

It's all about the format – unleashing the power of RAW aerial photography

G. J. J. VERHOEVEN*

Department of Archaeology, Ghent University, Sint-Pietersnieuwstraat 35(110.016),
B-9000 Ghent, Belgium

(Received 28 March 2008; in final form 31 October 2008)

Current one-shot, handheld Digital Still Cameras (DSCs) generally offer different file formats to save the captured frames: Joint Photographic Experts Group (JPEG), RAW and/or Tag(ged) Image File Format (TIFF). Although the JPEG file format is the most commonly used file format worldwide, it is incapable of storing all original data, something that also occurs, to a certain extent, for large TIFF files. Therefore, most professional photographers prefer shooting RAW files, often described as the digital photography's equivalent of a film negative. As a RAW file contains the absolute maximum amount of information and original data generated by the sensor, it is the only scientifically justifiable file format. In addition, its tremendous flexibility in both processing and post-processing also makes it beneficial from a workflow and image quality point of view. On the other hand, large file sizes, the required software and proprietary file formats remain hurdles that are often too difficult to overcome for many photographers. Aerial photographers who shoot with handheld DSCs should be familiar with both RAW and other file formats, as their implications cannot be neglected. By outlining the complete process from photon capture to the generation of pixel values, additionally illustrated by real-world examples, the advantages and particularities of RAW aerial photography should become clear.

1. Introduction

Since the advent of the first truly affordable small format Digital Single-Lens Reflex (D-SLR) Canon EOS 300D/Digital Rebel in August 2003, an increasing number of (aerial) photographers worldwide have been converted to the digital approach of 35 mm photography, applying one-shot Digital Still Cameras (DSCs) in a rich variety of photographic solutions (Petrie 2003). While far from being a smooth transition for many people, it is safe to state that the majority of digital shooters (both small and medium format) already know more about enhancing a digital image in photo editing software than they have ever known about darkroom techniques. The direct approach (there is no 'preview' button in the darkroom), the ability to work in daylight with 'clean' computers instead of juggling with toxic products in darkrooms and the relative simplicity, as well as low cost, are only some of the advantages that digital image acquisition and (post-)processing enjoy. Nevertheless, both film and digital photography perfectionists are trying to accomplish the same thing: getting the maximum out of their initially acquired data. This article is about the latter: the

*Email: Geert.Verhoeven@UGent.be

originally captured or RAW information. In addition to being beneficial for image quality, it will be shown that RAW is the only format that (remote-sensing) scientists should use in their research, as it offers quantitative and qualitative possibilities that in-camera-generated Joint Photographic Experts Group (JPEG) and Tag(ged) Image File Format (TIFF) files do not.

2. RAW – a definition

In accordance with common terminology from the digital world, one would think RAW is an acronym. However, the word is an exception and signifies just what it sounds like: raw data. Although most texts describe RAW as the unprocessed data from the sensor of the DSC (e.g. Andrews *et al.* 2006), it is more accurate to consider a RAW file as the analogue sensor information that has been amplified and converted to digital data, without being subjected to any major processing by the camera's embedded software (firmware). Because this RAW file holds all data with only a minimal change compared to the data from the camera's digital sensor, RAW can be seen as the digital negative: it will never degrade, allowing an infinite number of digital prints (as JPEG or TIFF files) to be made in the future. It might even be better to consider it the equivalent of the latent image, as a RAW file holds all captured information without any digital development carried out afterwards.

To completely understand the nature of digital RAW capture and master its full (remote-sensing) potential, it is best to delve deeper into the process of actual image capture that completely takes place inside the one-shot DSC, which is defined as a photo camera equipped with both a digital image sensor for capturing full photographic data in one exposure, as well as a storage device for digitally saving the obtained image signals (Toyoda 2006). In 2009, all small- and medium-format D-SLRs offer the possibility to shoot RAW, while even most hybrid and some compact models can (the latter two being known as the consumer DSCs).

3. RAW – the creation

3.1 Photodiodes

Whether a digital sensor is one of the Charge-Coupled Device (CCD), Complementary Metal Oxide Semiconductor (CMOS), or Junction Field Effect Transistor (JFET) types, all image sensors of today's one-shot DSCs are silicon chips containing a two-dimensional array of photosites in order to produce the final image (figure 1). Each of those photosites contains a light-sensitive area made of silicon, a photosensitive detector or *photodiode* (Theuwissen 1995, Holst 1996, Nakamura 2006, Yamada 2006). When the number of effective pixels in a DSC is, for example, said to be 2560×1920 , the camera's sensor has at least 2560 by 1920 photodiodes. As, typically, one photosensitive element of the array contributes one pixel to the final image, the result is an image with 2560×1920 or 4.9×10^6 pixels (also denoted as MegaPixels or MP).

From the moment the exposure begins, these photodiodes will start to collect photons (figure 1) that are gathered by the lens. After the exposure, each diode contains a certain number of these photons, just as buckets would contain a certain quantity of raindrops after a rainstorm (Janesick and Blouke 1987). By collecting these photons, DSCs sample ElectroMagnetic (EM) radiation, both in a spatial (the location of the photodiodes) and tonal way (the amount of photons captured), as well as in time (the exposure time).

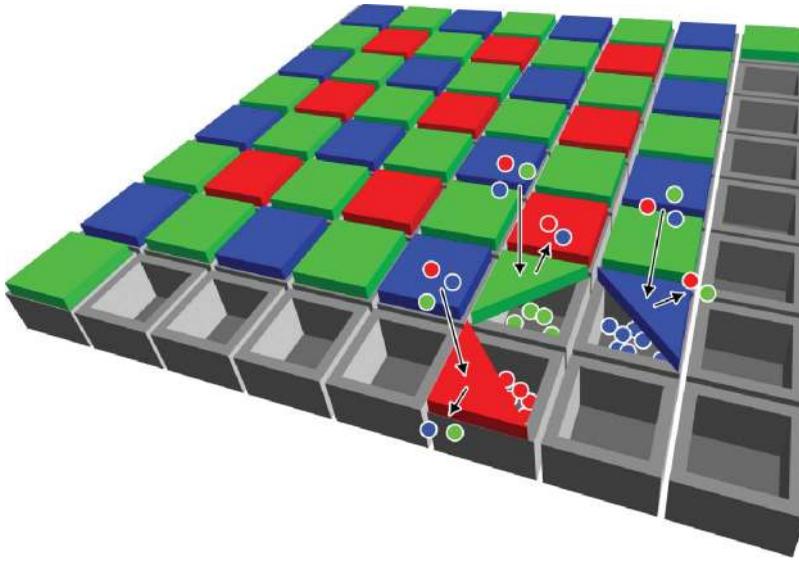


Figure 1. Layout and working of a photodiode array.

However, photodiodes simply count photons; they are monochrome devices, unable to tell the difference between different wavelengths. Consequently, such a construction would only be able to create greyscale photographs without further adaptations. Therefore, the most widespread method to give colour sensitivity to a one-shot DSC image sensor is the use of a *Colour Filter Array* (CFA). This mosaic pattern of coloured filters is positioned on top of the sensor, allowing only particular spectral components of the incident EM radiation to be collected (figure 1) (Holst 1996, Nakamura 2006).

Almost all DSCs use a three-colour Red–Green–Blue (RGB) pattern, in which the coloured filters are arranged as shown in figure 1. This arrangement, called a *Bayer pattern*, typically features a repeating group of four photodiodes, in which two have green filters – to mimic the higher sensitivity to green light of the Human Visual System (HVS) and enlarge the perceived sharpness of the digitally recorded scene (Hunt 1999, Parulski and Spaulding 2003) – while the remainder are either red or blue (Bayer 1976). Although this Bayer arrangement is almost constantly used in digital photography, other RGB patterns exist as well, as do designs with the three complementary colours cyan–magenta–yellow (CMY) (Holst 1996, Eastman Kodak Company 2001, Lukac 2009) along with four-colour systems such as the RGBE (E indicating Emerald) CFA introduced by Sony (figure 2).

A completely different approach was patented by Foveon Inc. With their innovative so-called X3[®] direct image sensor, this privately held corporation created in 2002 a particular kind of three-layered CMOS image sensor with a stack of three photodiodes at each photosite, enabling the capture of all three R, G and B wavebands at the same location by exploiting the wavelength-dependent absorption of silicon (Lyon and Hubel 2002, Foveon 2008). As this unique design is only implemented into Sigma's D-SLR SD9, SD10 and SD14, as well as into Polaroid's x530

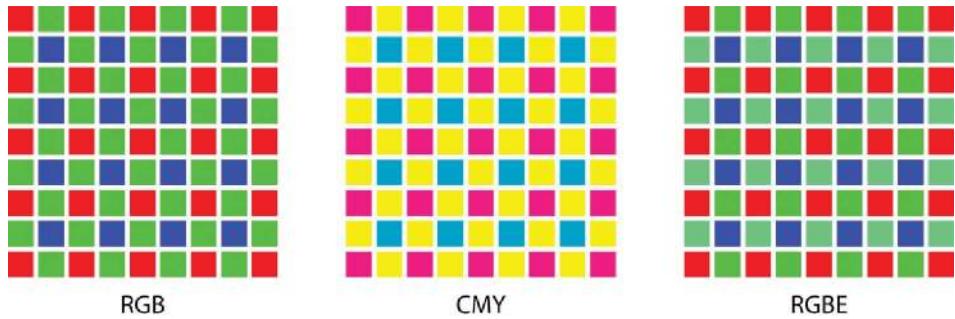


Figure 2. Different CFAs.

point-and-shoot camera (Foveon 2008), the remainder of this article will focus on the abundant Bayer CFA approach, unless otherwise indicated.

3.2 From analogue to digital

In spite of the different CFA designs, every single photodiode will capture only one spectral band (one colour component in the case of visible imaging), which is stored by a digital intensity value that is proportional to that particular incident EM radiation. As an example, consider a red-filtered photosite. Only the red part of the incoming light will pass through the filter, subsequently creating a charge in the silicon photodiode due to the photo-electric effect (Janesick 2001). This photo-electric effect, explained by Albert Einstein (1879–1955) in 1905, makes the silicon release electrons when exposed to EM radiation (Walker 2004), with the latter process obeying a linear relationship (Theuwissen 1995, Holst 1996, Eastman Kodak Company 1999, Janesick 2001, Yamada 2006). Even though only a fractional number of incident photons, denoted by the term Quantum Efficiency (QE), will effectively be converted by the photo-electric effect, more photons will always generate more free electrons. Referring back to the bucket analogy, it is this electrical charge that gets trapped and collected in a potential well, as long as the integration time lasts (Janesick and Blouke 1987). As soon as the shutter closes, this electrical charge is shifted to the output sense node (figure 3) and converted to a voltage (Holst 1996, Eastman Kodak Company 2001, Janesick 2001). Afterwards, these small analogue voltages are amplified by the read-out amplifier with a certain gain corresponding to the DSC's specific ISO value set (Koren 2001, White 2005). Promulgated as a standard by the International Organization for Standardization, this ISO value expresses the sensitivity of photographic film and sensors on a numerical scale. It goes without saying that a higher sensitivity setting (e.g. ISO 3200) needs more amplification than a lower value (e.g. ISO 200).

Finally, once the real-world signal is sampled by the diodes and captured in the form of voltages, it must be quantized to Digital/Data Numbers (DNs) or Analogue-to-Digital Units (ADUs) by the Analogue-to-Digital converter (A/D converter or ADC). The ADC therefore classifies the total possible range of continuously varying analogue voltages into a finite number of levels/gradations, subsequently assigning a DN to each level. The total range of different tones or quantization values an ADC can create is termed the

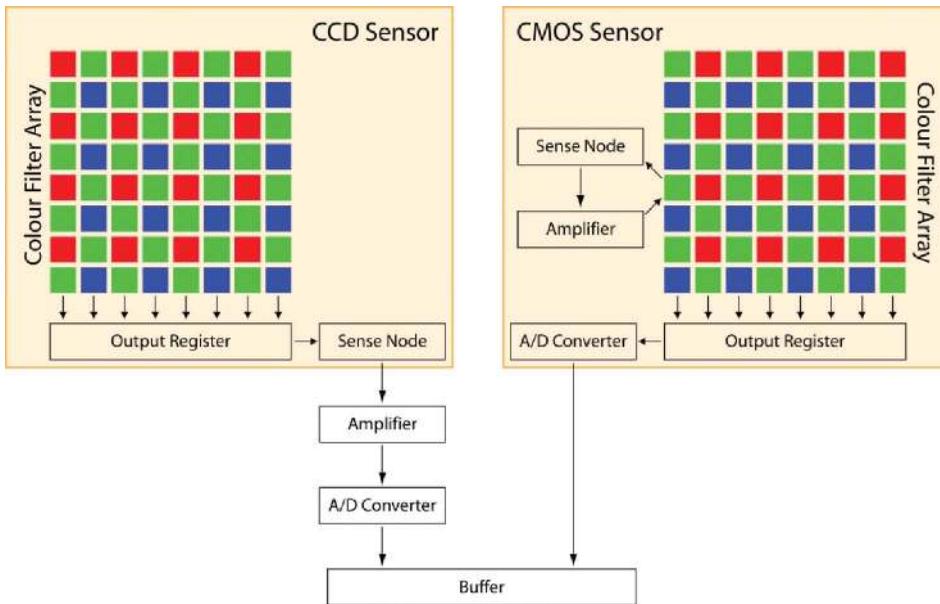


Figure 3. Creation of a RAW file (adapted from Bockeaert (2003–2009), with permission).

tonal range (Bockeaert 2003–2009), and is completely determined by its sample/bit depth: quantization with N bits rounds all possible voltage levels to these 2^N values (Gonzalez and Woods 2002). As each additional bit results in a doubling of the number magnitude, more bits used in quantization means that more shades can be encoded, which in turn leads to a smoother transition between each tone (figure 4). A typical consumer DSC therefore uses 8 bits / 1 byte, hereby allowing 2^8 or 256 distinct values. Most high-end D-SLR cameras use 12, 14 or 16 bit ADCs (Watanabe 2006, Adams and Hamilton 2009), yielding a wide tonal range of 2^{12} (i.e. 4096), 2^{14} (i.e. 16 384) or 2^{16} (i.e. 65 536) gradations, respectively. These high bit depths are important in avoiding *posterization*/banding, a phenomenon where abrupt changes between tones become apparent, often due to serious post-processing (e.g. histogram stretching) and is first discernable in regions with gradual tonal transitions such as skies and clouds (figure 4).

At this stage, the DN originating from a particular filtered photodiode still only refers to a greyscale radiation intensity value. Taking all diodes into account, a complete array of DNs is the generated outcome of each digital sensor, notwithstanding the sensor-related differences in this whole electron-to-DN chain (El Gamal and Eltoukhy 2005). Some DSCs also apply a pre-processing or camera compensation step on these DNs. Even though there is no real convention in the execution of this operation (neither in the steps executed nor the algorithms used), a few simple operations might be applied: defective pixel correction (to estimate the value of the defective diode), a linearization step to counteract any non-linearity introduced by the DSC's electronics, and some noise compensation (Ramanath *et al.* 2005, Adams and Hamilton 2009). Ultimately, this minimally processed array of DNs is sent to the camera's local buffer (figure 3), together with important information about the RAW file, the metadata.

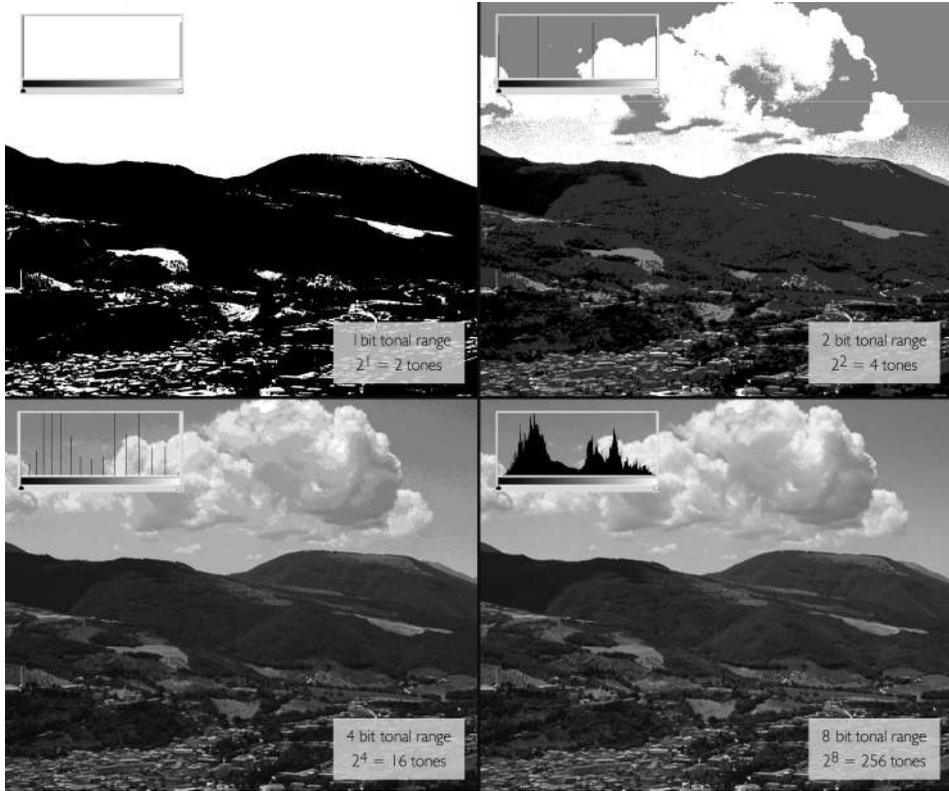


Figure 4. Four different tonal ranges and the occurrence of posterization.

3.3 Metadata

In addition to the DNs that encode the real-world scene, metadata are also generated. Literally meaning ‘data about data’, these metadata describe the content, quality, condition, owner rights and other characteristics of the data. In the world of digital photography, different standards are used to store information about digital frames, with the Exif (EXchangeable Image File) metadata standard probably being the most commonly known. Created by the Japan Electronic Industry Development Association (JEIDA), this Exif specification provides a rigid format to record shooting data (e.g. the serial number and model of the DSC, the aperture, shutter speed, focal length, possible flash compensation, colour space and date and time of shooting) in mandatory, recommended and optional tags stored in a separate segment of the file. If the camera is Global Positioning System (GPS)-enabled, tags can also hold the latitude, longitude and altitude of the geographical location that the particular photo was taken in. Moreover, also new vendor-defined metadata can be added (JEITA 2002, Parulski and Reisch 2009). These Exif-defined tags are created and stored simultaneously with the DNs, making it possible to analyse them afterwards. In addition, RAW files also hold some tags to define the CFA data (e.g. the pattern used), and additional image reconstruction parameters such as white balance, sharpening and noise settings (Parulski and Reisch 2009).

4. RAW – file details

Once it is created, such a RAW file can be seen as a container, holding two separate parts: the metadata stored in a separate header and a bunch of samples in the X and Y directions, with every sample characterized by one DN (except in Foveon's solution), and its location expressed in the image coordinate system. In this collection of DNs, each individual number represents one spectral value, generated by one photodiode. Consequently, this RAW image still has a greyscale character (Fraser 2005, Andrews *et al.* 2006) with embedded CFA pattern (apparent in figure 5). In addition to the aforementioned amplification and A/D conversion, no further adjustments are performed, making the amount of data processing very small. A RAW file thus truly is the most pure form of generated digital photographic data.

When RAW files of different manufacturers are compared, it becomes obvious that this file type lacks a general standard. Although this absence of a common structure is often considered to be a serious drawback, all RAW files store the original DNs and the metadata, as was outlined above.

5. RAW – processing

Initially, every DSC takes RAW pictures, but whether these are (1) directly saved for processing afterwards, or (2) instantaneously 'developed' by the camera to a JPEG or TIFF file depends on both the photographer and the DSC, as the latter needs to offer this option. When option two is chosen, the DSC's firmware will process the RAW data based on a mix of default parameters and certain user settings, such as sharpening, brightness, White Balance (WB), exposure adjustments, and so on. By contrast, the first choice allows the processing of the data at a later stage and gives the photographer almost total control over the further processing of the original DNs, only requiring a computer and suitable software, all in exchange for more visual

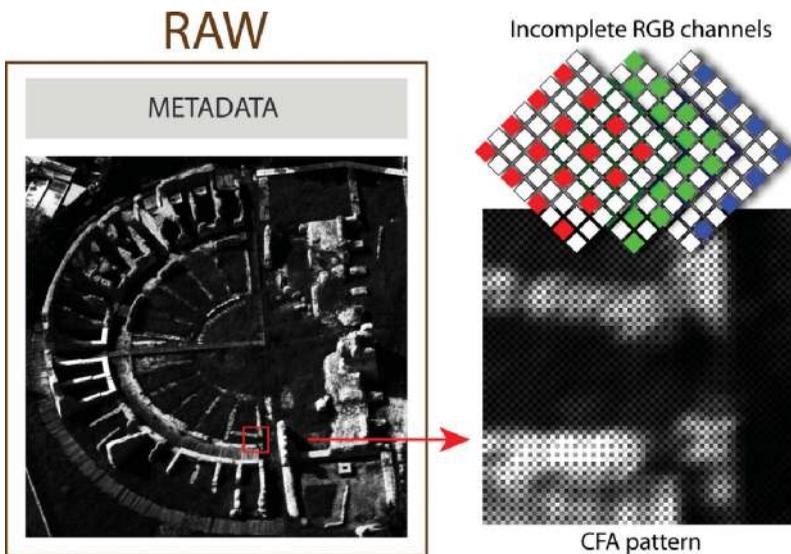


Figure 5. A RAW file.

quality (with the extra benefit of a data source that can be revisited and reprocessed endlessly without any quality loss) and the scientifically very important possibility of addressing the most pure data form generated. In comparison with the film-based approach, a JPEG or TIFF file that is written on the memory card is equivalent to the development and enlargement of the latent image by a photo laboratory inside the

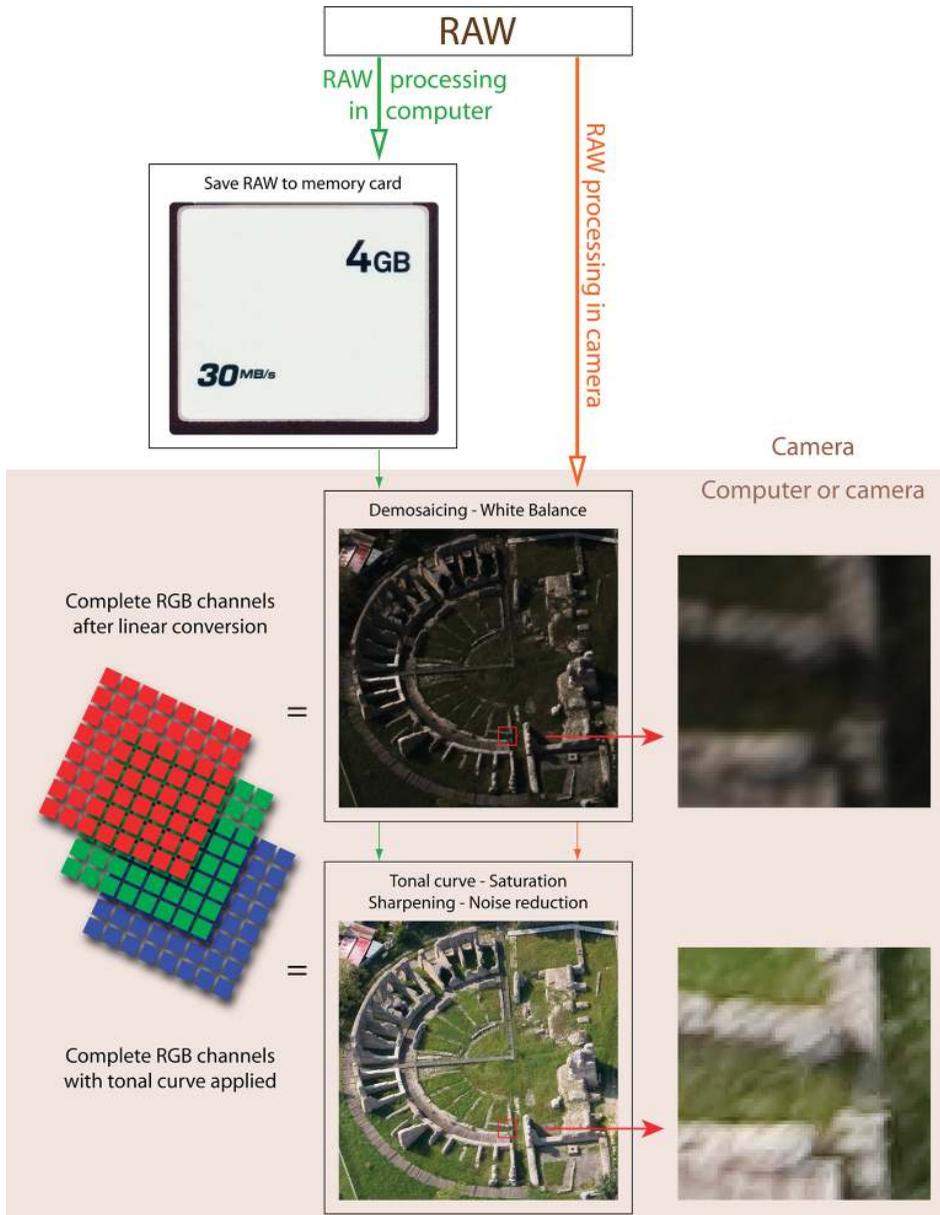


Figure 6. Main processing steps of a RAW file by the camera's firmware or on the computer.

camera. On the other hand, using the RAW file and subsequently processing it on the computer is equivalent to performing all darkroom work yourself (although now in a digital environment), with the additional benefit of reading out the initially captured values, which is of the utmost importance in scientific image processing and analysis.

By unfolding some of the individual processing steps the firmware performs, the opportunity is seized to compare this development procedure with the choices one has in a computer-based RAW conversion. Figure 6 depicts a flowchart of all individual development steps, hereby serving as a kind of visual guideline for the whole processing chain. Real-world remote-sensing examples will allow the important differences between both approaches to become clear.

5.1 White balance

Because the channel-specific DNs are generally unequal when photographing a spectrally flat object (white, black or grey), the values in each channel must be multiplied by a certain scaling factor to yield the expected identical channel numbers and tackle the unequal spectral responses (Giorgianni and Madden 1998, Stone 2003, Lam and Fung 2009). This is visually shown in figure 7. The RAW 12 bit values are yielded by photographing a grey, completely spectrally neutral WhiBal™ White Balance Reference Card (PictureFlow LLC 2007), which is displayed on the left-hand side. To generate a perfectly neutral grey card on screen, the red and blue channels are normalized here to the green channel.

However, it is also clear from the illustration that these multipliers change according to the EM source used to illuminate the object. A perfect white wall might reflect more blue wavelengths than red radiation when photographing on a cloudy, overcast day. On the other hand, several artificial light sources abundantly generate red wavelengths (Parulski and Spaulding 2003). This large variety of generated radiation attributes, to each EM source, a certain *Correlated Colour Temperature* (CCT): a number expressed on the Kelvin temperature scale, relating the specific spectral output of that EM source (which is perceived as a certain colour by the HVS) to the same colour perceived by heating a blackbody (i.e. an idealized dense object that absorbs all incident energy). The higher the temperature at which a blackbody is heated, the more its colour shifts to shorter wavelengths (from red to orange to bluish white), and the more intense is

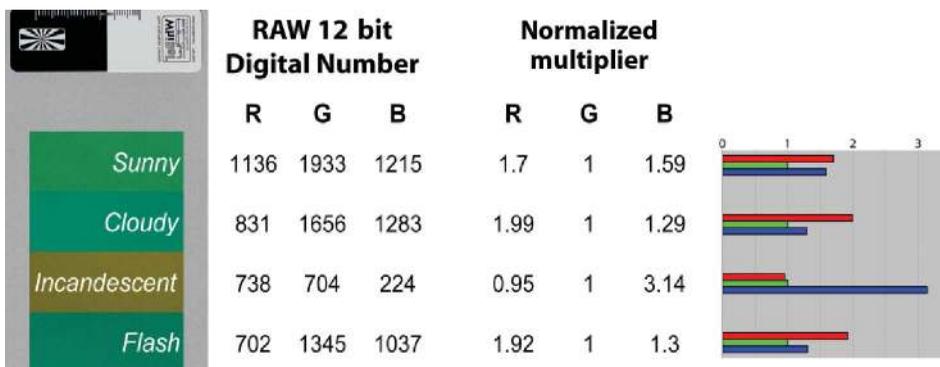


Figure 7. Channel specific DNs and calculated normalized multipliers retrieved from a WhiBal™ White Balance Reference Card photographed under different illumination conditions.

the emitted light (Walker 2004). A good example to illustrate this is heated iron. In a first stage, there will be a deep-red glow. By raising the temperature, the iron radiates brighter, reddish–orange light. Increasing the temperature even more yields a brilliant blue–white light. In other words, the emitted spectral radiation of a blackbody is only a function of its absolute surface temperature (as described by Planck’s law; Walker 2004), hence the term Colour Temperature (CT).

As a blackbody is an idealized object and most EM sources are far from ideal blackbody radiators – apart from the Sun (ca. 5800 K), halogen tungsten lamps (ca. 3200 K) and tungsten filament lamps (ca. 2850 K) – these sources cannot be described solely as a function of their temperature. This led to the concept of CCT: the blackbody temperature that yields the same chromacity experience as the EM source under consideration (Borbély *et al.* 2001, Fraser *et al.* 2004).

Human eyes constantly adjust to such CCT changes and will therefore be able to tell a wall is white, irrespective of the illumination conditions (Giorgianni and Madden 1998, Livingston 2002, Hung 2006). Digital sensors and film are unable to do so. In the analogue era, one had to change the type of film and/or use appropriate filters to avoid colour casts. In digital photography, the DSC only needs to know the wall is supposed to be white so it can accordingly calculate the correct multipliers (Lam and Fung 2009). This is also explained in figure 7. Within the WhiBalTM, four differently coloured patches are displayed, showing the WhiBalTM’s spectrally flat grey surface, but photographed under different illumination conditions (flash light, incandescent bulb, cloudy and open sky) and without any WB applied. By reading out the particular red, green and blue values of these reference pictures, the different channel multipliers were calculated by normalizing everything to the green channel. Both these patches and multipliers (which are also graphically displayed on the right) obviously demonstrate incandescent light to emit much more red wavelength than the other sources, indicated by the much lower red multiplier and the orange–yellowish colour cast of the patch. A cloudy sky, on the other hand, creates a bluish cast, indicated by the high DN in the blue channel (relative to the red channel) and the low blue multiplier.

At the time of capture, the WB setting is determined using the DSC’s automatic or manual WB setting and stored in the metadata. It has no effect on the generated DNs until the specific normalization values are effectively applied in the final calculation of the complete pixel values. In the case the RAW file is processed by the firmware, the multipliers are used to recalculate the initially generated DNs of all channels, hence making sure that the spectrally neutral zones – and by extension also all the other colours in the digital image – appear without serious colour casts, irrespective of illumination condition. In the case of the WhiBalTM, all three channels will ultimately have almost identical DNs for the spectrally flat surface (Adams and Hamilton 2009), yielding a picture of the grey card that looks very neutral to the HVS. Therefore, this card (and similar utilities) can be of great benefit to correctly calculate the WB, as even the automatic WB determination of professional DSCs can be fooled to a certain extent. In such cases of incorrect WB, the in-camera-processed JPEG or TIFF file will have a colour cast and although this can be mostly dealt with in post-processing, the quality of the picture will degrade to a certain extent and the cast is sometimes difficult to remove completely.

When storing RAW files, the WB can be set after the frame has been taken. Because this information is only stored as metadata, normal RAW conversion software reads this tag and applies it to the image when opening it. The user can always override this setting by applying other correction values, using a dedicated WB tool as the WhiBalTM

(and others) for accurate determination or arbitrarily choose values (in which case, the captured WB serves as a reference point to adjust the colours further). Independently of which approach is used, the WB is altered without altering the original RAW file or destroying any of the initially captured information. Therefore, RAW is an ideal solution when the photographer is not completely sure about the CCT of the light (as in aerial photography) and/or maximum possibilities in post-processing have to be maintained. Additionally, using software more suited for scientific purposes (e.g. MATLABTM by Mathworks, dcrwTM and IRISTM) the image can be processed without any WB being applied, hence yielding the originally captured DNs (which, in fact, is equivalent to multiplying every channel by a factor of 1.0). This is extremely important when the spectral response of a DSC needs to be determined (as in Moh *et al.* 2005, Verhoeven *et al.* 2009) or in case non-visual imaging is performed. Verhoeven *et al.* (2009) and Verhoeven (2007, 2008) describe the use of a modified D-SLR, capable of taking pure Near-InfraRed (NIR) aerial photographs. As there is no need to get true colour in this invisible range, white balancing can be omitted, generating files that clearly show the different spectral response of every channel. In figure 8(a), an NIR Nikon Electronic Format (NEF; Nikon's proprietary RAW format) aerial photograph was linearly processed (see §5.3) to a white balanced, 16 bit TIFF using dcrwTM, a free RAW decoder (Coffin 2008). Figure 8(b) shows the same image without any WB applied. The inset histogram clearly indicates the different responses of the three filter sets, showing that the red filters are most transparent to the NIR (hence the image's accordingly dominant colour). The different response in the red and blue channels of figures 8(a) and 8(b), is again shown in the lower part of figure 8. These results show that omitting WB can be very important in cases when only one or two spectral channels are needed or absolute spectral intensity measures have to be taken.

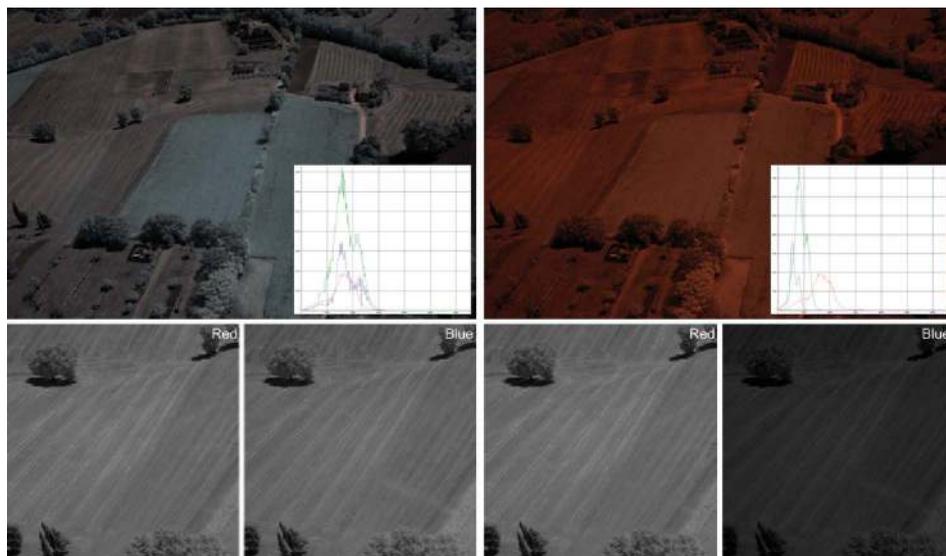


Figure 8. (a) NIR NEF file with custom channel multipliers and (b) the same image without any WB. The lower part shows the red and blue channels of (a) and (b), respectively.

5.2 Demosaicing

Apart from the unique, three-layered Foveon X3 sensor, most DSCs are single-shot models using one CCD, CMOS or JFET image sensor, where each photosite contains its own spectrally selective coloured filter. Because each photodiode senses only one spectral component, an algorithm is needed to fill in the corresponding DN for the other two bands (figure 9). To end up with a valuable three-channel image, the incomplete RGB values of the RAW file need to undergo a double operation:

- the greyscale DN (which corresponds to the intensity of a certain waveband) has to be converted to a matching colour value and
- the other two primary colours must be approximated to achieve a complete RGB image.

Both operations are done simultaneously in a process called *demosaic(k)ing*, CFA interpolation, colour reconstruction or de-Bayering (in cases where a Bayer array is used). To accomplish this process, a very important piece of information is read at the

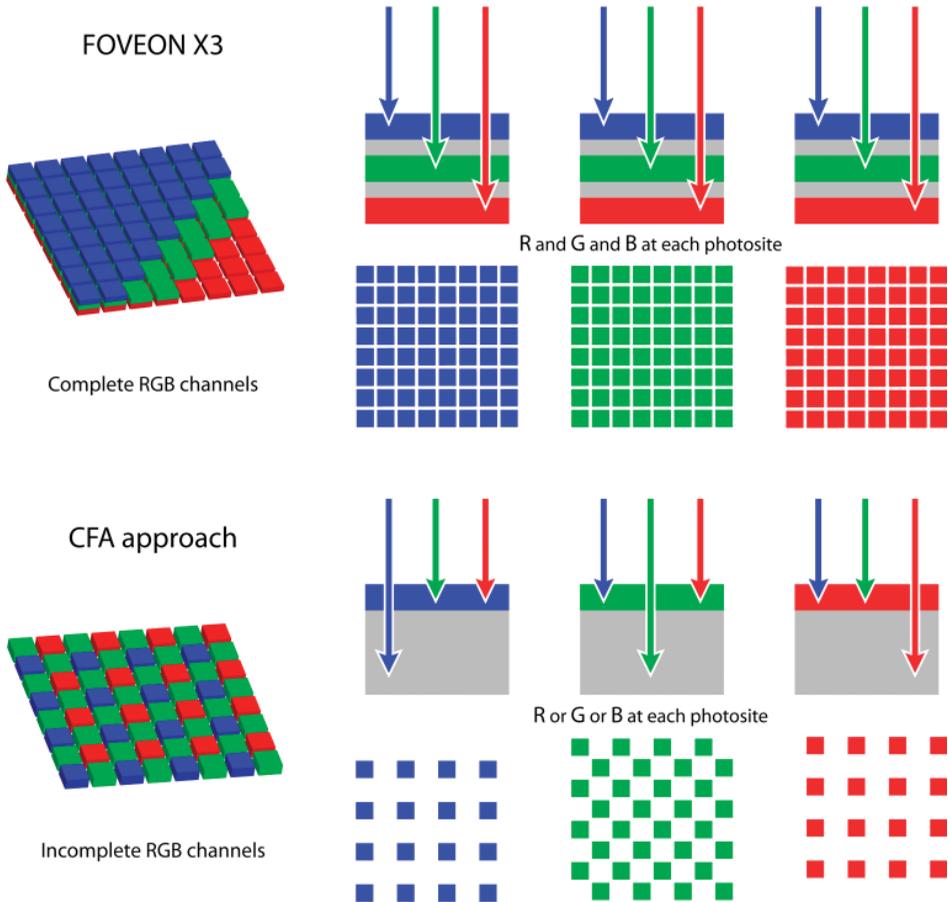


Figure 9. Difference in photosite-specific spectral information acquired by the Foveon sensor and the CFA solution (adapted from Foveon (2008), with permission).

Downloaded By: [Verhoeven, Geert Julien] At: 20:35 28 April 2010

beginning. This information, another piece of metadata included in the RAW file, is called the decoder ring (Fraser 2004). It is crucial in the processing of a RAW file since it stores the arrangement of the CFA, enabling the link between each intensity value and one of the three primary colours (Parulski and Reisch 2009). Once the software is aware of this, the missing spectral components can be filled in.

As a general guideline, a better reconstruction can be obtained if more actual spectral measurements are taken into account to estimate the specific missing channels (Sato 2006), although the situation is slightly more complex. To tackle the various artefacts that can be introduced during demosaicing, an immense range of algorithms (linear and non-linear), all varying in complexity and sometimes specified for particular CFA patterns, have already been proposed in recent years (e.g. Parulski 1985, Brainard and Sherman 1995, Adams *et al.* 1998, Gunturk *et al.* 2002, Lu and Tan 2003, Ramanath and Snyder 2003, Chang and Tan 2004, Lukac and Plataniotis 2004, Alleysson *et al.* 2005, Muresan and Parks 2005, Chung and Chan 2006, de Lavarène *et al.* 2007, Lian *et al.* 2007, Menon *et al.* 2007).

Most DSCs apply a Bayer CFA with a ratio of 2:1:1 among the green, red and blue filters (Bayer 1976). In addition to the abundance of information in the green channel, important correlations between red, green and blue DN_s exist (Kimmel 1999, Gunturk *et al.* 2002, Lu and Tan 2003, Wu and Zhang 2004, Li 2005). While old, well-known linear techniques such as nearest-neighbour, bilinear and bicubic interpolation do not use these characteristics (hence yielding rather bad to mediocre performances by blurring fine details and producing artefacts around edges), more sophisticated, adaptive techniques do exploit the diode's spatial and/or spectral correlations, yielding algorithms that completely interpolate the green channel before tackling the remaining red and blue spectral components, applying edge-directed interpolations as well as pattern recognition, pattern matching and even combinations of those techniques, all in complexity and computationally varying methods. However, less straightforward demosaicing requires more processing power to create a full colour image in an acceptable time span, power that often cannot be delivered by DSCs due to practical reasons: the processor must be small, light and may not ask too much of the batteries (Watanabe 2006). Therefore, the best interpolations can only be implemented in dedicated (proprietary) RAW converters run on modern computers, while DSCs (certainly compacts) apply quality-compromising algorithms.

To compare demosaicing quality, a minimally compressed, in-camera-generated JPEG file is compared to a post-processed JPEG file, generated out of the simultaneously stored NEF file (figure 10). By using Capture NX as a proprietary RAW converter, all parameters were kept equal, so as to only compare the difference between the demosaicing methods. It is obvious from enlarged portions of the in-camera-generated JPEG file that fewer details are present in both the structure of the leaf and the words 'DIGITAL' and 'Start'. Although this might not be a serious issue in the case of illustrative pictures, identifying small objects on aerial imagery certainly benefits from a RAW workflow.

In addition to the additional advantage of working with the uninterpolated data (which may be beneficial if only true spectral measurements are to be made), more technically oriented software packages might allow to choose/implement a specific demosaicing algorithm according to the data source. Consider aerial imaging in the invisible wavebands again. All the existing CFA interpolations are applicable in the visual domain, but modification might be expected to obtain a complete, reconstructed or demosaiced UltraViolet (UV) or NIR image with a minimum of various

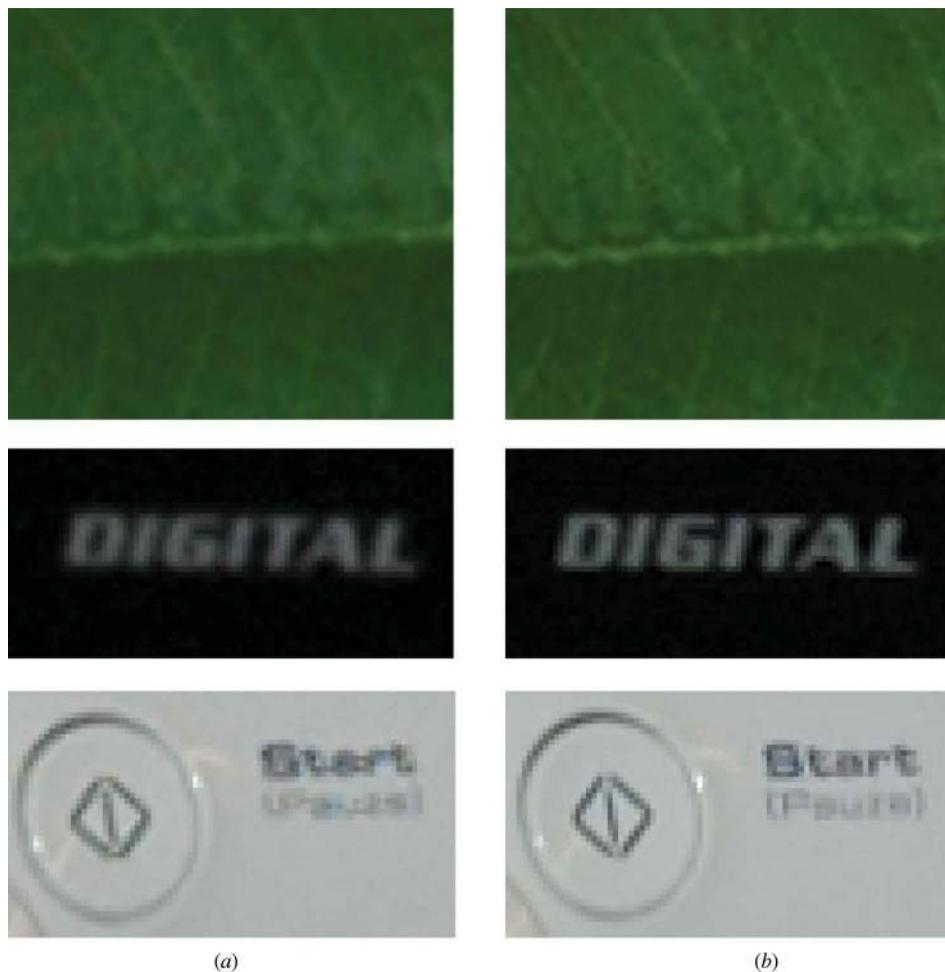


Figure 10. Comparison between demosaicing of (a) an in-camera-generated JPEG image and (b) a desktop-generated JPEG image of a Nikon D70s RAW frame.

artefacts. As an example, Miao *et al.* (2006) concluded that spectral correlations, based on colour ratio (Kimmel 1999) or colour difference (Lu and Tan 2003), are less in multi-spectral images compared to colour images. To solve this problem, they proposed the Binary Tree-based Edge-Sensing (BTES) approach, a generic demosaicing technique that establishes the interpolation order of different spectral bands, after which it determines the interpolation order of pixel locations within each spectral band, hereby using a binary tree-based scheme and edge-sensing (Miao *et al.* 2006). A quantitative comparison to assess reconstruction accuracy (using the Root Mean Square Error; RMSE) of mosaiced images created by their BTES method with the algorithm proposed by Lu and Tan (2003) showed the latter to perform better when only three channels (one visible and two IR) were taken into account. As a result, the Adaptive Homogeneity-Directed (AHD) demosaicing algorithm (Hirakawa and Parks 2003) was chosen to demosaic the author's NIR aerial imagery. Although the BTES method has not yet been directly compared to the AHD

algorithm, the latter algorithm yields fewer artefacts when compared to the Lu and Tan method, making it even more suited for NIR interpolation.

AHD demosaicing selects the direction of interpolation in order to minimize artefacts by applying a homogeneity map. Tested against even more recent methods, this non-linear iterative procedure still has to be considered an extremely well-performing algorithm, and is very good in reducing noise, hereby yielding sharp edges (Gunturk *et al.* 2005, Chung and Chan 2006, de Lavarène *et al.* 2007, Lian *et al.* 2007, Menon *et al.* 2007). A possible drawback is the unidirectional interpolation (only in horizontal or vertical directions), sometimes yielding artefacts in high-frequency components (Chang and Tan 2006).

Given this, the AHD algorithm was tested and proved to be the best in a range of possible algorithms for digital NIR imagery (figure 11). Using drcrawTM, demosaiced images were used in arithmetic channel operations in an attempt to calculate a Vegetation Index (VI) with the response of only one modified D-SLR (Verhoeven *et al.* 2009). Hence, a good calculation of the missing channel values was of the utmost importance and would be impossible without a RAW-based photography workflow.

5.3 Tonal curve

When it comes down to the response to EM radiation, there is a big difference between shooting with film and shooting digitally: film tries to mimic the light response of the HVS, while digital image sensors do not. The HVS, just like all sensations, functions in a non-linear way (Stevens 1961a,b). This can be illustrated by many examples: two light bulbs (stimulus) do not make the room seem twice as bright (sensation); moreover, it is easy to discriminate a 30 W light source from one of 60 W, but almost impossible to perceive the difference between a 500 W and 530 W lamp, although both differ by the

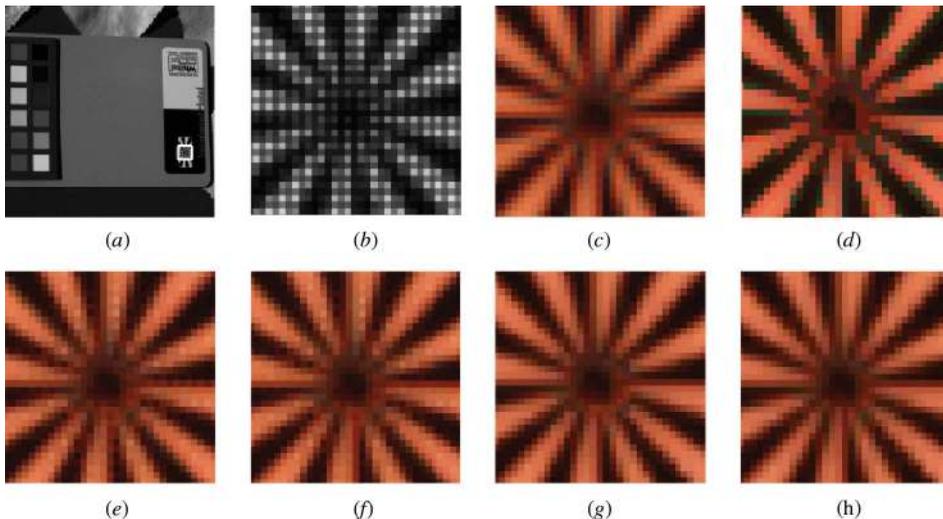


Figure 11. Effect of different demosaicing algorithms on an NIR image: (a) original NIR photograph, (b) portion of (a) showing the CFA pattern, (c) bilinear interpolation, (d) median interpolation, (e) gradient interpolation, (f) variable number of gradients algorithm, (g) patterned pixel grouping interpolation and (h) adaptive homogeneity-directed demosaicing.

same amount. Ernst Heinrich Weber (1795–1878) was the first person to try to describe the relationship between physical magnitudes of stimuli and their perception (Wolfe *et al.* 2006). Some 30 years after Weber, the German psychologist Gustav Theodor Fechner (1801–1887) elaborated on Weber’s law and came up with a logarithmic relationship to relate the magnitude of perceived sensation to the intensity of the stimulus (Fechner 1860, Stevens 1961a). His law, which is known as Fechner’s law (sometimes also denoted as Weber–Fechner’s law), states mathematically that $S = k \log(I)$, where S is the magnitude of the perceived psychological sensation, I is the physical intensity of the stimulus and k is a specific sensory constant previously defined by Weber (Fechner 1860). In the middle of the last century, Stanley Smith Stevens (1906–1973) showed Fechner’s law to be sometimes inaccurate. He therefore proposed a different relationship between sensory magnitude and stimulus magnitude: not a logarithmic one, but one that followed a power law. According to this power law, sensory or subjective magnitude (brightness in the case of light) grows in proportion to the physical intensity (luminance) of the stimulus raised to a power. Stevens therefore introduced a formula with one additional parameter (γ), $S = k I^\gamma$, where k is an arbitrary constant determining the scale units used and the power exponent gamma (γ) is a constant that is dependent on the sensory dimension (Stevens 1961a,b). For the sensation of brightness, the exponent in Stevens’ law (also called Stevens’ power law or the *psychophysical power law*) varies between 0.33 and 0.5 according to the conditions (Stevens 1961a,b), therefore resulting in a concave curve (figure 12). To double the

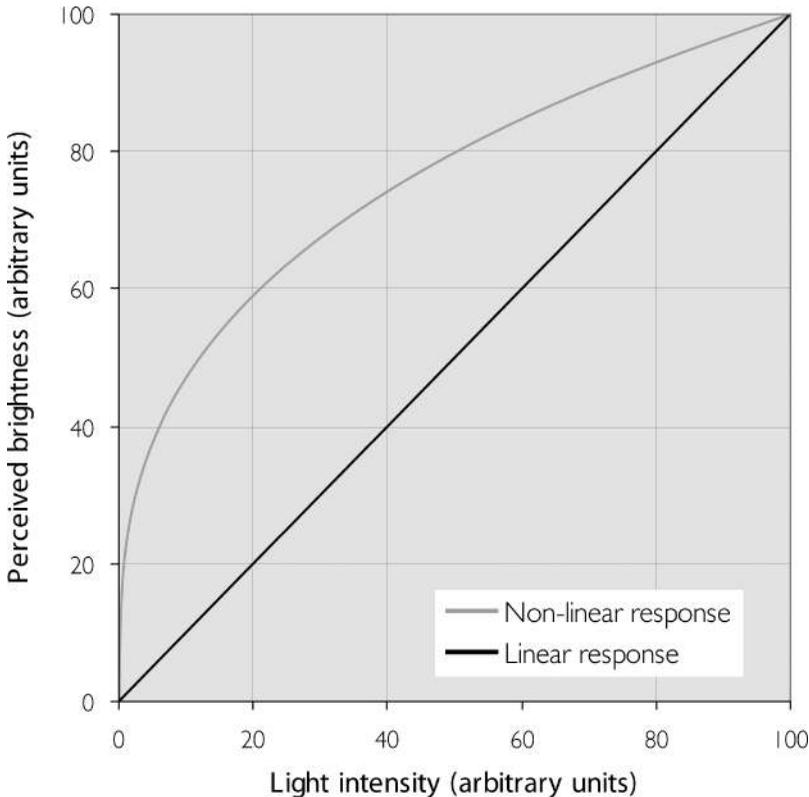


Figure 12. Stevens’ power law illustrated for brightness.

brightness sensation of a light source, a considerable amount of light intensity must be expended: more exactly eight times ($2 = I^{0.33} \rightarrow I = 2^3$). Due to this perceptual compression, the HVS perceives smaller steps in dark than in light, as the absolute stimulus increase in the latter must be a lot larger to perceive a brightness variation, compared to the stimulus intensity needed in dark areas.

However, a DSC's image sensor lacks this built-in compression. It just functions linearly: twice the amount of captured photons produces twice the sensor response. Digital sensors are therefore said to be linear recording systems, with a gamma equal to 1: output pixel value = k (input value) $^\gamma$, with $\gamma = 1$ (see figure 13).

However, this linear response yields two unpleasant effects:

- The darker areas (to which the HVS is most sensitive) are described by just a few tones (Koren 2001, Fraser 2005). As an example, consider the graph and table in figure 13. The x -axis of the graph indicates the whole light intensity range that a particular sensor is capable of displaying (i.e. its Dynamic Range; DR). This input is related to an 8 bit output by two curves. The linear curve indicates the way a digital image sensor maps the intensity range to 256 different tones. The brightest possible value the sensor is capable of will be captured by DN 255. If this maximal amount of light (100%) is halved (i.e. the quantity of photons divided by two), the remaining brightest value will correspond to a DN of 127 (due to the linear input–output relationship of a sensor). As halving the light quantity is known in photography as a photographic stop (f-stop), it is obvious that the brightest f-stop uses 128 or 50% of all available tonal values, while it only corresponds to the range where human vision is least sensitive. As additional halving yields 64 levels, 25% of all DNs are used to describe the consequent photographic stop (Koren 2001, Fraser 2005). Correspondingly, the third f-stop is represented by 12.5% of all levels (32 tones), etc. The higher the photographic

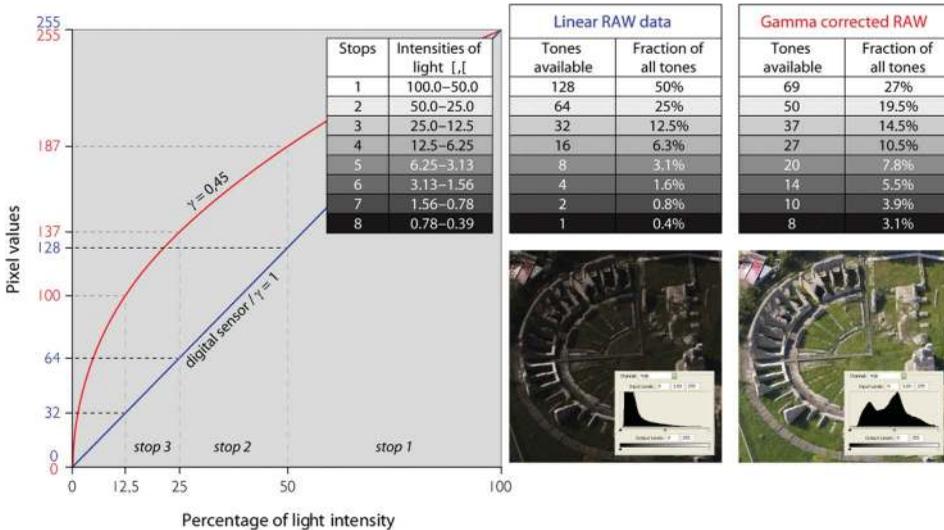


Figure 13. Gamma correction of RAW data.

stop, the smaller the amount of light. However, these smaller amounts of light indicate zones where the HVS becomes more sensitive; the more sensitive the human vision, the smaller the amount of available tones. In this example, the darkest shadows are captured in the eighth stop and represented by only one tonal value: black;

- The uncorrected output of a digital sensor looks very dark to the non-linear working human eye (Fraser 2005, Bockaert 2003–2009). Because humans perceive the first 12.5–20% of a complete intensity scale as middle grey (a value that depends upon the particular gamma value used), all remaining intensities (> 80%) are perceived as middle grey to white. In a linear converted RAW file, the lower part of this intensity range (< 20%) will be converted to very dark tones. Here, middle grey is only reached at an intensity value of 50% (figure 13). Due to the linear relationship, all pixels falling in this 0–50% intensity interval are now attributed with tonal values ranging from black to middle grey, whereas they should be perceived as black to quite bright. Consequently, a linear processed RAW image is perceived as very dim (see figure 8), which is additionally revealed by the histogram to be seriously skewed to the right (figure 13).

To solve these issues, RAW converters will apply a gamma curve (mostly $\gamma = 1/2.2 = 0.45$) to redistribute all tonal values and mimic the HVS (Koren 2001, Fraser 2005, Lukac 2009). This non-linear correction allocates more tones to the shadows and fewer to the brighter areas, yielding a more equal distribution of levels. In the end, only 69 different levels (or 27.1%) out of the same range of 256 discernable tonal values will be attributed to the first stop, while the shadow areas (stop eight) are now captured by eight levels (figure 13).

However, rather than a pure gamma curve, RAW converters and DSCs apply a *gradation / tonal / tone curve* to compensate for the non-linear human perception (Bockaert 2003–2009, Yoshida 2006). To a large extent, such a tonal curve equals a gamma curve. The difference between both can be seen in a log–log plot (figure 14); the gamma curve forms a straight line, while the tonal curve is more S-shaped, enlarging the overall contrast of the image (Adams and Hamilton 2009). As a matter of fact, this tonal curve strongly mimics the characteristic curve (D log E curve) known from film, hence making the digital image look as ‘normal’ as frames shot on conventional emulsions.

Although in-camera-processed images are subjected to a curve determined by the manufacturer (typically contrasty S-curves), setting the optional contrast parameters of the DSC (usually with the low, medium and high as options) allows for this tonal curve to be altered, generating a different output (e.g. reduction or enhancement of the midtones). This is where the advantage of RAW comes into play. RAW conversion within the DSC gives the photographer only a limited amount of control over the final tonal curve, whereas storing RAW files yields enormous possibilities concerning the redistribution of the available tones afterwards, as most RAW converters have several tonal curves embedded: one for emphasizing the shadows, another for better general contrast, etc. Additionally, each tonal curve can be further modified to a great extent by advanced settings such as contrast, brightness and shadows (the specific terms and settings all depending on the software used), thus creating the best tonal distribution for a particular image.

As RAW can be developed over and over again, one can even go back to the original file and use particular tonal curves of other RAW converters. The

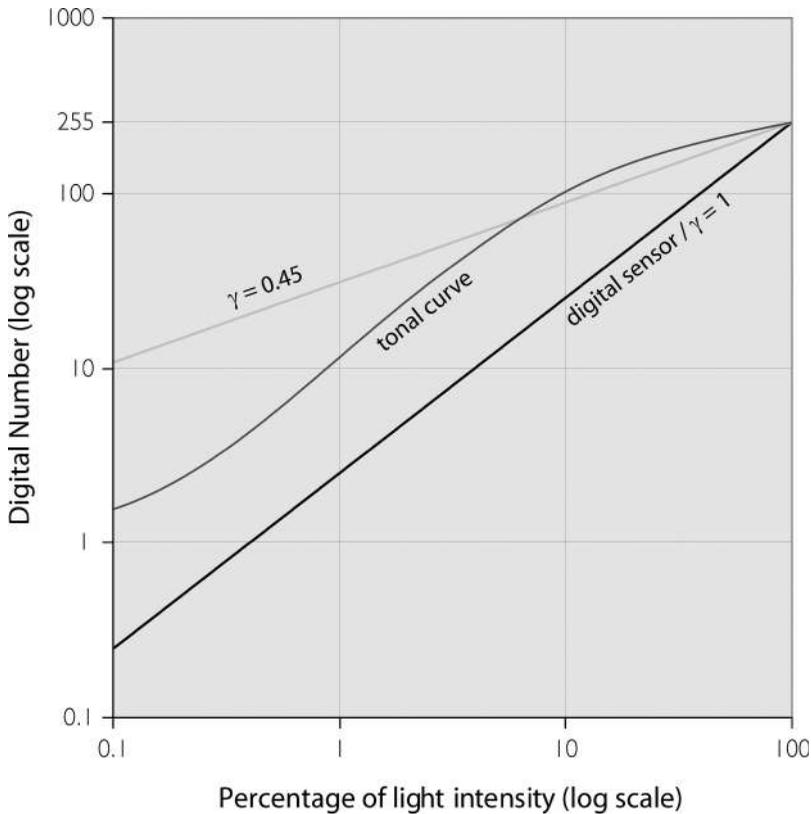


Figure 14. Comparison between normal gamma curve and generally applied tonal curve (adapted from Bockeaert (2003–2009), with permission).

possibilities are almost endless, which is not the case when shooting TIFF or JPEG files, as these formats have already had an in-camera tonal curve applied. Of course, an additional tonal redistribution can be executed in image-editing software, but will degenerate the final image quality and invite posterization, certainly in 8 bit images (Fraser 2005).

One major option offered by some software packages is the linear conversion (Hoffmann 2007), hereby displaying the dark image by omitting any tonal curve. Using this linear converted frame is believed by some photographers to yield even better results, since maximum flexibility is offered in retaining all possible highlight and shadow details (Wienke 2006). In addition to better image quality, a completely linear converted RAW is scientifically very important in research that needs the originally generated DNs: spectral characterization of the DSC or photographed features, mathematical operations between channels, spectroscopy and astronomical photometry (e.g. Howell 2006).

Figure 15 shows an aerial image of a faint negative crop mark, indicating a possible Roman road in central Adriatic Italy. Three versions of the same NEF file are displayed. All were converted to an 8 bit TIFF, but conversion settings differed on two points: the WB, which was only applied in figure 15(c), and the tonal remapping, which was suppressed in figure 15(a). From the larger absolute differences in DNs (the



Figure 15. Relative (%) and absolute (DN) reflectance increase in stressed vegetation canopy collected from a single, but differently processed RAW frame: (a) linear / no WB, (b) non-linear / no WB and (c) non-linear / WB.

latter acquired by taking the mean value of a zone in the healthy and one in the adjacent stressed canopy), it is obvious that most tonal curves seriously increase contrast in the middle input value range (DN divergence rises from 7 to > 17), while the relative reflectance difference decreases (from 17.5% to ca. 11%). Although the red and blue values are often scaled to the green channel, this example shows that white balancing might, on some occasions, also alter the values of the green channel to yield equal RGB values for neutral objects. However, scaling does not have an influence on the relative reflectance increase (the small difference between tile 15B and 15C is due to rounding errors in the 8 bit space). Consequently, a linear converted RAW file (WB being omitted or not) is the only image that can give accurate reflectance information (if the DN's are expressed in absolute units), whereas visual interpretation is generally best performed on imagery that took the WB, as well as the dissimilarity between the sensor's response and the HVS, into account (example 15C). Only in this way, does the yellow discolouration of stressed plants (*chlorosis*) due to the lost chlorophyll dominance over the carotenoids (Hendry *et al.* 1987, Adams *et al.* 1999) of the leaves becomes visibly apparent.

6. RAW – saving

Apart from these three major processing steps, capturing RAW allows us to assign a specific colour space, such as sRGB or Adobe RGB 1998, to the image (note that most of these colour spaces use a $\gamma = 1/2.2$, which explains why this particular gamma value was used in the previous step). In addition, there are particular possibilities to tackle noise (rather than some automatic noise abatement inside the DSC) and perform sharpening, which is a process that alters the DN's again, with the additional risk of generating image artefacts (jaggies, halos, etc.). But finally, both the in-camera and computer workflow end with the choice of file format in which to ultimately save the processed RAW file.

File formats are containers, created to hold digital data permanently and securely in files by offering a particular way of encoding the information. In order to define the

format used for a specific file, a filename extension is often used, being a text string of usually three or four letters that comes in the file name after the final period. As there are a lot of different data types in circulation, many file formats exist. In essence, three fundamentally different graphical data types occur: raster data (as photographs), vector or geometry data (e.g. computer aided design (CAD) drawings), and latent image data (such as the sensor's RAW information). As mentioned, the RAW information stores both incomplete intensity information, as well as metadata, that need to be processed into a raster image file format, holding complete colour or greyscale data. Both TIFF and JPEG files are raster file formats specifically designed to store still (or static) images as photographs.

6.1 JPEG

Rather than being a specific file format, JPEG is a large and complex compression standard defined in 1992 by the Joint Photographic Experts Group that was created in 1986 (Pennebaker and Mitchell 1993). Instead of encompassing one algorithm, JPEG can be seen as a toolbox. In this toolbox, several algorithms reside together with optional capabilities (Wallace 1991). When an application wants some of this standard to be implemented, a particular subset of the JPEG standard is selected to meet the best requirements.

As a matter of fact, the JPEG standard allows for two compression methods: lossy and lossless, meaning 'throwing away information' and 'storing all information', respectively (Gonzalez and Woods 2002). Lossless compression always yields smaller files than the original ones, but the compression gained by a lossy technique can be much higher. Because it typically can only compress full colour data by 2:1 (i.e. the resulting file is two times smaller), the lossless JPEG standard was never that popular. The lossy JPEG standard, on the other hand, is a completely different story. With a typical compression of 10:1, files can become very small when compressed with the baseline lossy JPEG algorithm, hereby almost visually indistinguishable from the input file. More compression can be achieved, but as more data need to be thrown away, the quality of the file will gradually deteriorate. Due to the possibility to trade-off quality against file size, this baseline JPEG subset (which also features a few optional extensions) became implemented worldwide in most applications dealing with photographs (Wallace 1991, Pennebaker and Mitchell 1993).

To store this JPEG-encoded stream of pixel data, together with the header containing all compression parameters (i.e. quantization tables and Huffman coding tables), a file format is needed. The JPEG File Interchange Format (JFIF) standard was therefore created by Eric Hamilton (Hamilton 1992, Pennebaker and Mitchell 1993). When one talks about a JPEG file (recognizable by one of the possible extensions *jpe*, *jpg*, *jpeg*, *jif*, *jfif* and *spf*) in reality a JFIF or Still Picture Interchange File Format (SPIFF) file is meant, with the latter being a more advanced substitute for JFIF, but without ever achieving significant adoption (Murray 1996, Parulski and Reisch 2009). However, this situation changed in 1996 when the JEIDA approved version 1 of the Exif standard, a file format that was specifically designed for storing image data originating from DSCs. In essence, this format stores fully processed images using the baseline subset of the JPEG compression standard, together with metadata embedded in various tags. On the memory card, these Exif-JPEG files are organized in particular directories and named according to rules defined by the Design rule for Camera File System (DCF). The Exif-JPEG image files themselves use TIFF tags to

store the required metadata (such as camera model, capture date and time) in one or more specific segments near the beginning of the image file, together with optional information on exposure values, lens settings and GPS coordinates. In addition, the first segment also contains a thumbnail image of 160 pixels by 120 pixels (Parulski and Reisch 2009). The structure of the file was cleverly designed so existing JPEG/JFIF file readers can process these new files without a problem. Consequently, virtually all consumer DSCs currently (i.e. end 2008) store JPG compressed images in this standard Exif-JPEG file format with the JPG extension (Parulski and Reisch 2009).

However, in addition to throwing away original information by the provided lossy compression, the Exif-JPEG file format is also a scientifically unjustifiable file format because its limited tonal range often has implications on the displayed Dynamic Range (DR). To explain this, the dissimilar, but often intermingled, concepts of DR and tonal range must be explained. If a DSC is capable of capturing a DR of 11 photographic stops, it means that the sensor can simultaneously discern intensities that are 2^{11} or 2048 times higher than the smallest amounts of light detectable by that sensor during a normal exposure (Pertierra and Reel Stream LLC 2005). Besides expressing DR on a logarithmic scale in f-stops, it is also possible to express DR linearly: if the amplitude of the smallest detectable signal is said to be 1, the sensor's DR can be expressed as the contrast ratio 2048:1. The bigger the difference between the faintest and most luminant objects a sensor is able to capture in one exposure, the larger its DR (Janesick and Blouke 1987, Yoshida 2006). A DSC with a high DR is preferable in order to capture weak light signals without washing out the highlights. Outside this DR, everything becomes absolutely white (in which case the sensor is said to be saturated) or hidden in noise. The tonal range, on the other hand, is nothing but the number of tones a digital image has to describe the whole DR. A useful analogy may be a staircase: while the height of the staircase is the DR, the number of staircase steps is the bit depth (Fraser 2005).

Consider again the example of the sensor with approximately 11 stops of the DR, in which case the brightest region of the digital image can be 2^{11} times more luminous than the darkest region. In order to fully render these approximately 2000 different intensities, the ADC needs to have at least a bit depth of 11, enabling the discrimination of 2048 (2^{11}) tones. However, if this 11 bit image would be remapped to an 8 bit output only 256 tonal levels would be possible in the final image. Although the DR could be almost unaltered, the image would suffer from serious posterization due to the loss of tones that could not be stored (Clark 2005a). In this example, all eight units of linear intensity (2000/256) will be grouped into one of these 256 categories, disregarding 87.5% (7/8) of all tonal values.

By depicting two sets of images with a different DR, figure 16 illustrates these concepts. In the upper images, the tones range from dark shadows on the forested slopes to detailed highlights in the clouds. Only a small portion of the clouds is completely white without any detail. In the second row of images, a lot of highlight detail is gone and large parts of the clouds are pure white. Both narrow and wide tonal ranges can, however, be found in small, as well as high, DR imagery. So, notwithstanding its intrinsic properties, using the aforementioned DSC to capture an 11 stop scene by saving a JPEG file will cause tonal values to be lost because a JPEG file can maximally store 8 bits/colour channel/pixel. The final DR will, however, depend on the workflow followed. Dealing with an in-camera-generated JPEG file, firmware generally sacrifices a portion of the DR originally present in the RAW data by using a highlight-clipping tonal curve, consequently attributing enough tones to the shadow areas where the HVS can discern most gradations. What remains is a smaller



Figure 16. Dynamic versus tonal range.

(e.g. 7–8.5 stops), but better rendered, DR, yielding a smoother image with good detail in the darker areas using the available tonal range. The largest DR most in-camera-produced JPEG files can attain is about nine stops for ISO 100 (even when applying user defined curves), a value that will decrease at higher ISO settings (Digital Photography Review 2000, Clark 2006, 2009).

When shooting in RAW mode, no tone curve is applied to the data, thus storing the sensor's full DR. Note that due to the linear input-output relationship of a sensor, the ADC's bit depth puts a theoretical upper limit to a RAW's maximum achievable DR. As an example, 12 bit encoding allows at the most 4096 tonal levels or 12 f-stops: 2048 values in the first stop, 1024 levels in the second stop, and only one tonal value to represent the twelfth stop. Because almost all current prosumer and professional DSCs utilize a 12- or 14-bit ADC, the shortage of possible tonal values as well as the omni-present noise limits their DR to a maximum of about 10 to 12 stops (Imaging Resource 2007, Clark 2009). As a result, a RAW workflow can extend the DR by at least one stop (Clark 2005b). In this respect, capturing RAW enables some 'exposure' flexibility, as shooting JPEG files would ask for two different exposures of the same scene (i.e. bracketing) to capture the RAW file's DR. This is visualized in figure 17, which displays two enlarged portions ((a) and (b)) of the same low-altitude oblique aerial photograph. Although a standard JPEG conversion (a) disregarded all highlight details, the latter were brought back (and selectively darkened to make them more pronounced) by processing the RAW file, using the same WB, demosaicing algorithm and JPEG settings as (a).

In addition to the possibility of recovering highlight detail that would be lost forever if the camera was set to JPEG mode, the wider tonal range that RAW possesses offers more freedom to play around with tones, and lowers the

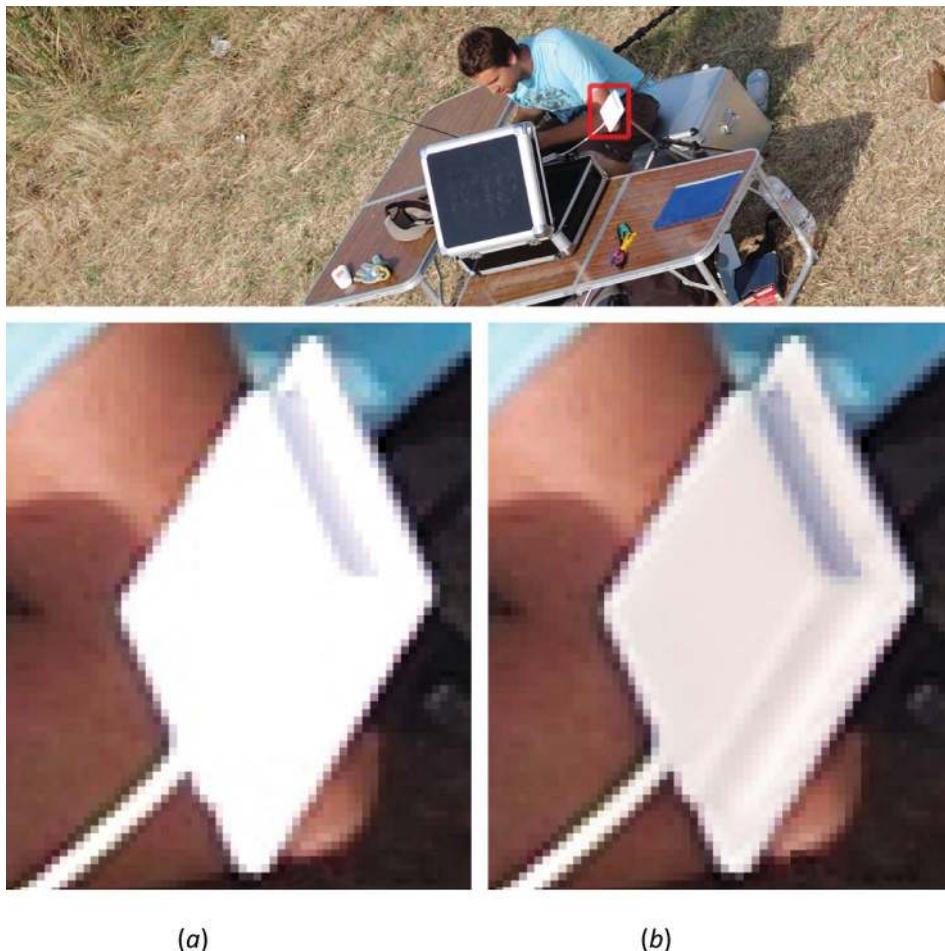


Figure 17. Comparison between highlight detail of (a) an in-camera generated JPEG image and (b) a JPEG image, resulting from a RAW workflow.

posterization risk in case of image manipulation (Steinmueller and Steinmueller 2004, Fraser 2005, Sheppard 2005). In both respects, JPEG files (certainly in-camera-generated ones) truly reduce the creative and scientific options. In addition to disregarding most of the acquired tonal values (which is scientifically indefensible), the saved image will always be worse than the original data (perceivable or not), making JPEG files totally unsuitable for aerial imaging. Moreover, resaving a JPEG file (e.g. after post-processing) generally introduces compression errors, meaning the file's quality is diminished every time it is opened and saved again (Hass 2007). Therefore, Exif-JPEG files should only be used at the end of the complete imaging workflow, serving purposes as such small preview files, e-mail attachments, imagery to be embedded in databases and/or in presentations. In the case where the DSC is only capable of initially storing JPEG files, one must make sure that the largest file size and the lowest compression possible are used, subsequently converting every JPEG file into a TIFF file from the moment all frames are downloaded from the memory card.

6.2 TIFF

First published by Aldus Corporation in 1986, but currently maintained by Adobe Systems Incorporated, this Image File Format describes and stores raster images using Tags (Adobe Systems Incorporated 1992). All the pixels that make up an image are stored in the body section of the TIFF file, while these tags (which would also be used in the Exif standard) hold information on width and depth of the image, acquisition date and time, copyright data, colour profiles, etc. TIFF was designed with the flexibility to define new tags in the future (Parulski and Reisch 2009), which has led, for instance, to the development of a tagset for carrying georeference information, enabling the TIFF to be located somewhere on the Earth's surface. This new standard was called GeoTIFF (Ritter and Ruth 2000).

In contrast to JPEG files, TIFF files can handle 16 bit/colour channel data and serve as a container for both uncompressed as well as compressed images. In the latter category, the lossless Lempel–Ziv–Welch (LZW; named after Abraham Lempel, Jacob Ziv and Terry Welch) is often offered (Adobe Systems Incorporated 1992, Gonzalez and Woods 2002), in addition to the lossy JPEG compression algorithm. The virtue of a TIFF file (recognizable by the extension TIFF or TIF) is its ability to store all data in the original order, hereby containing all captured colours and other pixel-related information (if no lossy JPEG compression is applied). Uncompressed TIFFs are insensitive for the aforementioned accumulative data loss. Additionally, its support for 48 bit imagery and the fact that it is portable (i.e. supported across different platforms such as Windows, Macintosh, and UNIX and has no favour for particular file systems) make it the world's number one preservation format for master copies, and hence the standard to save developed RAW frames (although the latter should never be thrown away). These characteristics, of course, have their drawbacks. Due to the fact that all possible data are stored, large file sizes occur (with a maximum of 4 GB). In addition, the flexible set of information fields or tags sometimes leads to problems concerning the correct opening or interpretation of the image file, because TIFF-enabled software packages do not always support the new tags added by the wide variety of (scientific) users. Although this issue created a new interpretation of the acronym 'Thousands of Incompatible File Formats', it is a problem that is not often encountered.

Finally, shooting in-camera-generated TIFFs is far from ideal, as most DSCs only output a 24 bit image (Steinmueller 2003). Even though most of the sensor's DR can be captured and the tonal range offers large editing headroom in case a 16 bit/colour channel TIFF can be saved, this file format presents, at the capturing state, no advantages to RAW, as its file size will be much larger and the interpretation of the scene is already performed (just as in the case of JPEG files), with no chance of addressing the originally captured DN's. Consequently, few DSCs provide TIFF files as an option.

7. RAW – workflow and software

Under most circumstances, the best processed and least compressed JPEG files are comparable to RAW images from a perceptual point of view. As long as the real-world scene being photographed is characterized by a limited DR, the exposure is spot-on and no (or little) editing is required, 8 bit JPEG files straight from the DSC can be just good enough (Koren 2001). This is possible in environments such as studios, where the light can be completely controlled, but the question remains whether one can be sure all these

demands are fulfilled when shooting aerial photographs with uncontrolled and ever changing lighting. Consequently, a RAW workflow always pays off, even when aerial photographs are only needed for visualization and interpretative purposes: after all, most RAW developing software enables a fast way to create TIFFs/JPEG files from the stored RAW files. All it takes is tweaking the WB and tonal distribution once, before a batch process can apply all the settings to the RAW frames, converting them to qualitative JPEG files with good demosaiced spectral channels. This approach is hugely superior to the JPEG workflow, even if it would be possible to adjust all camera settings on an image-by-image basis when flying.

An alternative might be the capture of a RAW + JPEG file combination, enabled by most DSCs: the JPEG image is there for immediate use, while the RAW can be processed at a later time and/or serve as back-up in case the JPEG file is corrupted (or vice versa). However, the storage space and processing power needed are the major drawbacks of this workflow, with the stored JPEG file quantitatively inferior to one generated by a batch process initiated a few minutes after the (aerial) shoot has ended. Some RAW formats (e.g. NEF by Nikon) already include a processed Exif-JPEG file that can be extracted by software which is not capable of processing RAW files, hence improving the workflow.

From a scientific point of view, it is impossible to disregard RAW since no other means are available to work with all the acquired spectral information in its most pure form. However, to create something meaningful out of these raw data, both scientists and photo artists need RAW conversion software to construct an image out of the recorded DN_s, or at least a programme that can decode the specific data format. Different software packages exist, all with their own specific advantages and drawbacks. Each package uses more or less proprietary techniques for demosaicing, noise reduction and tonal distribution, with the final aim of allowing for the best possible image to be created by the data available. Therefore, different converters produce different results.

In the early RAW days, most camera manufacturers provided some basic converter software with the purchase of the DSC. Nowadays, some of the best RAW converters are typically stand-alone programmes or plug-ins created by third-party companies. Due to the amount of software and the options available, a detailed comparison will not be given. The list provided in table 1 is just an enumeration of a few multi-camera

Table 1. RAW software.

RAW software	Manufacturer
ACDSee Pro	ACD Systems
Aperture	Apple
Bibble	Bibble Labs
BreezeBrowser Pro	BreezeSystems
Camera Raw	Adobe
Capture NX	Nikon
Capture One Pro	Phase One
DiMAGE Master	Konica-Minolta
Digital Photo Pro	Canon
Lightroom	Adobe
UFRaw	Open-source
dcraw	Open-source

capable converters, mostly optimized for a complete, colour managed RAW workflow, enabling the efficient processing of thousands of photographs. The two last pieces of software (UFRawTM and ddrawTM) are worth mentioning because they are free. The Unidentified Flying Raw (UFRawTM) package is Open-source software. It functions on its own or as a plug-in for the Open-source image-editing package GIMP (an acronym for the GNU Image Manipulation Program).

The second programme, ddrawTM, is an excellent free tool and the digital child of Dave Coffin who tried to 'write and maintain an ANSI C program that decodes any raw image from any digital camera on any computer running any operating system' (Coffin 2008). In its present form, ddrawTM supports over 300 different DSCs. As it is completely free and offers the possibility to process RAW frames without any WB and/or tonal curve and with a choice of demosaicing options, ddraw was implemented by many other RAW programmes (ACDSee ProTM, BreezeBrowserTM, UFRawTM and a lot of others), hereby making use of its excellence and offering the possibility to call ddrawTM from a graphical interface. In addition to ddrawTM, several functions also exist to read RAW imagery into the technical computing environment MATLAB (from Mathworks), after which an unlimited amount of Digital Image Processing (DIP) techniques are free to explore.

A last, often overlooked, advantage of RAW is the possible future improvement of RAW converters. As software mostly gets better with each new version, RAW converters can still be the subject of major development, delivering even improved imagery created by superior demosaicing and noise-suppression algorithms. Capturing RAW allows for the exploitation of these improvements, as the 'digital latent image' can be developed again in the future. In this respect, RAW presents the greatest long-term flexibility.

8. RAW – disadvantages

There are only three reasons for not using RAW: the DSC is not capable of storing RAW; it is a compact or hybrid model that freezes for some seconds after shooting and saving a RAW image; and the lack of storage space. Although it should not happen, it can occur. In those situations, there is only one rule: buy more memory, but capture the shot as a JPEG file, as it is far better than missing the photograph. In the other cases, RAW is the format to choose, even though a RAW workflow might saddle (aerial) photographers with some (minor) disadvantages. Generally speaking, the problems with RAW are threefold.

8.1 File size

As most RAW formats are not compressed, RAW files are much larger than similar JPEG files. A typical file from a 6 MP camera is about 7 MB to 9 MB (6 000 000 pixels x 12 bits/pixel = 8.6 MB), whereas a similar JPEG file with low compression is about 2 MB to 3 MB. Although the size of a JPEG file largely depends on the content of the image and the amount of compression applied, RAW files are often three to five times larger than the same image saved with JPEG compression. To tackle this issue, some DSCs use a compression algorithm to decrease their RAW size. To avoid mortgaging the quality however, these compressions are lossless, meaning no information is lost (except for some Kodak and Nikon models, which sometimes create slightly lossy, but visually lossless compressed RAW files). However, RAW compression algorithms can never achieve the level of compression that JPEG files can. Therefore, fewer RAW

images than JPEG images will always fit onto a particular memory card. Moreover, more Random Access Memory (RAM) and computing power are needed in comparison to working with JPEG images. Both issues should, however, largely be resolved by the fact that RAM and memory cards have significantly dropped in price during the last few years, making the file size less of a problem. However, the size of RAW files can still remain a concern in the cases of a DSC's maximum frame rate, burst rate and buffer capacity. The maximum frame rate of a camera quantifies the amount of consecutive images/frames it can maximally take in 1 s (e.g. 5 frames per second). Burst rate or burst frame rate equals the number of frames the DSC is able to shoot at the maximum frame rate. Because the DSC is not able to transfer all these images instantaneously to the memory card, they need to be temporarily stored in the buffer. If the latter is full, the DSC will stop shooting or reduce the frame rate significantly (White 2005). Bigger RAW files will limit this burst rate since they fill the camera's buffer faster and take longer to be written onto a memory card. Here, the differences between shooting JPEG files and capturing RAW files can be really significant. To give an example: the Canon EOS-1Ds Mark III shoots in continuous mode up to 12 RAW frames, while it enables the capture of 56 low-compressed JPEG images (Canon 2007).

Finally, a simple comparison with an 8 bit 6 MP TIFF file clearly shows RAW to be the winner concerning file size, as this TIFF will approximately be 18 MB (= 6 000 000 pixels \times 3 bytes/pixel).

8.2 Proprietary file format

As stated before, no two RAW formats from diverse companies are alike. Even though the structure of RAW files is, in most cases, a specific attribute of the TIFF 6.0 or TIFF / Electronic Photography (EP) standard, the formats are highly proprietary and unstandardized (Fraser 2005, Sheppard 2005). Due to constant improvements, even different RAW formats are created by the same DSC manufacturer. As a result, about 300 RAW formats are in existence as of 2008 (Coffin 2008), some of them with their own extension (table 2). These proprietary file formats bring along a major drawback: software that has the ability to handle these particular RAW files is needed. Instead of releasing one version every 2 or 3 years, all RAW software converters need to be updated every moment a manufacturer brings out a DSC that

Table 2. Extensions of some proprietary RAW formats.

Extension	Manufacturer
crw	Canon
cr2	Canon
dcr	Kodak
dng	Adobe
mrw	Minolta
nef	Nikon
orf	Olympus
pef	Pentax
raf	Fuji
srf	Sony
x3f	Sigma

creates a new RAW file format, just to offer support to a RAW range as wide as possible. Additionally, certain RAW files possibly will not be supported any more by RAW converters within 5 years, making them impossible to access.

Moreover, most big camera companies, such as Canon and Nikon, even try to hide or encrypt parts of their metadata to make it impossible (or at least harder) for others to decode the format. Nikon's D2X came, in 2005, with encrypted WB data to force Nikon consumers to only make use of Nikon software (which can of course decrypt the data). Just as was the case with the encrypted RAW files Sony used in 2003, it was only a matter of days before programmers cracked these files. However, big companies such as Adobe and Phase One do not dare to implement this crack into their software, because it could expose them to liability. As a reaction, the OpenRAW Working Group was founded, which proposed two possible solutions to this proprietary file format issue: the adoption of a universal RAW standard or the public documentation of RAW formats. In an attempt to contribute to the former solution (and of course further strengthen their role in the image business), the Digital NeGative (DNG) format was announced by Adobe in 2004. It was their attempt to standardize the RAW file format. To make sure everybody would use the new concept, Adobe launched a free, simple batch DNG converter for both Windows and Macintosh, capable of handling most current RAW formats. Given that companies are always doubtful regarding a standard created and owned by a single corporation, only four DSC manufacturers have implemented the DNG format as their native RAW format since 2004: Hasselblad, Leica, Ricoh and Samsung (Adobe Systems Incorporated 2008). Due to the resistance from many manufacturers even to incorporate this DNG specification into their RAW software, DNG often seems to be just another RAW format. The OpenRAW group therefore opts for the other solution, compelling the camera manufacturers to publicly document their RAW image formats (Specht 2006). By using only openly documented RAW formats, digital photographs created now and 5 years ago should also be readable over 30 years.

8.3 Time consuming

The flexibility RAW gives in processing each individual image can often lead to relatively long 'development' times. However, the time needed depends on both the final purpose of the file, as well as on the photographer. Artists trying to make the best image possible can literally devote hours to this RAW conversion, tweaking every possible setting in the software and combining layers of differently developed RAW files into one final piece of art. These photographers are the Ansel Adams of the digital age and their aim to create arty images compels complicated workflows.

However, RAW conversion does not need to be so time-consuming. In cases where a JPEG file is sufficient for viewing or printing purposes, only some initial settings need to be made, after which batch processing can be applied. Post-processing RAW images can even be a real time winner in certain cases. Imagine WB adjustments are required for hundreds of aerial images. Using a RAW workflow, it is much easier and more time-efficient to apply these corrections to RAW images than to any other file format, certainly if this can be integrated into a batch conversion.

9. Conclusion

With the advent of RAW digital photography, (aerial) photographers are now given the chance to get the most from their photography, as a RAW workflow enables

enormous control over the final output. Unprocessed files can always be accessed and developed, while converted files may be reprocessed if needed. Even when the high-end results and (scientific) flexibility offered by RAW are not needed, capturing RAW frames can, in fact, speed up the entire workflow if corrections need to be made. Moreover, batch processing the whole RAW series with the converter's default values generally delivers an output that will match at least the result of the in-camera-generated developed photograph. For strictly scientific photography, such as airborne remote sensing, shooting RAW is mandatory practice because it is the only way to assess the originally captured DNs, of the utmost importance in mathematical channel operations, intensity measurements, calibration and spectral characterizations.

A last statement often heard is that slide film, loved and (still) used by so many photographers, also yielded an image with a limited DR that was completely finished at the time the shutter button was pressed. Although that assertion stands up to scrutiny, it should not restrain photographers (and certainly not scientists) from exploiting new camera techniques and possibilities, neither because of these reminders of an analogue past, nor by the fact that some technical photographic knowledge is often a necessary prerequisite to take full advantage of new imaging technologies.

Acknowledgements

This paper arises from the author's PhD, which studied the application of remote sensing in archaeological surveys. As a PhD fellowship of the Research Foundation – Flanders (FWO), the doctoral study was conducted with the permission and financial support of the FWO and supervised by Professor Dr Frank Vermeulen (Department of Archaeology and Ancient History of Europe, Ghent University). Finally, Karen Ryckbosch and Ronny Verhoeven are acknowledged for proofreading the article and correcting the English where needed. All errors and misconceptions remain, of course, the author's own responsibility.

References

- ADAMS, J.E. and HAMILTON JR, J.F., 2009, Digital camera image processing chain design. In *Single-Sensor Digital Color Imaging Fundamentals*, R. Lukac (Ed.), pp. 67–103 Boca Raton, FL: CRC Press).
- ADAMS, J., PARULSKI, K. and SPAULDING, K., 1998, Color processing in digital cameras. *IEEE Micro*, **18**, pp. 20–30.
- ADAMS, M.L., PHILPOT, W.D. and NORVELL, W.A., 1999, Yellowness index: an application of spectral second derivatives to estimate chlorosis of leaves on stressed vegetation. *International Journal of Remote Sensing*, **20**, pp. 3663–3675.
- ADOBE SYSTEMS INCORPORATED, 1992, TIFF™, Revision 6.0. Available online at: partners.adobe.com/public/developer/en/tiff/TIFF6.pdf (accessed 30 September 2007).
- ADOBE SYSTEMS INCORPORATED, 2008, DNG hardware and software support. Available online at: www.adobe.com/products/dng/supporters.html (accessed 16 February 2008).
- ALLEYSSON, D., SÜSSTRUNK, S. and HÉRAULT, J., 2005, Linear demosaicing inspired by the human visual system. *IEEE Transaction on Image Processing*, **14**, pp. 439–449.
- ANDREWS, P., BUTLER, Y.J. and FARACE, J., 2006, *RAW Workflow from Capture to Archives: A Complete Digital Photographer's Guide to Raw Imaging* (Oxford, UK: Focal Press).
- BAYER, B.E., 1976, US Patent 3971065. *Color Imaging Array*, 20 July 1976. Available online at: www.freepatentsonline.com/3971065.pdf (accessed 12 March 2007).
- BOCKAERT, V., 2003–2009, *The 123 of digital imaging Interactive Learning Suite*. Available online at: www.123di.com

- BORBÉLY, Á., SÁMSON, Á. and SCHANDA, J., 2001, The concept of correlated colour temperature revisited. *Color Research and Application*, **26**, pp. 450–457.
- BRAINARD, D.H. and SHERMAN, D., 1995, Reconstructing images from trichromatic samples: from basic research to practical applications. In *Proceedings of 3rd IS&T/SID Color Imaging Conference: Color Science, Systems and Applications*, Scottsdale, USA, 1995, pp. 4–10.
- CANON, 2007, EOS-1Ds Mark III. Digital SLR. Available online at: www.usa.canon.com/consumer/controller?act=ModelInfoAct&fcateoryid=139&modelid=15710 (accessed 29 December 2007).
- CHANG, L. and TAN, Y.-P., 2004, Effective use of spatial and spectral correlations for color filter array demosaicking. *IEEE Transactions on Consumer Electronics*, **50**, pp. 355–365.
- CHANG, L. and TAN, Y.-P., 2006, Hybrid color filter array demosaicking for effective artifact suppression. *Journal of Electronic Imaging*, **15**, p. 013 003.
- CHUNG, K.-H. and CHAN, Y.-H., 2006, Color demosaicking using variance of color differences. *IEEE Transactions on Image Processing*, **15**, pp. 2944–2955.
- CLARK, R., 2005a, Dynamic range and transfer functions of digital images and comparison to film (intensity detail of images). Available online at: clarkvision.com/imagedetail/dynamicrange2/ (accessed 29 May 2006).
- CLARK, R., 2005b, Dynamic range of an image (intensity detail of an image). Available online at: clarkvision.com/imagedetail/dynamicrange/ (accessed 29 May 2006).
- CLARK, R., 2006, Procedures for evaluating digital camera sensor noise, dynamic range, and full well capacities; Canon 1D Mark II analysis. Available online at: www.clarkvision.com/imagedetail/evaluation-1d2/ (accessed 30 May 2006).
- CLARK, R., 2009, Digital camera sensor performance summary. Available online at: <http://clarkvision.com/imagedetail/digital.sensor.performance.summary> (accessed 1 December 2009).
- COFFIN, D., 2008, Decoding raw digital photos in Linux. Available online at: cybercom.net/~dcoffin/dcrawl/ (accessed 11 February 2008).
- DE LAVARÈNE, B.C., ALLEYSSON, D. and HÉRAULT, J., 2007, Practical implementation of LMMSE demosaicking using luminance and chrominance spaces. *Computer Vision and Image Understanding*, **107**, pp. 3–13.
- DIGITAL PHOTOGRAPHY REVIEW, 2000, Dynamic range in RAW mode. Available online at: www.dpreview.com/news/0011/00111701dynamicrange_raw.asp (accessed 30 June 2005).
- EASTMAN KODAK COMPANY, 2001, Charge-coupled device (CCD) image sensors. Available online at: www.kodak.com/US/en/digital/pdf/ccdPrimerPart2.pdf (accessed 21 March 2005).
- EASTMAN KODAK COMPANY, 1999, Conversion of light (photons) to electronic charge. Available online at: www.kodak.com/US/en/digital/pdf/ccdPrimerPart1.pdf (accessed 21 March 2005).
- EL GAMAL, A. and ELTOUKHY, H., 2005, CMOS image sensors. *IEEE Circuits and Devices Magazine*, **21**, pp. 6–20.
- FECHNER, G.T., 1860, *Elemente der Psychophysik* (Leipzig: Breitkopf und Härtel).
- FOVEON, 2008, Foveon. Available online at: www.foveon.com (accessed 27 December 2007).
- FRASER, B., 2004, Understanding digital RAW capture. Available online at: www.adobe.com/digitalimag/pdfs/understanding_digitalrawcapture.pdf (accessed 13 June 2005).
- FRASER, B., 2005, *Real World Camera RAW with Adobe Photoshop CS2. Industrial-strength Production Techniques* (Berkeley, CA: Peachpit Press, Adobe).
- FRASER, B., MURPHY, C. and BUNTING, F., 2004, *Real World Color Management. Industrial-Strength Production Techniques* (Berkeley, CA: Peachpit Press).
- GIORGIANNI, E.J. and MADDEN, T.E., 1998, *Digital Color Management: Encoding Solutions* (Reading, MA: Prentice Hall).
- GONZALEZ, R.C. and WOODS, R.E., 2002, *Digital Image Processing* (Upper Saddle River, NJ: Prentice-Hall).

- GUNTURK, B.K., ALTUNBASAK, Y. and MERSEREAU, R.M., 2002, Color plane interpolation using alternating projections. *IEEE Transactions On Image Processing*, **11**, pp. 997–103.
- GUNTURK, B.K., GLOTZBACH, J., ALTUNBASAK, Y., SCHAFER, R.W. and MERSEREAU, R.M., 2005, Demosaicking: color filter array interpolation. Exploring the imaging process and the correlations among three color planes in single-chip digital cameras. *IEEE Signal Processing Magazine*, **22**, pp. 44–54.
- HAMILTON, E., 1992, JPEG file interchange format. Available online at: www.jpeg.org/public/jfif.pdf (accessed 19 July 2006).
- HASS, C., 2007, JPEG compression, quality and file size. Available online at: www.impulseadventure.com/photo/jpeg-compression.html (accessed 05 March 2007).
- HENDRY, G.A.F., HOUGHTON, J.D. and BROWN, S.B., 1987, Tansley review no. 11. The degradation of chlorophyll – a biological enigma. *New Phytologist*, **107**, pp. 255–302.
- HIRAKAWA, K. and PARKS, T.W., 2003, Adaptive homogeneity-directed demosaicing algorithm. In *Proceedings of IEEE International Conference on Image Processing*, 2003, *S.n.* (Piscataway, NJ: IEEE), pp. 669–672.
- HOFFMANN, S., 2007, Digital camera RAW converter comparisons. Available online at: www.sphoto.com/techinfo/rawconverters/rawconverters.htm (accessed 06 November 2007).
- HOLST, G.C., 1996, *CCD Arrays, Cameras and Displays* (Washington, DC: SPIE – The International Society for Optical Engineering).
- HOWELL, S.B., 2006, *Handbook of CCD Astronomy* (Cambridge, UK: Cambridge University Press).
- HUNG, P.-C., 2006, Color theory and its applications to digital still cameras. In *Image Sensors and Signal Processing for Digital Still Cameras*, J. Nakamura (Ed.), pp. 205–221 (Boca Raton, FL: Taylor & Francis).
- HUNT, R.W.G., 1999, Why is black and white so important in colour? In *Colour Imaging, Vision and Technology*, L.W. MacDonald and M.T. Luo (Eds), pp. 3–15 (Chichester, UK: Wiley).
- IMAGING RESOURCE, 2007, Nikon D200 Imatest results. Available online at: www.imaging-resource.com/PRODS/D200/D200IMATEST.HTM (accessed 11 December 2007).
- JANESICK, J.R., 2001, *Scientific Charge-Coupled Devices* (Bellingham, WA: SPIE).
- JANESICK, J.R. and BLOUKE, M.M., 1987, Sky on a chip: the fabulous CCD. *Sky and Telescope Magazine*, **74**, pp. 238–242.
- JEITA, 2002, Exchangeable image file format for digital still cameras: Exif Version 2.2. Available online at: www.exif.org/Exif22-2-.PDF (accessed 25 March 2006).
- KIMMEL, R., 1999, Demosaicking: image reconstruction from color CCD samples. *IEEE Transactions on Image Processing*, **8**, pp. 1221–1228.
- KOREN, N., 2001, Making fine prints in your digital darkroom. Tonal quality and dynamic range in digital cameras. Available online at: www.normankoren.com/digital_tonality.html (accessed 29 May 2006).
- LAM, E.Y. and FUNG, G.S.K., 2009, Automatic white balancing in digital photography. In *Single-Sensor Digital Color Imaging Fundamentals*, R. Lukac (Ed.), pp. 267–294 (Boca Raton, FL: CRC).
- LI, X., 2005, Demosaicking by successive approximation. *IEEE Transactions on Image Processing*, **14**, pp. 370–379.
- LIAN, N.-X., CHANG, L., TAN, Y.-P. and ZAGORODNOV, V., 2007, Adaptive filtering for color filter array demosaicking. *IEEE Transactions on Image Processing*, **16**, pp. 2515–2525.
- LIVINGSTON, M., 2002, *Vision and Art. The Biology of Seeing* (New York, NY: Harry N. Abrams).
- LU, W. and TAN, Y.-P., 2003, Color filter array demosaicking: new method and performance measures. *IEEE Transaction on Image Processing*, **12**, pp. 1194–1210.
- LUKAC, R., 2009, Single-sensor digital color imaging fundamentals. In *Single-Sensor Digital Color Imaging Fundamentals*, R. Lukac (Ed.), pp. 1–29 (Boca Raton, FL: CRC Press).

- LUKAC, R. and PLATANIOTIS, K.N., 2004, Demosaicked image postprocessing using local color ratios. *IEEE Transactions on Circuits and Systems for Video Technology*, **14**, pp. 914–920.
- LYON, R.F. and HUBEL, P.M., 2002, Eyeing the camera: into the next century. In *10th Color Imaging Conference: Color Science and Engineering Systems, Technologies, Applications*, Scottsdale, USA, S.n. (Springfield, MA: IS&T–SID), pp. 349–355.
- MENON, D., ANDRIANI, S. and CALVAGNO, G., 2007, Demosaicing with directional filtering and a posteriori decision. *IEEE Transaction on Image Processing*, **16**, pp. 132–141.
- MIAO, L., QI, H., RAMANATH, R. and SNYDER, W.E., 2006, Binary tree-based generic demosaicking algorithm for multispectral filter arrays. *IEEE Transaction on Image Processing*, **15**, pp. 3550–3558.
- MOH, J., LOW, H. and WIJNTJES, G., 2005, Characterization of the Nikon D70 digital Camera. Available online at: scien.stanford.edu/class/psych221/projects/05/joanmoh/index.html (accessed 17 April 2006).
- MURESAN, D.D. and PARKS, T.W., 2005, Demosaicing using optimal recovery. *IEEE Transactions on Image Processing*, **14**, pp. 267–278.
- MURRAY, J.D., 1996, SPIFF: still picture interchange file format. *Dr. Dobb's Journal of Software Tools*, **21**, pp. 34–41.
- NAKAMURA, J., 2006, Basics of image sensors. In *Image Sensors and Signal Processing for Digital Still Cameras*, J. Nakamura (Ed.), pp. 53–93 (Boca Raton, FL: Taylor & Francis).
- PARULSKI, K.A., 1985, Color filters and processing alternatives for one-chip cameras. *IEEE Transactions on Electronic Devices*, **32**, pp. 1381–1389.
- PARULSKI, K.A. and REISCH, R., 2009, Digital camera image storage formats. In *Single-Sensor Imaging. Methods and Applications for Digital Cameras*, R. Lukac (Ed.), pp. 351–379 (Boca Raton, FL: CRC Press).
- PARULSKI, K. and SPAULDING, K., 2003, Color image processing for digital cameras. In *Digital Color Imaging Handbook*, G. Sharma (Ed.), pp. 728–757 (Boca Raton, FL: CRC).
- PENNEBAKER, W.B. and MITCHELL, J.L., 1993, *JPEG: Still Image Data Compression Standard* (New York, NY: Van Nostrand Reinhold).
- PERTIERRA, J.P. and REEL STREAM LLC, 2005, Dynamic range, latitude and the quest for digital film. Available online at: www.reel-stream.com/DynamicRangeAndLatitude.pdf (accessed 29 May 2006).
- PETRIE, G., 2003, Airborne digital frame cameras. The technology is really improving! *GeoInformatics*, **6**, pp. 18–27.
- PICTUREFLOW LLC, 2007, WhiBal™ white balance reference card. Available online at: www.pictureflow.com/products/whibal/index.html (accessed 11 December 2007).
- RAMANATH, R. and SNYDER, W., 2003, Adaptive demosaicking. *Journal of Electronic Imaging*, **12**, pp. 633–642.
- RAMANATH, R., SNYDER, W., YOO, Y. and DREW, M., 2005, Color image processing pipeline. a general survey of digital camera processing. *IEEE Signal Processing Magazine*, **22**, pp. 34–43.
- RITTER, N. and RUTH, M., 2000, GeoTIFF format specification. GeoTIFF revision 1.0. Available online at: www.remotesensing.org/geotiff/spec/geotiffhome.html (accessed 04 March 2006).
- SATO, K., 2006, Image-processing algorithms. In *Image Sensors and Signal Processing for Digital Still Cameras*, J. Nakamura (Ed.), pp. 223–253 (Boca Raton, FL: Taylor & Francis).
- SHEPPARD, R., 2005, *Adobe Camera Raw for Digital Photographers Only* (Indianapolis, IN: Wiley).
- SPECHT, J., 2006, OpenRAW. Digital image preservation through open documentation. Available online at: www.openraw.org (accessed 19 September 2007).
- STEINMUELLER, U., 2003, The raw deal. Why you should use the raw file format with digital cameras. Available online at: www.jirvana.com/shutterbug/Aug_002_raw_deal.pdf (accessed 15 June 2005).

- STEINMUELLER, B. and STEINMUELLER, U., 2004, *Die digitale Dunkelkammer. Vom Kamera-File zum perfekten Print: Arbeitsschritte, Techniken, Werkzeuge* (Heidelberg, Germany: dpunkt.verlag).
- STEVENS, S.S., 1961a, The psychophysics of sensory function. In *Sensory Communication*, W.A. Rosenblith (Ed.), pp. 1–33 (Cambridge, MA: MIT Press).
- STEVENS, S.S., 1961b, To honor Fechner and repeal his law. *Science*, **133**, pp. 80–86.
- STONE, M.C., 2003, *A Field Guide to Digital Color* (Natick, MA: AK Peters).
- THEUWISSEN, A.J.P., 1995, *Solid-State Imaging with Charge-Coupled Devices* (Dordrecht, The Netherlands: Kluwer Academic Publishers).
- TOYODA, K., 2006, Digital still cameras at a glance. In *Image Sensors and Signal Processing for Digital Still Cameras*, J. Nakamura (Ed.), pp. 1–19 (Boca Raton, FL: Taylor & Francis).
- VERHOEVEN, G., 2007, Becoming a NIR-sensitive aerial archaeologist. In *Remote Sensing for Agriculture, Ecosystems, and Hydrology IX*, 17–19 September 2007, Florence, Italy, pp. 333–345.
- VERHOEVEN, G., 2008, Imaging the invisible using modified digital still cameras for straightforward and low-cost archaeological near-infrared photography. *Journal of Archaeological Science*, **35**, pp. 3087–3100.
- VERHOEVEN, G., DE SMET, P., POELMAN, D., VERMEULEN, F., 2009, Spectral characterisation of a digital camera's NIR-modification to enhance archaeological revelation. *IEEE Transactions on Geoscience and Remote Sensing*, **47**, 3456–3468.
- WALKER, J.S., 2004, *Physics* (New Jersey, NJ: Pearson Education).
- WALLACE, G.K., 1991, The JPEG still picture compression standard. *Communications of the ACM*, **34**, pp. 30–44.
- WATANABE, S., 2006, Image-processing engines. In *Image Sensors and Signal Processing for Digital Still Cameras*, J. Nakamura (Ed.), pp. 255–276 (Boca Raton, FL: Taylor & Francis).
- WHITE, R., 2005, *How Digital Photography Works* (Indianapolis, IN: Que).
- WIENKE, J., 2006, Hybrid raw conversion. Using B&W to improve your color images and increase dynamic range. Available online at: www.luminous-landscape.com/tutorials/hybrid-conversion.shtml (accessed 21 June 2006).
- WOLFE, J.M., KLUENDER, K.R., LEVI, D.M., BARTOSHUK, L.M., HERZ, R.S., KLATZKY, R.L. and LEDERMAN, S.J., 2006, *Sensation and Perception* (Sunderland, MA: Sinauer).
- WU, X. and ZHANG, N., 2004, Primary-consistent soft-decision color demosaicking for digital cameras. *IEEE Transactions on Image Processing*, **13**, pp. 1263–1274.
- YAMADA, T., 2006, CCD image sensors. In *Image Sensors and Signal Processing for Digital Still Cameras*, J. Nakamura (Ed.), pp. 95–141 (Boca Raton, FL: Taylor & Francis).
- YOSHIDA, H., 2006, Evaluation of image quality. In *Image Sensors and Signal Processing for Digital Still Cameras*, J. Nakamura (Ed.), pp. 277–303 (Boca Raton, FL: Taylor & Francis).