



City Research Online

City St George's, University of London

Citation: Hull, C. (2008). Digital Imaging and Screening for Diabetic Retinopathy. Optometry Today, pp. 28-35.

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/5100/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Digital Imaging and Screening for Diabetic Retinopathy

Christopher C. Hull, PhD

The National Screening Committee has stated that “All people with diabetes aged 12 years and older should be offered screening for sight-threatening diabetic retinopathy using digital photography.”¹ Digital imaging offers several advantages such as archiving, ease of viewing evidence of progression, quality assurance, patient education and immediate indication of ungradable images. Knowledge of key aspects of digital imaging technology and performance therefore underpin screening for diabetic retinopathy in the UK. This article aims to discuss aspects of digital imaging related to diabetic screening as well as provide an indication of how computers can be used to automatically screen for sight threatening retinopathy.

We live in a digital world but to start at the beginning, what is “digital”? Analogue signals vary continuously in that they can take any value e.g. sound or light. Photographic film is analogue in that the darkening of silver halide crystals increases continuously with light exposure. Digital devices represent signals using discrete levels and this is necessary if we are to exploit digital computers and display devices. Following some historical context, the remainder of this article will consider the steps in the process of imaging a diabetic fundus in some detail before finally examining automatic screening methods for diabetic retinopathy using image processing.

The process of converting an analogue signal such as light variations in an image and converting it to a digital representation is known as sampling (fig. 1). Sampling has two important parameters: the spatial resolution and how many digital levels or “steps” are used to represent the light. The former is determined by the sensor resolution and the latter relates to the bit-depth of the image. Both will be considered in separate sections below.

In the Beginning

Digital imaging can be traced back to the 1930s and Farnsworth’s “image dissector” and Zoworokin’s “iconoscope.”² Both were imaging tubes; electrical devices that could record and transmit images without the need for photographic film. In the image dissector, the image was formed on a layer of caesium oxide and the electrons, which were emitted in proportion to the local light intensity striking the caesium oxide, were focused with electro-magnets on to an electron collecting plate. The electro-magnets allowed the image to be scanned without the need for mechanical parts. The image dissector was therefore a cathode ray tube (CRT) in reverse. In contrast, Zoworokin’s iconoscope was much closer to a modern digital imaging sensor employing a layer of photosensitive particles (detectors) that had a common base plate. This created an array of capacitors (devices that store electronic charge) where the charge held by each capacitor was proportional to the light intensity falling on a local area of the photosensitive particles. The amount of charge held by each capacitor was “read out” by scanning an electron beam across the base plate, which acts as one terminal of the capacitor. The iconoscope, unlike the image dissector, could build up charge as the light continuously fell on the photosensitive layer before the capacitors were discharged during the read out process. It was

therefore an image storage device and more sensitive than the image dissector. The iconoscope became part of television cameras in the United States until the late 1940s.

The evolution of digital imaging since then has been intrinsically coupled with that of the digital computer. The modern digital computer dates back to the 1940s when important parallel work was carried out in Germany, the United Kingdom and the United States although only the latter was widely reported because of secrecy. Since then developments and minaturisation have been rapid; it is likely that you have more computing power in your mobile phone than a large desktop computer from 25 years ago. The combination of the ability to record images digitally and modern computers has led to exciting developments in digital image processing, an aspect relevant to diabetic screening.

Steps in the Digital Imaging Process for Diabetic Screening

Figure 2 shows the steps in the imaging process. There is a lot of emphasis placed on image sensor resolution yet this misses several other potentially important parts of the process that could limit resolution and the ability to screen for significant pathology. The steps listed in figure 2 will be explored in the following sub-sections:

Optics of the Examined Eye

We assume that sufficient light has been directed on to the fundus and that it is reflected/scattered back out of the eye. The quality of optical systems is commonly measured using the modulation transfer function (MTF); the change in image contrast with increasing spatial frequency for a 100% contrast sinusoidal grating object. The smallest clinically visible microaneurysm is generally considered to be $30\mu\text{m}$ although they vary from $12\mu\text{m}$ to over $100\mu\text{m}$. Therefore the optics must give sufficient contrast at 5cycles/deg. if the microaneurysm is to be resolved assuming the second nodal point of the eye lies 17mm from the retina. Although measurements of MTF in normal eyes vary, there is clear indication that this requirement is easily met⁴. In addition, 5cycles/deg. corresponds to an acuity of 6/36, a level not generally encountered in normal corrected eyes. The exception would be eyes with abnormally high levels of aberration and/or media opacity. Consequently, the optics of the eye being examined will rarely limit the ability to screen for diabetic retinopathy unless media opacity or significant corneal irregularity is present.

Fundus Camera Optics

It is possible to make an estimate of the quality of the fundus camera optics to derive a ball park figure that can be compared with the resolution of the image sensor. If we assume an f-number for the fundus camera of f/5.7 then this would produce a diffraction-limited spatial frequency cut-off of 316 cycles/mm for a wavelength of 555nm. One of the latest digital single lens reflex cameras, the Nikon D3x, has a 24.3Mp (24.3×10^6 pixels) sensor with 6,048 pixels across the 35.9mm width of the sensor. This gives just over 168 cycles/mm resolution, a factor of 2 less than the assumed fundus camera optics. It is highly unlikely that the fundus camera optics will be a limiting factor and experience of many users and different fundus cameras has not indicated anything to the contrary.

Imaging Sensors

At the heart of a digital imaging system is an image sensor that can turn analogue variations in light level into digital information. Modern digital image sensors rely on developments in the electronics industry and in particular the ability to fabricate and connect many transistors on a single silicon chip called an integrated circuit. This technology, in its modern refined form, is used to build image sensors as well as other integrated electronic components, such as the microprocessor at the heart of a computer. However, the basic requirement for an image sensor is very simple; convert incoming photons into electric charge that can be determined at points across an image.

There are two main technologies for image sensors: charge coupled device (CCD) and complementary metal oxide semiconductor (CMOS). CCD image sensors were developed from computer shift register memories and the first solid state camera based on a CCD sensor was developed by AT&T Bell Labs in the USA in 1969⁵. A CCD image sensor consists of a regular array of photodetectors which are essentially capacitors that store the electrons generated when light falls on them. Each one of these detectors is referred to as a pixel, a contraction of picture element. The electronic charge, which is proportional to the light falling on the pixel is “read out” into additional circuits using a shift register. A shift register transfers the charge to the neighbouring pixel and eventually on to a ‘conveyor belt’ that takes the charge to the amplifying and sampling electronics. It is akin to passing buckets of water down a human chain. This process is known as charge coupling. The consequences of the design of CCDs, including the shift register read out are:

1. A high density of pixels leading to good image quality since minimal additional circuitry is needed at the site of the pixel.
2. The power consumption is relatively high.
3. Sequential read-out of each pixel – time taken.
4. Off-board electronics – not fully integrated.
5. Good uniformity since all signals go through the same electronics so more uniform.
6. Leakage of charge from one pixel to the next – ‘blooming’
7. Good dynamic range.

CMOS image sensors developed from the manufacture of logic and memory chips. It has only been since the 1990s that the techniques were sufficiently fine to allow production of image sensors that could challenge CCDs. In a CMOS sensor, each pixel has its own associated electronics, (amplifiers, digital converters) fabricated on the chip next to it. There is no charge coupling and no shift register; the pixels can all be read out in parallel. This architecture has the following advantages:

1. It is faster – parallel read out.
2. Less of the sensor area is devoted to light sensitive detectors – lower image quality.
3. Lower power consumption.
4. Compact “camera on chip” design suitable for applications such as camera phones, web cams etc..
5. No charge leakage particularly if pixel saturated.

For low to moderate cost applications there is little to separate the technologies. For

example, in recent years, technological developments have allowed better fill factors for CMOS (the area of the chip occupied by light sensitive elements). Similarly it has been possible to lower the power consumption of CCDs. More technical information can be found in the articles by Litwiller^{6,7}.

Colour Image Sensors

There are 3 main methods for generating colour images with digital image sensors: coloured filter arrays (CFAs), 3 separate image sensors for red, green and blue channels and temporal multiplexing using a spinning filter disc. Most low to moderate quality sensors use CFAs. A common pattern for the filters is known as a Bayer pattern where each group of 4 pixels has 2 green, 1 blue and 1 red sensitive pixel (fig. 3). This arrangement is in part because the human eye is more sensitive in the green.

A colour can be assigned to each pixel in a process known as demosaicing since it removes the mosaic pattern of the CFA. The need for demosaicing becomes obvious if we consider imaging a plain red background; If this produces a response from only the red pixels then there will be a large number of gaps in the image, an obviously undesirable feature. A colour is ascribed to a pixel using information from its neighbours (interpolation) or from assumptions about the colour variation for local groups of pixels in the image (spatial/spectral correlation). The algorithms for demosaicing are complex and go beyond the scope of this article. The aim is to reproduce the colour in the original image without introducing any colour artefacts. Fundus images contain a limited range of colours hence it is even less likely that colour artefacts will appear in the image that could potentially affect clinical decisions.

Image Sensor Resolution

Initial estimates of the pixel resolution needed to screen for diabetic retinopathy were based on the smallest clinically visible microaneurysm estimated to be 30 μ m. For this size feature to be unambiguously resolved, the image of the microaneurysm must fall on at least 2 pixels i.e. each pixel corresponds to 15 μ m on the retina. For a nominal 45° field and assuming that the nodal point lies 17mm from the retina, this corresponds to an arc length of 13.4mm. This must be imaged by $13.4 \div 0.015 = 890$ pixels of size 15 μ m. Cameras rarely use all of the sensor so the original resolution requirement was stated as 1365 x 1000 pixels by the National Screening Committee. This is a 1.4Mp sensor; a very modest requirement these days and exceeded on most camera phones! The current resolution requirement is stated as 20 pixels/degree since it refers to a circular region common in fundus images. This is equivalent to the previously stated pixel count allowing for some redundancy.

Modern mid-range image sensors typically have pixel counts from 4 to 24Mp. Consequently, there is no longer any difficulty meeting the required image sensor resolution when screening for diabetic retinopathy.

Sampling and Bit-Depth

There are two major types of image files: vector graphics and bit-mapped. Vector graphics files are most commonly used in drawing/illustration packages and for CAD

(computer-aided design). They mathematically record the shapes in the image which allows more faithful rendering particularly when the image is zoomed. Most images of general objects taken with digital cameras will produce bit-mapped image files where each pixel represents light levels using fixed steps. The number of steps is known as the bit-depth of the image. Computers work in binary arithmetic. A common bit depth is 8 binary bits which corresponds to 256 discrete steps (00000000 binary = 0 decimal; 11111111 binary = 255 decimal). It is possible to reduce the bit-depth and fig. 4 illustrates this effect. In a colour image sensor, each pixel records a level for red, green or blue on a scale of 0 to 255. A colour is ascribed to each pixel by combining the 256 levels of red green and blue leading to the ability to have over 16million different colours ($256 \times 256 \times 256 = 16,777,216$). It is possible to get 12bit and 14bit sensors but 8-bit is perfectly adequate for diabetic screening.

Image File Size

Reducing file size is usually necessary for ease of storage and file handling. To give an example of the scale of the problem, most diabetic screening examinations produce 4 to 6 images. For an 8Mp sensor, that would require 144Mb of storage per patient if the files were stored as raw data. Optometrists must complete 500 screening episodes per year with a recommendation for 1500. The recommended number of screenings would produce 216Gb of data. This then needs to be scaled for larger eye departments along with archiving this amount of data and backing it up to a central server. Reducing file size is therefore necessary.

There are two major ways in which file size can be reduced: downsampling and compression. Downsampling reduces the pixel count in the image. It can be carried out within image processing software and a number of image handling packages. We have already shown that a sensor resolution of greater than 2Mp exceeds the screening requirements even allowing for some redundancy. An 8Mp sensor could be resampled to reduce the pixel count to 2Mp. There are various ways of achieving this: for nearest neighbour resampling, the new pixel grid is effectively overlaid on the original image and the value of the nearest pixel in the original image assigned to the new pixel. This is a fairly crude technique and it leads to jaggedness and loss of some of the original pixel values. A better method uses interpolation. Bilinear interpolation is an extension of linear interpolation to two directions and uses the surrounding pixels and a linear fit between their values. The preferred method of downsampling is bi-cubic interpolation which produces new pixel values by interpolating from a bi-cubic spline surface, (other mathematical polynomials can be used), a mathematically smooth surface that represents the change in pixel intensity values. Although this algorithm takes longer, bi-cubic sampling reduces image artefacts and produces a smoother image.

Alternatively, compression can be used to reduce the file size in either a lossless or a lossy way. Lossless algorithms look for redundancy in the images, information that can be coded in a more efficient way. For example, the black corners of a circular fundus image typically contain pixels of the same value and this can be coded perhaps by recording a sequence as 28 pixels of 0 intensity rather than 28 individual instances of zero. This simple approach is known as run length encoding. Details of the various lossless compression algorithms go beyond the scope of this article; the interested

reader is referred to reference 8 for technical details or the various sources on the web.

The important point to note is that lossless compression algorithms allow the original image data to be retrieved. The most common image file types that use lossless compression are TIFF (tagged image file format), BMP (bitmap) and GIF (graphics interchange format). The latter is likely to cause some confusion since the file sizes are often small and GIF is the default image format for use on the web. GIF only allows 256 colours and hence it is not strictly a lossless format since the original image colours can not be reconstructed from the file data. The original recommendation of the National Screening Committee was that TIFF files should be used for fundus images when screening for diabetic retinopathy but this has changed.

Lossy image file formats lose information and the original image data can not be recovered from the data stored. However, if this is achieved without loss of visibly significant information, the file sizes can be much smaller. The most common image file type that uses lossy compression is JPEG. JPEG stands for the Joint Photographic Experts Group, a group set up in 1986 that produced the first JPEG standard in 1992. The JPEG compression algorithm has several stages and again details go beyond the scope of this article. The interested reader is referred to reference 8. There are three practical points of note⁹:

1. JPEG was designed to make minimal visible difference to photographic images i.e. those that have much smoother gradations in colour and tone. It does not work well on images where there is sharp contrast.
2. Part of the JPEG algorithm works on blocks of usually 8 x 8 pixels. The edges of these blocks can become visible.
3. Successive decompression when the image is opened and subsequent compression when it is saved will cause a progressive reduction in image quality. This is because the decompressed image contains artefacts that are compressed differently to the original.

JPEG is very effective and the compression ratio can be altered so that you can control the trade-off between image quality and file size. The key factor is the application the images are used for. The current statement on image compression for diabetic screening is that there should be no loss of clinically significant information. However, it should be noted that many cameras already use compression and produce a compressed (typically JPEG) image for grading. In this case further compression should be avoided¹.

Compression settings of 12:1 on JPEG are considered acceptable when screening for diabetic retinopathy whereas 20:1 is almost certainly not. As a guide, images over 2.5Mb in size are unlikely to offer any advantage whereas those below 400kB contain insufficient information for screening¹ although it is not clear to the author where the evidence for this comes from. The file sizes can be reduced by downsampling, compression or a combination of the two. Some of the more recent work in this area has looked at just noticeable differences in compressed and uncompressed images once they have been “viewed” by a standard spatial observer (SSO). This is a computer model where the 2-dimensional contrast sensitivity function has been modelled for the SSO. The two images, once processed by the SSO, can then be subtracted to see the differences. Using this approach, Helen Li and colleagues at the University of Texas have found that compression ratios as high as 37:1 do not produce just noticeable differences. A new

JPEG standard, JPEG 2000, which uses wavelet compression to produce a smoother image, (it also has a lossless option), can achieve even greater compression ratios. It is clear that more research on downsampling and compression is needed and the National Screening Committee supports this view.

Display Requirements

At some point an image needs to be displayed. Traditionally this was by printing a negative on to photographic paper, projecting a transparency or taking a Polaroid. Digital images are initially displayed using display screens of which there are two main types: cathode ray tubes (CRTs) and flat panel displays. CRTs comprise a cathode producing electrons that are accelerated by an electric field before striking a phosphor. The phosphor emits light and the electron beam is raster scanned across the phosphor to create the image. In recent years flat panel displays have become increasingly common due to their size and lower power consumption. There are a number of different display technologies that can be used in flat panels but the most common for small to medium size is liquid crystal displays. These displays are still inferior to CRTs and have a lower contrast ratio (the blacks aren't as black as on a CRT), smaller range of colours, are less robust and have a limited viewing angle. All of these limitations are of little relevance to diabetic screening. The key feature of flat panel displays that is relevant is that they can not change their screen resolution to match the image being displayed. The screen resolution is fixed unlike for a CRT. As a result, if the image doesn't match the screen pixels then interpolation has to be used possibly resulting in image artefacts.

Consequently, the National Screening Committee has recommended that a minimum of a 17" display be used (preferably 19") with a resolution of at least 1600 x 1200. This matches the image resolution requirement of 20 pixels/degree provided no more than 33% of the screen vertically and 45% horizontally is occupied by menu bars and window frames within the software. Importantly, it is a requirement that flat panel displays are used at their native resolution with care taken over settings for brightness, contrast and colour to maximise the appearance of significant pathology.

Image Quality

The National Screening Committee has specified the requirements for good image quality¹: For the macula image, vessels must be visible across 90% of the image, the centre of the fovea must be less than 1DD from the centre of the image and vessels must be clearly visible within 1DD of the centre of the fovea. The specifications for a good quality disc centred image are the same but centred on the disc and with fine vessels clearly visible on the surface of the disc. Adequate images are those where the central feature, macula or disc, is at least 2DD from the edge of the image with vessels still visible within 1DD of the centre of the fovea (macula image) and fine vessels visible on the disc (disc centred image). Good quality images are shown in fig. 5. Ungradable images don't meet these requirements unless there is referable retinopathy anywhere else in the eye.

PaSA Approved Systems

Supply of camera systems and related screening management software is handled through the NHS and their Procurement and Supply Authority (PaSA). Cameras

available for use within the National Screening Programme in England have been carefully evaluated against fixed criteria¹⁰. A re-procurement was carried out in 2006 with the framework agreement running until April 2009. There are currently 7 systems listed with the commercial digital camera back listed in brackets:

- i. Nidek AFC-210 (Canon EOS 5D)
Auto-focus fundus camera with the Canon EOS 5D DSLR incorporating a 12.8Mp near 35mm format CMOS sensor. The camera has the ability to store lossless (RAW) and lossy (JPEG) files
- ii. Canon CR DGi (Canon EOS 30D)
Canon's non-mydiatic, 45degree field of view fundus camera, with an image size of 17mm at the sensor plane. The Canon EOS 30D contains a 8.2Mp sensor measuring 22.5 x 15mm and again the ability to store lossless (RAW) and lossy (JPEG) file formats.
- iii. Kowa non-my 7 (Nikon D100)
Non-mydiatic fundus camera with the ability to switch between 20 and 45 degree fields. The Nikon D100 has a 6Mp CCD sensor giving a maximum pixel size of 3008 x 2000 over a 23.7 x 15.5mm format. This camera is now discontinued
- iv. Topcon TRC NW6s (Nikon D80)
Digital fundus camera offering 30 and 45 degree fields and peripheral fixation points for flexibility. The Nikon D80 offers a 10.2Mp (3872 x 2592 pixels) CCD sensor measuring 23.6 x 15.8mm.
- v. Topcon 3D OCT-1000
This is Topcon's Fourier Domain OCT coupled with a non-mydiatic fundus camera offering a 45 degree field and a working distance of 40.7mm. Topcon do not state the sensor resolution in their literature although it is integrated from the appearance of the instrument
- vi. Topcon TRC NW8 (Nikon D80)
The NW8 is the successor to the NW6s and offers increased ease of focusing and 30/45 degree fields. It can take a number of digital camera backs. The performance of the Nikon D80, which was approved with the NW8 as part of the framework agreement has been stated above under the NW6s.
- vii. Zeiss Visucam NM Pro
Zeiss' non-mydiatic fundus camera with 45/30 degree fields, internal and external fixation and the ability to cope with 3.3mm pupils although only for a 30degree field (most other cameras can't go below 3.7mm). The integrated CCD sensor offers 5Mp resolution coupled with Zeiss' telecentric optics.

Automatic Detection of Diabetic Retinopathy

There are approximately 2.5million diabetics in the UK all at risk of developing sight threatening retinopathy and requiring screening on an annual basis if over the age of 12. The potential benefits of automatic screening systems that can remove the clear "no referral" patients would be huge in terms of time and cost. Efforts to use digital image analysis for detecting specific ophthalmic signs in diabetic eye disease have been reported for over 20 years initially using photographic transparencies that were scanned to provide the digital images. Over 10 years ago systems were reported that gave better than 80% sensitivity and specificity so it is somewhat surprising that automatic screening systems are not available that meet the National Screening Framework requirement of

80% sensitivity and 95% specificity, the so-called Exeter Standard set by a meeting of Diabetes UK in 1984. Research suggests that *no* screening method meets these levels and it has remained in place for guidance only. Studies looking at screening using mydriatic digital photography report sensitivity and specificity values in excess of 85% indicating that a required specificity of 95% could have been set too high. Consequently, any automated screening method needs to prove that it meets the Exeter Standard or comes as close as digital photography and manual grading.

One of the first points to address with diabetic screening is what is required of an automated screening system? Ideally it would replace the grader and perform at least as well. This means it must first distinguish retinopathy from no-retinopathy and then grade the retinopathy if present. Early workers realised that it would first be necessary to identify the major components of the fundus: the optic disc, macula and vessels¹¹. Images could then be analysed for the location of lesions to find, for example, perifoveal exudation as well as indications of other sight threatening conditions such as glaucoma and age-related macula degeneration. Sensitivity and specificity values were near 100% for the optic disc falling to 83% and 91% respectively for the blood vessels. Different techniques are required for the different features: the optic disc, which was detected with the highest sensitivity and specificity, has a rapid change in contrast and colour at its boundary and is of a constant size. A clear definition of the feature being located helps the image processing algorithms probably explaining why the sensitivity and specificity scores were so high for the optic disc. An artificial neural network was used to define the blood vessels. An artificial neural network can learn from experience or “inputs.” As such it develops its own rules for deciding whether a pixel in a fundus image is vessel or non-vessel. This approach is taken because it is clearly more difficult to write a definition that would allow vessels to be found unambiguously. The same workers then examined the ability to automatically detect haemorrhages and microaneurysms (treated as one group) and exudates¹². Sensitivity and specificity for exudate detection was 89% and 100% and 78% and 89% respectively for haemorrhages and microaneurysms. Again these results may be predicted from the clearer appearance of the exudates with their higher contrast and colour difference from the background. One criticism of these early studies is that they didn't always use a large number of images to test the system. For example, in the last study, only 14 images contained haemorrhages and microaneurysms. It is perhaps surprising that 10 years on from these early promising results that automatic screening at least for no-refer/refer has not been adopted. Some of the possible reasons are listed below:

1. The lack of a gold standard. It would be ideal if all workers (there are approximately 30 individuals/institutes throughout the world working in this area) had a standard set of images for testing algorithms. In addition, manual grading is a useful comparison but should not be considered a gold standard.
2. Any screening system must be able to cope with eyes showing signs of other pathology.
3. An acceptance that the Exeter standard sets the bar too high and that a useful contribution can be made to disease/no-disease grading with current methods. A recent study from the Aberdeen group¹³ reported 90.5% sensitivity and 67.4% specificity for automated disease/no-disease screening using 14,406 images. However, the automatic system did pick up 99.8% of technical failures (better than manual grading) and

97.9% of patients with referable observable maculopathy/retinopathy. It can be argued that this has been achieved by having a much lower specificity but the workload of graders has still been reduced.

New algorithms continue to be published for lesion detection^{14,15} since each forms an important part of the detection process. However, evaluation of the whole process including how to potentially combine the results of 4 images from 1 patient to improve sensitivity and specificity are fundamentally important¹⁶. The results of these studies still produce vigorous debate in the literature between the various groups and there is no clear convergence of opinion that suggests automatic screening will be adopted in the near future even though it offers clear advantages particularly as a pre-screener to remove the obvious no disease cases. Such a system would have a lower specificity than the Exeter standard but it would still reduce significantly the workload of diabetic screening and grading and this is largely the view taken by the Aberdeen group who are prominent in the UK in this area.

Conclusions

Digital imaging offers many advantages when screening for diabetic retinopathy and technology is more than adequate for the task. Current issues include the amount of image compression that can be applied without affecting detection of clinically significant features and whether high resolution image sensors help compensate for image errors such as defocus or low light levels. There is still vigorous debate over the potential for automated screening for diabetic retinopathy although it could significantly cut down the workload of graders particularly for disease/no disease screening.

Figure Captions

[pic]

Fig. 1. Fundus image showing a line profile across the optic disc (indicated on the image by a black line) of pixel intensity (graph top right) and an expanded graph showing spatial and intensity sampling (bottom right)

[pic]

Fig. 2. Main steps in digital fundus imaging.

[pic]

Fig. 3. Bayer pattern used in a majority of colour filter array image sensors

[pic]

Fig. 4 (a-c) Fundus image showing the effect of bit-depth: (a) 8-bit (256 levels); (b) 4-bit (16 levels) and (c) 2-bit (2 levels)

[pic]Fig. 5. Examples of good quality macula (a) and disc (b) images (see text for details - from reference 1 with kind permission).

[pic]Fig. 6. Example of automatic detection of haemorrhages and microaneurysms from ref. 12: (a) original image; (b) contrast enhancement to improve detection; (c) classification using recursive region growing; (d) neural network identification of vessels; (e) subtraction of vessels from image; (f) overlay of potential haemorrhages and microaneurysms on original image.
(from fig. 6, reference 12 with kind permission of John Wiley and Sons)

References

1. UK National Screening Committee. *Essential Elements in Developing a Diabetic Retinopathy Screening Programme*. Workbook 4 (release 4.2 19th March 2008). (<http://www.retinalscreening.nhs.uk/pages/> accessed 29/4/09)
2. Davies A and Fenessy P. *Digital Imaging for Photographers*. Oxford, Focal Press, 1998. (3rd ed) p6.
3. Olk RJ and Lee CM. *Diabetic Retinopathy – Practical Management*. Philadelphia: JB Lippincott Co, 1993.
4. Atchison DA and Smith G. *Optics of the Human Eye*. Butterworth-Heinemann, Oxford (2000). Chapter 18.
5. http://en.wikipedia.org/wiki/Charge-coupled_device (accessed 29/4/09)
6. Litwiller D (2001). CCD vs. CMOS: Facts and Fiction. *Photonics Spectra*. January 2001.
7. Litwiller D (2005). CMOS vs. CCD: Maturing Technologies, Maturing Markets. *Photonics Spectra*. August 2005.
8. Gonzalez RC and Woods RE. *Digital Image Processing*. Prentice-Hall, 3rd ed. (2007).
9. <http://en.wikipedia.org/wiki/JPEG> (accessed 1/5/09)
10. <http://www.retinalscreening.nhs.uk/pages/> (accessed 1/5/09)
11. Synthanayothin C et al. (1999). Automated localisation of the optics disc, fovea and retinal blood vessels from digital colour fundus images. *Br. J Ophthalmol*, **83**:902-910
12. Synthanayothin C et al. (2002). Automatic detection of diabetic retinopathy on digital fundus images. *Diabet Med*. **19**:105-112
13. Philip S et al. (2007). The efficacy of “disease/no disease” grading for diabetic retinopathy in a systematic screening programme. *Br J Ophthalmol*. **91**:1512-1517
14. Quéllec G et al (2008). Optimal wavelet transform for the detection of microaneurysms in retinal photographs. *IEEE Trans Med Im* **27**(9):1230-1241.
15. Sánchez CI et al. (2008). A novel automatic image processing algorithm for the detection of hard exudates based on retinal image analysis. *Medical Engineering & Physics* **30**:350–357
16. Niemeijer M et al. (2009). Information Fusion for Diabetic Retinopathy CAD in Digital Color Fundus Photographs. *IEEE Trans Med Im* **28**(5):775-785