a video editor, which increases the noise

level in the black patch to about -38dB, which is probably unacceptable, but the black patch is clearly distinguishable from its surround. Therefore, it seems possible to apply about +10dB of gain in post-production, 1.4 stops.

Now, the Colorchecker chart has a contrast range of 28.76:1 (white to black) which is about 4.85 stops. Add to this the $5\frac{2}{3}$ stops of lens range deemed acceptable in the tests, and the 1.4 stops which can be comfortably stretched in post-production, the total exposure range comes to about 11.9 stops, which is rather good.

7 Motion portrayal, rolling shutter

Since this is a CMOS camera, the image data is probably read from the sensor by scanning, rather than by taking a very brief, global, reading of the values into a store for later scanning. This process, called 'rolling shutter', is the same as scanning in a vacuum tube camera or CRT television set, and can cause severe geometrical distortions when there is significant motion in the image.

To test for this, a small fan was set up, rotating at a speed designed to cause strobing of the fan blades, 6 black sectors on a light disc. Fig.7 shows the results.

With the shutter set to 1/1,000, each frame shows a sharp image of the rotating blades, and it is clear that the distortion due to the scanning process is quite small. Blades on the right, travelling downwards, are widened by the scanning process, blades on the left move against the scanning and are narrowed. At the top and bottom, when motion is across the direction of scanning, the edges are bent to the left. This effect is independent of the shutter duration.

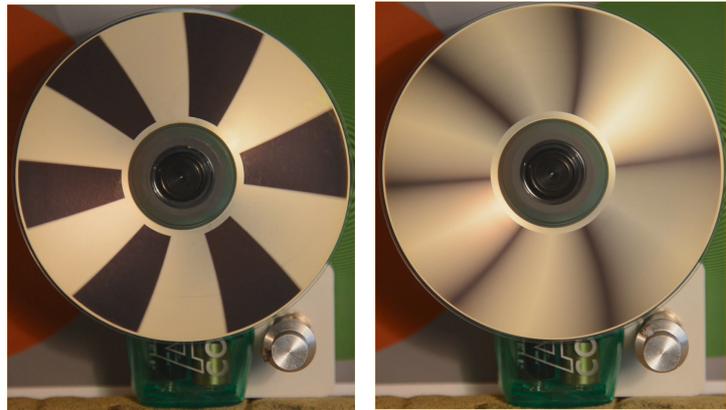


Figure 7 Rolling shutter (a) 1/1,000 (b) 1/100

8 Infra-red response (IR)

No video camera should respond to IR. The simple test for this is to take a remote control unit and point it into the lens, then start pressing buttons. If the LED is seen lit or flashing, then the camera must be responding to IR since all consumer remote controls use LEDs which emit 'light' at about 930nm, well beyond the red extreme of the visible range.

This camera showed no response to IR. Any camera which does show a response to IR will show odd colour behaviour under some lighting conditions, and will potentially have unstable black level, due to this light pollution.

9 Camera settings

The camera menus have few items which directly affect the video imaging performance. Only Sharpening was investigated, since all the other controls which modify the image appear to be special effects rather than the normal image controls found in conventional video cameras.

With Sharpening set to 0 (minimum value, Fig.8a), the image is decidedly soft but the coloured aliasing is still at high level. Setting Sharpening to 9 (maximum value, Fig.8b) greatly enhances contrast at low frequencies and emphasises the coloured spatial aliasing somewhat, both of which are unacceptable. Since there are no controls other than the sharpening level, nothing can be done other than to select a setting which is the least unacceptable.

Lower values are more promising. Fig.9 shows levels 2 and 4. Of these, the lower obviously emphasises the aliasing less and is more 'restful'. Level 4 should be regarded as the absolute maximum value for video shooting, with level 3 or 2 used for preference. However, none of these settings makes the pictures really suitable for serious HDTV programme-making, that is limited by the effect of the down-scaling.

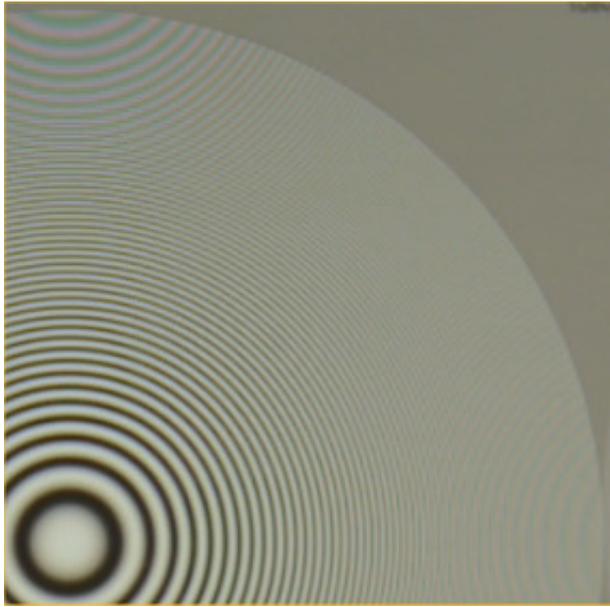
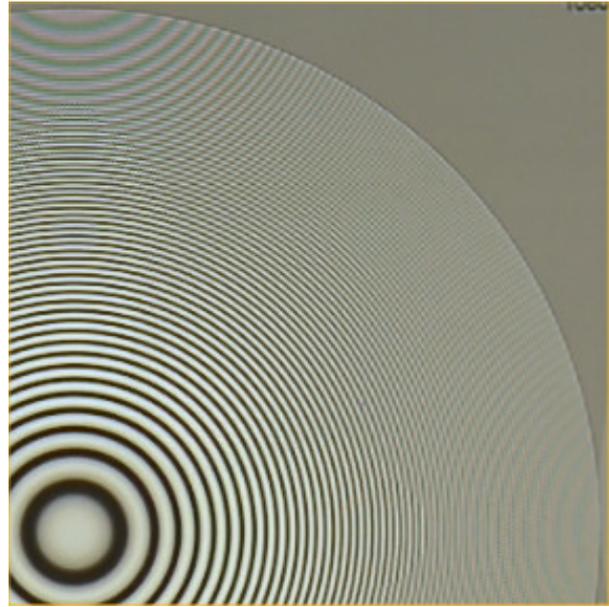


Figure 9 Sharpening (a) level 0 (minimum)



(b) level 9 (maximum)

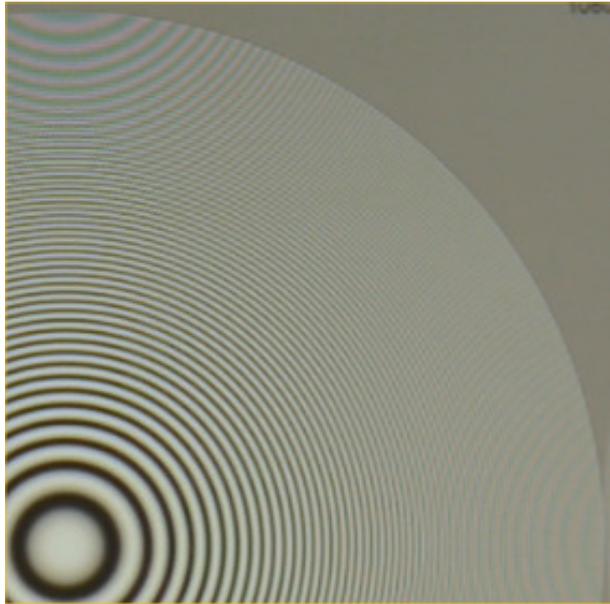


Figure 9 Sharpening (a) level 2



(b) level 4

10 Conclusion

The camera has very limited controls when in video mode, but has reasonable connectivity, allowing full-resolution external monitoring and recording. Sound facilities are sparse, the internal microphone is adequate for note-taking or guide sound, and the microphone connector is a 3.5mm jack offering only un-balanced input.

The test procedures were as described and recommended by the EBU, in Tech 3335 (<http://tech.ebu.ch/docs/tech/tech3335.pdf>). Video performance is not really acceptable at 1080p, much less so at 720p. Even though the sensor has 36.8 million photo-sites, it achieves only about 1,355x764, little better than 1,280x720.

Noise levels are very low, even with ISO settings up to 6,400. Significant coloured spatial aliasing is always present, and is clearly visible even on the camera's LCD display (921k pixels, about 1,176x784). In theory, meticulous control of the shooting style can reduce this but is unlikely to eliminate it, by using only

motivated pans together with fairly short depth of field, such that detail out of the focused plane is always soft and therefore can never provoke aliasing. Exposure range is, potentially, as high as 12 stops, although this will be limited by the acceptability of the noise levels near black.

Colour performance is good, and the camera does not respond to infra-red illumination.

Motion portrayal is good, the effects of the rolling shutter are nicely suppressed.

This camera cannot be recommended for serious programme-making.