

High dynamic range imaging for archaeological recording

David Wheatley¹

Archaeology, Faculty of Arts and Humanities,

University of Southampton,

Southampton, SO15 1BF

dww@soton.ac.uk

+44 2380 594779 (office)

+44 7801 435297 (mobile)

¹ Archaeology, University of Southampton, Southampton, UK

Abstract

This paper notes the adoption of digital photography as a primary recording means within archaeology, and reviews some issues and problems that this presents. Particular attention is given to the problems of recording high-contrast scenes in archaeology and High Dynamic Range imaging using multiple exposures is suggested as a means of providing an archive of high-contrast scenes that can later be tone-mapped to provide a variety of visualisations. Exposure fusion is also considered, although it is noted that this has some disadvantages. Three case studies are then presented (1) a very high contrast photograph taken from within a rock-cut tomb at Cala Morell, Menorca (2) an archaeological test pitting exercise requiring rapid acquisition of photographic records in challenging circumstances and (3) legacy material consisting of three differently exposed colour positive (slide) photographs of the same scene. In each case, HDR methods are shown to significantly aid the generation of a high quality illustrative record photograph, and it is concluded that HDR imaging could serve an effective role in archaeological photographic recording, although there remain problems of archiving and distributing HDR radiance map data.

Keywords: archaeology; photography; HDR; recording

Archaeological record photography

Photography has been a fundamental part of archaeological recording for well over half a century. By the 1950s, for example, Cookson published a dedicated guide to archaeological photography in which he reports that ‘*one cannot imagine any archaeologist today cutting even the simplest trial trench without photographing it from three or four positions*’ (1954:11). The aim of record photography is to capture as much of the variation in texture and colour within a scene or surface as possible. Where this is part of an excavation recording process, the aim is to create a record that complements the other primary records such as plan and section drawings, and which may subsequently be used to effectively illustrate reports and some considerable attention has been given over the years to the technical detail of how this can be best achieved (BAJR 2006; Cookson 1954; Conlon 1973; Dorrell 1989; Fischer 2009a, b; Howell and Blanc 1992; Schlitz 2007).

In recent years digital photography has almost completely replaced film-based photography in many contexts and has begun to be used as a recording tool in archaeology. With a few exceptions, however (see for example Woolliscroft 2010), archaeology has not yet given serious consideration to the methodological implications of this change. At a trivial but practical level, digital photography provides the ability to check results in the field and to record large numbers of images without changing or wasting film, whereas film produces a relatively robust physical result for archiving and does not require batteries. More profoundly, though, images produced by digital sensors differ from film-based photographs in many ways. Most significantly there are differences in *archival properties*, and in the *quality* of images that can be obtained in terms of both resolution (acuity) and dynamic range. For detailed information on the capabilities of digital sensors see, for example, Clark (2010).

Although archiving of digital collections of images has been discussed for at least fifteen years (e.g. Ester 1996), there remains a deep conservatism within archaeology that manifests itself in concerns over the most appropriate formats and metadata standards for archiving digital images, and this may be one reason why archaeology has been reluctant to adopt digital photography as a replacement for film. The *Archaeological Archives Forum Guide to Best Practice in Field Archaeology* notes that digital photographs are increasingly used, but that there needs to be a ‘*clearly established procedure for long-term preservation*’ and states that ‘*black and white film processed to BS5699 is the archival ideal*’ (Brown 2007: 13) and

this view is reflected by many museums and archives, who will not yet accept digital images as archival records. Issues of archiving digital material, including images, are largely outside the remit of this paper (but see e.g. Digital Preservation Coalition 2010; Hunter 2000; Kenney and Rieger 2000; Marty 2009; Parry 1998; Richards and Robinson 2000) but it is interesting to note that, despite advice from JISC Digital Media, formerly TASI (JISC 2010), practitioners (e.g. Andrews et al. 2006) and aerial photography specialists (e.g. Verhoeven 2010) all of whom outline the advantages of camera RAW data and the open – albeit proprietary – Adobe Digital Negative format, at the time of writing the Archaeology Data Service (ADS) expects deposition of digital photographs in TIF format and currently has no policy on archiving of RAW images in *any* format. This is, arguably, the digital equivalent of a physical archive refusing to accept negatives, only prints.

Ultimately, however, this may be of little relevance to the adoption of digital photography as a recording method by many fieldworkers. Indeed it seems increasingly likely that archaeology may have no choice but to ‘go digital’ for primary recording, at least for colour photographs. In 2009, for example, Kodak announced that they were ceasing production of Kodachrome slide film (Topping 2009) so that when existing stocks are exhausted, it will no longer be available for colour recording. Like vinyl records, black and white photography is unlikely to suffer the same fate immediately because there is a sufficiently large enthusiasm for its particular characteristics, but in general it seems very probable that ‘born digital’ photographs will make up the majority of archaeological record images in the near future.

Fortunately, the issue of image quality is rather more tractable, and may even offer positive benefits to archaeology rather than solely present us with problems. The resolution of digital photography is decreasingly a serious concern because, while early digital sensors did not offer recording resolution close to film, most modern digital cameras in use for archaeological recording offer ample resolution. There is, however, still a problem over the more limited dynamic range that most electronic sensors offer over, particularly, black and white film. Additionally, many photographers find the response curves of digital sensors, which typically have a rather abrupt graduation in the highly saturated (bright) regions of the image, rather less useable than the more gradual transition to zero density (the ‘shoulder’) found in the response curves of most film and photographic papers.

Dynamic range in archaeological photographs

Dynamic range is of particular concern in this context because the majority of archaeological field recording takes place out of doors, where it can be difficult to control the lighting in any given scene. If the contrast in a scene is too great – where direct sunlight falls on part of a surface while another part is in shadow, for example – it can be almost impossible to create a single photograph that retains sufficient detail in all areas of the scene. Dorrell explains this problem very clearly:

‘The factor that can make or mar the effectiveness of site photographs is the strength and direction of the natural light ... Direct strong sunlight, particularly if it falls cross the scene diagonally, is probably the worst possible lighting. Not only will it raise the contrast to unacceptable levels, giving solid black shadows, or burnt-out highlights, or both, but the pattern of light and shade may form an outline of black stripes and masses far more obvious to the eye than the shape of the walls themselves’ (1994: 127).

Archaeological photography presents many situations in which the contrast can be extremely high. Almost any photograph that attempts to show the interior of a cave, chambered tomb or other unlit structure as well as the view through the entrance or window will likely fall into this category. Typically a digital camera sensor may have the ability to register a contrast ratio of around 1,000:1, while the scene to be recorded may exceed 70,000:1 in some circumstances (Fig 1).

The task of the photographer is to adjust either the time of the exposure (normally using the camera's shutter speed) or the intensity of light permitted to fall on the sensor (normally using the lens's aperture) so that the range of reflected light intensities in the scene fall within the dynamic range of the sensor. The problem with high contrast scenes is that photographic films and digital sensors are only able to respond to a limited range of radiances and the range of reflected light intensities within a high contrast scene quite often exceeds that range, a problem that is most obvious with colour transparency films which have notoriously limited dynamic range sensitivity. If the photographer (or camera metering system) chooses an exposure so that the luminance values within shaded areas of the scene correspond to the available dynamic range of the sensor, then regions in direct sunlight will exceed the sensor's saturation point and will be recorded as uniform maximum value. Conversely, if the

photographer chooses an exposure for which the sunlit parts of the scene are correctly mapped to the dynamic range of the sensor, then much of the regions in shade will fall below the minimum sensitivity of the film or sensor where they will be recorded as uniform black (or the signal-to-noise ratio becomes unacceptably high). There may still be an 'optimal' exposure for the scene, which can be estimated using either the automatic metering system in the camera or a separate light meter and some method such as the 'zone' system advocated by Adams (1970) but 'optimal' in this sense means only that the loss of information at either end of the dynamic range is minimised, not eliminated altogether.

We tend not to notice this problem when we look at a scene, because the human visual perception system does not work like a camera. Although there are some well documented similarities between human eyes and cameras, the human visual perceptual system as a whole consists of far more than a single 'camera' generating a single static image. Our eyes can move around the different areas of the scene while changing their sensitivity (by contracting the pupil) which allows simultaneous perception of luminance levels that vary over a range of 3.7 log units (Kunkel and Reinhard 2010). The sensitivity of the human visual system is also adjusted in other ways such as the bleaching of photopigment in the rod and cone cells and adaptation of photoreceptor mechanisms (Reinhard et al. 2010: 243-251). Although these latter are far less rapid, taking from a few seconds to a few minutes to fully adjust, they allow the human visual system to be sensitive to up to 10 orders of magnitude of sensitivity in total (Fig 1). Digital imaging systems typically respond to variations of not much more than 2 orders of magnitude (Reinhard et al. 2010: 4-5).

The problem is very familiar to both amateur and professional photographers and also to archaeological fieldworkers and it is therefore unsurprising there are a number of established ways to mitigate it. Simply avoiding situations in which lighting is from the side and casting strong shadows by ensuring that photographs are taken with the sun behind the camera is one approach, but is generally rather undesirable as *'this would mean that any surface facing the camera would be in direct, straight-on, light, and its surface texture would therefore be lost'* (Dorrell 1994: 127). In a typical archaeological situation, mitigation usually involves reducing the difference in luminance between regions in the scene to make it easier to 'map' the scene to the available dynamic range of the sensor or film. This can be achieved by reducing the absolute luminance of those areas in direct sunlight by shading them with – variously – tarpaulins, gazebos, vehicles or even students. Alternatively, regions in shadow

can be 'filled' with additional light so as to raise their luminance, a task most easily achieved using either a reflector or electronic flash. Both of these approaches can be highly effective but neither is without disadvantages: shading can be both time consuming and labour intensive (and is sometimes impossible due to the height or position of the sun) while few field archaeologists have both the expertise and the equipment to deploy effective fill lighting.

Partly for these reasons, archaeologists have also long recognised the need to 'bracket' record photographs. This means that the photographer establishes the 'optimal' exposure as above and takes one photograph using these values. They then also record at least one additional image that is intentionally over-exposed, and one that is under-exposed to make it more likely that both bright and dark areas are correctly exposed in at least one image. Although these are sometimes discarded, there is a significant legacy archive of archaeological record photographs recorded in this way.

High Dynamic Range (HDR) imaging

High dynamic range photography is an increasingly popular branch of photography and a very active area of research that deals with the recording and representation of scenes with extended dynamic ranges – in other words, high contrast scenes such as those described above. Although it is possible to foresee cameras that record high dynamic range information directly, the most popular approach is currently based on the approach of Debrevic and Malik (2008) who described a method of using multiple images taken at different exposures to create a single 'High Dynamic Range radiance map'. A full description of the method is beyond the scope of this paper, but in summary it involves utilising two or more images of the same scene taken in the same lighting conditions to reconstruct the 'response function' of the imaging process. Alternative methods of deriving the response function do exist, such as that proposed by Mitsunaga and Nayar (1999) and a full treatment of the derivation of the response function and various other processing steps that may be needed such as noise removal, image alignment, control of lens flare and automatic removal of 'ghost' elements – which arise when an object does not appear in all of the source images or moves between exposures – can be found in Reinhard et al (2010:171-197). The response function is then used to combine multiple exposures into a single High Dynamic Range (HDR) 'radiance map' whose pixel values are proportional to the original luminance values in the scene, rather than (as in a conventional photograph) the result of some non-linear function.

Two implications of the methodology may be relevant in archaeological photography. Firstly, the approach uses the conventional photographic assumption of *reciprocity*: the property of a film or sensor that ensures that there is a reciprocal equivalence between the time of exposure and the intensity of illumination – in other words if we double the exposure time and half the intensity of the light falling on it then the sensor or film will provide the same response. The second assumption is that the images used to reconstruct the response function are recorded in the same lighting conditions. The first of these provides no significant problems for archaeological situations, as reciprocity can reasonably be assumed to hold within the range of exposure times and intensities (expressed as aperture stops or f-stops) that would be used in archaeological recording. This would only become relevant were we to require exposure times of more than around 30 seconds, or encounter intensities of light well beyond those we would expect. At these extremes then 'reciprocity failure' could occur, and the method may not produce reliable results (although see discussion of 'exposure fusion' below). The second assumption is of more practical concern, as it is not uncommon – in Britain at least – for lighting conditions to change quite rapidly during field recording. It follows that care should be taken to record the various exposures quickly and, where possible, to choose periods of relatively stable light. In practice, this means taking the series of images in quick succession, perhaps taking advantage of modern Digital Camera's automation functions for bracketing, and/or waiting for a large cloud or a large gap in the clouds to pass across the sun.

Software to calculate the response function and/or to combine source images into HDR radiance maps is readily available and is widely used within the photographic community to enable photographers to record high contrast scenes and to obtain interesting pictorial effects. Various methods can be used to 'tone map' the values in the HDR file to the LDR of conventional monitors, output devices and colour spaces.

Tone mapping

The HDR radiance map is obviously of interest as a useful record of the illumination of a scene, but clearly it does not contain values that fall within the low dynamic range (LDR) gamut of an output device such as a monitor or printer and so the data cannot be viewed directly. In order to view an HDR radiance map, therefore, some processing or *tone-mapping* needs to take place between the radiance values in the HDR file and the LDR of an output device.

There are a very wide variety of methods and approaches that can be used for tone mapping, many of which are described in Reinhard (2010). Often these take the human visual system as a model, and attempt to generate representations of the full (HDR) luminance scale in such a way as to mimic the response of a human observer in some form. Some approaches compress or select ranges from the HDR histogram which can then be mapped through some linear or – more commonly – sigmoidal function to the range of values in the output gamut. Broadly, these can be referred to as ‘*tone compression*’ methods but they can often lead to a final image with poor overall contrast or in which parts of the dynamic range in the HDR radiance map are poorly represented. For these reasons region-based and adaptive local contrast operations are often used for tone mapping (referred to below as ‘*detail enhancement*’ methods) as these are capable of producing outputs that feature good local contrast, while also retaining the local variations in luminance that best evoke the texture and tonality of the original scene.

Alternatives to High Dynamic Range methods

HDR is a physically-based method for combining information from several exposures. Where the HDR radiance map is derived using the method of Debevec and Malik (2008), it requires that the relative exposure values (EV) for each image are known, or at least can be estimated (see Avebury Cove example below). Where the HDR radiance map is not required then it may be more efficient to use alternative one-step procedures for combining multiple-exposure sequences. Manual and ‘ad hoc’ methods for doing this have always existed such as ‘sandwiching’ of negatives together to make single prints, and the use of image processing software to combine images in similar ways. Automated or semi-automated methods for doing this are referred to as ‘*exposure fusion*’ (Mertens et al. 2009; Reinhard et al. 2010: 400-404). Exposure fusion methods usually rely on computing perceptual quality measures for each pixel, and combining those areas of each image that produce the ‘best’ results to form the final image, and they have the advantage that they are computationally less intensive because there is no need to deduce a physically-based radiance function, and they do not require that we know the relative EV of each exposure in advance. This may be particularly relevant to the potential use of legacy archaeological material, which may consist of sequences of photographic negatives or slides whose shutter speeds and apertures have not been recorded and for this reason exposure fusion was used, and is illustrated in the Avebury Cove example, below.

One minor disadvantage of this approach is that it ‘*cannot extend the dynamic range of the original picture*’ (Mertens et al. 2009: 161) but a more significant disadvantage in the context of archaeological recording may be that it negates one of the long-term potential benefits of HDR because all decisions as to how to make the final image are made by the publisher/creator and not – as is possible if HDR radiance maps are archived or distributed – by the researcher/reader. The real potential for HDR imaging may ultimately lie in the publication and distribution of HDR radiance maps themselves. Although they do not contain information that can be viewed directly, it should be possible to distribute viewers (as standalone software or browser plugins) that perform tone mapping ‘on the fly’. In this way, HDR images might be useable in archaeology in a similar way to ‘bubbleworld’ methods such as Quick Time VR (Jeffrey 2001) or Polynomial Texture Maps (Earl et al. 2010a; Earl et al. 2010b) in that the researcher or reader is permitted to determine the optimum tone mapping parameters just as they determine the view position in the case of QTVR, and the lighting conditions for Polynomial Texture Maps. In short, it is difficult to entirely predict in advance what aspects of the image a future researcher/reader will wish to explore, and so it would be better to create and distribute physical HDR radiance maps where possible, so as to permit them to manipulate the final image appearance to suit their own interests and needs.

The potential for HDR in archaeological recording

The potential benefits to archaeological field recording of this approach should be fairly clear by now. The raw data to construct HDR radiance maps is relatively easy to obtain: it requires only that multiple exposures are made of the scene that is being recorded, which most archaeologists already do, and it complements rather than replaces existing photographic recording methods. The immediate benefit is that it permits a far wider range of textures and colours to be accurately recorded and represented than is currently possible. We might also expect to achieve significantly improved recording outcomes where we have high contrast scenes but are unable to employ – for whatever reason – the conventional mitigation strategies described above. Moreover, the method should be applicable to a significant body of legacy photographic material: we may be able to use this approach to generate HDR radiance maps, and hence significantly more useful photographs, from existing bracketed slide and film records that exist within physical archaeological archives.

Three examples are presented here to explore the potential of HDR imaging methods in archaeological recording. The first (Cala Morell) illustrates what has become the

conventional application of HDR methods to a typically high contrast scene, the second (Itchen Abbas) is intended to illustrate how HDR recording can be incorporated into a 'born digital' approach to site photography, while the final example (Avebury Cove) explores whether HDR methods may be used with 'legacy' photographic materials to provide added value.

In all the case studies below, generation of HDR radiance maps and subsequent tone mapping was done using *Photomatix Pro v3.0* (HDR 2010). Although proprietary, this is a relatively inexpensive software solution that is widely used within the wider photographic community and provided all the functionality required for these examples. Usefully, it allows for batch processing large numbers of source images automatically, which proved extremely useful for the Itchen Abbas case study below. It also has an exceptional level of functionality for adjusting and manipulating the tone-mapping process, although this is not fully explored in these simple examples. Many other software solutions are available both for the generation of HDR radiance maps and for tone mapping and these include both commercial and Open Source platforms.

Cala Morell – a very high contrast photograph

The requirement in this example was to record the interior of one of the rock-cut tombs at Cala Morell, Menorca, with a doorway through which direct sunlight was visible. This presents a photographic problem for which HDR has been widely applied and which is common in archaeology. In this example, a wide angle shift lens was used on the camera.

To ensure that the full dynamic range of the scene was represented, seven exposures were made with a digital camera mounted on a tripod at 2EV intervals, centred on the exposure recommended by the camera's metering system (Figure 2) Although RAW images were recorded (shown here using settings as shot in the camera) it is clear that the recommended exposure provides acceptable detail in few areas of the image, and that none of the exposures provides an acceptable level of detail within the cave and through the entrance doorway. The exposure at +2EV provides an acceptable image of the interior of the cave – note the green algae growing on the walls, for example – but much of the detail of the surface texture of the floor is poorly rendered, the shape of the entrance and the floor inside it show no detail at all and – most obviously – the view through the doorway is also entirely 'burnt out'. Some additional detail in the original exposures could be recovered through careful processing of

the RAW image data, but nowhere near enough to achieve an acceptable result. The region of the image seen through the doorway appears to need an exposure of some 7 or 8EV less than the average in order to render acceptable detail. All seven images were used to generate an HDR radiance map which also allowed the contrast ratio in the scene to be quantified as approximately 60,000:1 (estimated by the Photomatix software).

The HDR image was then tone-mapped using the ‘detail enhancement’ option of Photomatix software. Some experimentation with the parameters quickly enabled an acceptable image to be produced (Figure 2, bottom right) that rendered all areas of the scene within the gamut of a computer monitor and enabled adequate representation in print. The quality of the image has a slightly strange tonality that is typical for images that have been tone mapped using local contrast optimisation methods, although the overall colour balance of the image appears reasonable.

The resulting image is a useful record that illustrates well the various textures and colours of the scene. Unlike any of the source images, the tone-mapped image shows good textural detail on the ceiling, walls and floor of the cave. It also renders the floor in front of the door, the door jambs and the view through the entrance itself. Although the image appears slightly unnatural, it is arguably a better rendering of the perceptual experience of being inside the tomb, from where the human visual system is capable of adapting to the extreme luminance levels outside. Alternative renderings of the HDR radiance map are possible using different parameters of the tone mapping. It is worth noting that the entire HDR recording process took less than 2 minutes. The only other method of achieving an acceptable result would have been to use artificial lighting to raise the light level inside cave which would have been extremely difficult given the limited space available and the absence of mains electricity – it is unlikely that portable flash would have provided sufficient additional light. It would also have been both time consuming and expensive.

Itchen Abbas – ‘born digital’ photography for HDR

For the second example, field recording was undertaken with the prior intention of making HDR radiance maps from the images. The University of Southampton field school at Itchen Abbas, Winchester was used for this purpose as the problems faced were typical of archaeological photographic recording: weather conditions varied from overcast

(occasionally rainy) to bright sunlight and the contrast of some images was also significantly increased where test pits revealed white chalk features against dark red/brown clays.

Reinhard *et al* (2010: 182) provides guidelines for successfully recording HDR images and these proved straightforward enough to incorporate into a conventional archaeological recording workflow. Record photographs were taken using a high-quality digital SLR camera mounted on a tripod, and for each record photograph a sequence of images were taken at fixed EV intervals. The ISO setting was set to a constant 200 ISO, white balance of the camera was fixed prior to the excavation to a suitable value, and a standard Gretag/Xrite colour chart was included in each image to enable identification of the photograph in the written registers, and provide colour calibration information. The process was therefore almost unchanged from a conventional archaeological recording workflow in which exposures would be 'bracketed' for safety. The only significant change was that the 'bracketing' was always done using the camera's shutter speed because altering the lens aperture has the effect of altering the depth of focus, and hence would complicate the process of combining images later, and care was taken not to alter ISO and white balance values. Photographs were recorded as RAW images, and – as is conventional – a blackboard with metadata (at least site code, test pit number and north arrow) was incorporated in each image.

Some experimentation was required to establish the appropriate number of exposures and EV intervals to use (each EV step corresponds to a halving or doubling of the exposure time or an increase or decrease of one aperture stop). Initially five images were recorded at an interval of one EV but it was apparent that this failed to fully cover the dynamic range in some of the higher contrast scenes, so a general policy of recording five images at two EV steps was then adopted. This was a relatively small field project resulting in sixty-eight record photographs of test pits, which required 340 raw images totalling 984Mb of data. Post excavation, each sequence of five RAW files was batch processed using Photomatix 3.0 Pro software to generate HDR radiance maps and tone-mapped images using both *tone compression* and *detail enhancement* methods. Because EV information is recorded as metadata within digital photographs, the generation of the HDR radiance map is entirely automated and so batch processing of the images was relatively straightforward. In this case it took around six hours of processing for all 68 sets of five images using a personal computer.

Many of the record images exhibited relatively low contrast and in these examples a single RAW file was fully capable of rendering an acceptable record photograph. In several cases,

the recommended exposure settings of the camera proved not to be suitable, and a published image would have been selected from one of the other nearby exposures. In around 12 of the 68 cases, however, the contrast was sufficiently high in the scene that some mitigation would have been required in the field to ensure an acceptable photograph.

Figure 3 illustrates one of the higher contrast scenes, although not the most extreme. The photograph shows part of a rammed-chalk floor, partly excavated, adjacent to a dark red/brown clay. Although an exposure slightly above the camera meter's recommendation would provide an acceptable result, either the areas of shadow or the stark white of the chalk floor surface would be difficult to render effectively without considerable post-processing of the image. The tone compression method (using only default values) shown as Image A in Figure 3 shows some improvement in this respect, with both the side of the cutting and the surface of the chalk floor exhibiting textural detail, but the detail enhancement method (shown as Image B, also using default values) is a considerable improvement, with good textural detail visible in all areas of the photograph. As with the previous example, the colour and lighting appear slightly unnatural and – for publication purposes – it would probably benefit from experimentation with the tone mapping parameters to achieve a more natural result. As a final point of interest, Image C in Figure 3 shows an automated exposure fusion which has also produced a record image that is very comparable to the results obtained using HDR-based methods.

Avebury Cove – using legacy photography for HDR

The final example concerns the possibility of recovering HDR radiance maps from legacy archaeological photography. Four colour slides were selected from a record made of excavations at Avebury Cove in 2003 (Pollard 2008). These had been taken as a bracketed sequence of images using a tripod in difficult lighting conditions – vertically downwards into a fairly deep cutting adjacent to one of the stones of the cove. The range of dynamic variation had proved too great for the slide film, and use of any of the original images for illustration purposes would have represented a compromise between good representation of the base of the cutting (Figure 3, image 2) or the sides of the cutting (Figure 3, image 3).

The slides were scanned using a film scanner, resulting in the digital images shown as images 1 to 4 in Figure 3. The EV settings of the sequence was not recorded, and so some experimentation was required in order to achieve a plausible reconstruction of the HDR

radiance map. The procedure adopted here was simply to try a reconstruction with a particular set of values, examine the HDR radiance map, and to reject it if the result appeared incorrect in the HDR viewer or the histogram exhibited odd effects. After a few attempts, a sequence of relative EV estimates was arrived at that provided an apparently reasonable reconstruction of the HDR radiance map. One further complication arises with the use of legacy material in that, although these sides had been photographed using a tripod, the scanned images were not perfectly registered (this can be seen in the positions of the black frames surrounding the images of Figure 4) and so it was necessary for the software to perform an additional image alignment step that required re-sampling of the images, resulting in some slight loss of sharpness.

The image shown as Figure 4 image A was then tone mapped from the HDR radiance image using the detail enhancement method. It should be clear that it provides a considerable improvement in the level of textural detail that is represented, permitting a single effective illustration of both the base of the cutting and the sides (which include objects in the section on the left of the image, and a brick wall on the right). Figure 3 image B was created directly from the four scanned slides using exposure fusion. This was a considerably quicker process that did not require any estimation of the relative EV stages, and it should be clear that the result is in many ways equally effective as the tone-mapped image A.

Conclusions and future considerations

The brief experiments reported here suggest that HDR approaches could be of considerable use to archaeological field recording. Although the examples here are all of colour images, the methodology is equally applicable to digital monochrome photography and to ultraviolet or infrared imaging – which is becoming increasingly available through modification of conventional digital cameras (see e.g. Verhoeven 2008, Verhoeven and Schmitt 2010) – and so further archaeological applications in areas such as rock art recording are more than likely.

In future, it is quite possible that the design of digital sensors will be improved in various ways to improve their dynamic range sensitivity, and cameras may become more widely available that are capable of producing a proportional response to the entire range of luminance values that may be encountered in any photographic situation. In this case, the capture of HDR radiance maps may be possible directly rather than requiring the processing of several differently exposed images. Similarly, the availability and development of monitors

capable of displaying a wider range dynamic range is also likely to improve in coming years (Seetzen et al. 2003). Good quality conventional LCD displays are capable of a contrast ratio of around 2,000:1 but this is an area of considerable commercial development (because it impacts on television and home cinema products) and monitors with a contrast ratio of around 200,000:1 (capable of rendering luminance levels from 0.015 to 3000 cd/m²) have been demonstrated. These should ultimately enable a far wider range of luminance levels to be experienced than is currently possible, and their potential widespread availability is a further argument in favour of recording and using HDR radiance maps as part of archaeological archives.

If it is appropriate to store HDR radiance maps as components of archaeological archives, then decisions need to be taken with respect to which formats are the most appropriate and what metadata standards – and indeed content – should be archived. There are several candidate ‘open’ formats that could be used for archiving HDR radiance data including Radiance (.hdr) format and OpenEXR (.exr), although both involve some loss of precision. Floating-point TIF may be the best option in the short term, although this may require as much as 96 bits per pixel resulting in very large files. As intimated above, however, there are similarities here with the need to archive Polynomial Texture Map data and even QTVR-type image data for use in ‘bubbleworld’ virtual reality viewers because – in all these cases – the image depends not only on the archived data but also on the viewer used to interpret that data. Because of this, it would seem prudent to establish new archival policies that permit a far wider range of image data, including all of these, to be deposited and documented.

This paper has suggested that there could be a role in future for HDR-based photography in archaeological recording. The case studies each provide evidence that there are some benefits to be gained from post-processing record photographs (both ‘born digital’ and legacy) to recover the HDR radiance maps. The range of benefits achieved is clearly proportional to the contrast in the original scene, with the first case study providing the clearest benefit. The Itchen Abbas case study suggests that the benefits in many cases may not be significant, although in a few high-contrast scenes the use of HDR is a perfectly adequate alternative to the traditional approach of mitigating contrast using either shading or fill lighting. Although it may provide little benefit for low contrast scenes, the additional cost – both of capture in the field, and of post-processing the results – is surprisingly modest and so incorporating an ‘HDR-friendly’ recording strategy into an existing workflow may be regarded as worthwhile

for the minority of scenes that would significantly benefit. The Avebury Cove example also demonstrates that this approach can achieve useful results on existing legacy material, although this may prove more difficult to process. Clearly further exploration of a wider range of legacy material, including different film types, may provide more reliable conclusions about the future potential.

Acknowledgements

A number of colleagues have commented on this manuscript and I am particularly grateful to Graeme Earl and Tom Goskar for their useful input. Thanks are also due to Geraldine Joffre of the Photomatix Engineering team and to the Southampton University students and supervisors for their interest in and help with the HDR recording experiments during the Itchen Abbas excavation project.

References

- Adams A (1970) *Basic photo*, 1st revision edn. Morgan & Morgan, Hastings-on-Hudson, N.Y.
- Andrews P, Butler Y, Farace J (2006) *Raw workflow from capture to archives : a complete digital photographer's guide to raw imaging*. Focal, Oxford
- BAJR (2006) *Short guide to Digital Photography in Archaeology*. BAJR Guide. British Archaeological Jobs Resource
- Brown DH (2007) *Archaeological Archives: A guide to best practice in creation, compilation, transfer and curation*. Institute of Field Archaeologists, London
- Clark, R. N. (2010) *Digital Camera Sensor Performance Summary* (<http://www.clarkvision.com/articles/digital.sensor.performance.summary/>) Retrieved 15/10/2010
- Conlon VM (1973) *Camera techniques in archaeology*. J. Baker, London,
- Cookson MB (1954) *Photography for archaeologists*. Parrish, London,
- Debevec P, Malik J (2008) *Recovering high dynamic range radiance maps from photographs*. SIGGRAPH '08: ACM SIGGRAPH 2008 classes. ACM, Los Angeles, California, pp 1-10
- Digital Preservation Coalition (2010) *Digital Preservation Coalition Web Site* (<http://www.dpconline.org/>) Retrieved 15/10/2010
- Dorrell PG (1989) *Photography in archaeology and conservation*. Cambridge University Press, Cambridge ; New York

- Dorrell PG (1994) *Photography in archaeology and conservation*, Second edition edn. Cambridge University Press, Cambridge ; New York
- Earl G, Beale G, Martinez K, Pagi H (2010a) Polynomial texture mapping and related imaging technologies for the recording, analysis and presentation of archaeological materials. *Commission V Symposium. International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, Newcastle Upon Tyne, pp 218-223
- Earl G, Martinez K, Malzbender T (2010b) Archaeological applications of polynomial texture mapping: analysis, conservation and representation. *Journal of Archaeological Science* 37
- Ester M (1996) *Digital image collections: issues and practice*. Commission on Preservation and Access, Washington, DC
- Fischer LJ (2009a) *Photography for archaeologists. Part I: Site Specific*. BAJR Practical Guides
- Fischer LJ (2009b) *Photography for archaeologists. Part II Artefact Recording*. BAJR Practical Guides
- HDR (2010) HDR Photomatix website (<http://www.hdrsoft.com/>). Retrieved 15/10/2010.
- Howell CL, Blanc W (1992) *A practical guide to archaeological photography*. Institute of Archaeology, University of California, Los Angeles
- Hunter GS (2000) *Preserving digital information: a how-to-do-it manual*. Neal-Schuman, New York ; London
- Jeffrey S (2001) A Simple Technique For Visualising Three Dimensional Models in Landscape Contexts. *Internet Archaeology* 10 (http://intarch.ac.uk/journal/issue10/jeffrey_index.html) Retrieved June 2010
- JISC (2010) JISC Digital Media Website (<http://www.jiscdigitalmedia.ac.uk/>) Retrieved 15/10/2010
- Kenney AR, Rieger OY (2000) *Moving theory into practice : digital imaging for libraries and archives*. Research Libraries Group, Mountain View, Calif.
- Kunkel T, Reinhard E (2010) A reassessment of the simultaneous dynamic range of the human visual system. *ACM symposium on applied perception in graphics and visualisation*, Los Angeles
- Marty PF (2009) *Digital convergence - libraries, archives, and museums in the information age*. New York: Springer
- Mertens T, Kautz J, Van Reeth F (2009) Exposure Fusion: A Simple and Practical Alternative to High Dynamic Range Photography. *Computer Graphics Forum* 28(1):161-171
- Mitsunaga T, Nayar SK (1999) Radiometric Self Calibration. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'99)*, Fort Collins, pp 374-380

Parry D (1998) *Virtually new : creating the digital collection; a review of digitisation projects in local authority libraries and archives*. Library and Information Commission

Pollard J (2008) *Archaeological excavation in the World Heritage Site*. In: Simmonds S (ed) *Avebury World Heritage Site: values and voices*. Kennet District Council, Devizes, pp 10-12

Reinhard E, Ward G, Pattanaik S, Debevec P, Heidrich W, Myszkowski K (2010) *High dynamic range imaging : acquisition, display, and image-based lighting*, 2nd edn. Morgan Kaufmann, Amsterdam ; Boston

Richards JD, Robinson D (2000) *Digital archives from excavation and fieldwork : a guide to good practice*, 2nd ed. edn. Oxbow, Oxford

Schlitz M (2007) *Archaeological Photography*. In: Peres MR (ed) *The Focal Encyclopedia of Photography: Digital Imaging, Theory and Applications History and Science*, Fourth Edition edn. Focal Press, Oxford

Seetzen H, Whitehead L, Ward G (2003) A high dynamic range display using low and high resolution modulators. *Society for Information Display Symposium*, pp 1450-1453

Topping A (2009) *Mama, they've taken away my nice bright Kodachrome colours: Kodak ends production of film loved by photographers and remembered in Paul Simon song*. The Guardian, Manchester (Newspaper article)

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(12):3087-3100

Verhoeven G (2010) It's all about the format - unleashing the power of RAW aerial photography. *International Journal of Remote Sensing* 31(8): 2009-2042.

Woolliscroft D (2010) Thoughts on the suitability of digital photography for archaeological recording. IFA Website (www.archaeologistsnet/modules/icontent/). Institute for Archaeologists. Retrieved 15/10/2010.

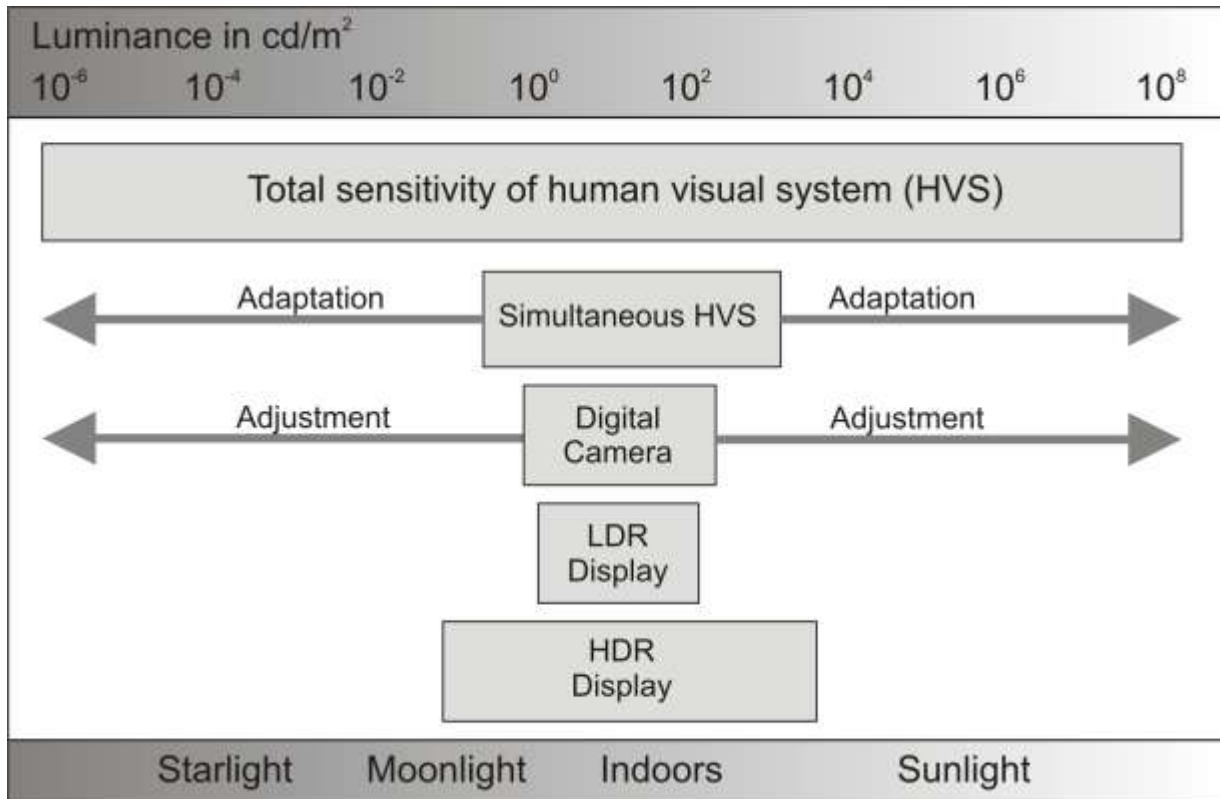


Figure 1. Comparison between the approximate sensitivity of the Human Visual System (HVS) and the typical sensitivity range of a digital camera, Low Dynamic Range (LDR) display and High Dynamic Range (HDR) display. The arrows show how the HVS can adapt its sensitivity within the range shown, while most cameras can be adjusted (by changing aperture and shutter speed settings) to at least the range shown, commonly considerably more. (Based on Kunkel and Reinhard 2010 with additions).



Figure 2. The interior of one of the rock-cut tombs at the necropolis of Cala Morel (Menorca) in which the dynamic variation between the dark, interior of the cave and the exterior presents a contrast ratio of around 60,000:1, far exceeding the available dynamic range of the camera sensor. The top six smaller images and the bottom left larger image were taken at two EV stop intervals, and none contains sufficient variation to fully represent the scene. All seven exposures were used to construct an HDR radiance map from which the final image (bottom right) could then be mapped. This image was generated using Photomatix Pro 3.0 software using the *details enhancement* method. (Photographs were taken with a Canon TS-E 17mm L using a Canon 5D MkII digital camera body, which has a very high quality 'full frame' (35mm) CMOS sensor.




Figure 3. Five digital photographs of a test pit recorded at two EV intervals (by altering shutter speed) recorded during investigations at Itchen Abbas, Hampshire in June 2010. The five images were used to generate an HDR radiance map from which image A was derived using the *tone compression* method, and image B was derived using the *detail enhancement* method. Image C was generated directly from the five source images using *exposure fusion*. All images were taken as RAW images with a Canon 5D Mk II camera, the colour calibration chart included was then used to create and apply a custom digital camera profile before the images were imported into Photomatix Pro software. Photographs were taken with a Canon 50mm 1.2L at f16 using a Canon 5D MkII digital camera body as in Figure 2. Default values were used for all other processing steps. As before, the histograms are included to illustrate the general shape of the distribution of values in each image and are not necessarily to the same scale.




Figure 4. The four images (numbered 1 to 4) were scanned from a sequence of bracketed Fujichrome colour slides. These are from excavations at the Cove, Avebury in 2003 but are fairly typical of ‘legacy’ photographic material that exists in many archaeological archives. Specific exposures were not recorded for these slides, so some experimentation with different estimates of their relative exposures was used to generate a suitable HDR radiance map. Image A (bottom left) was generated from that using the *detail enhancement* method. Image B was generated directly from the four scans without the need to construct an HDR radiance map using *exposure fusion*. Scanning was done with a Nikon 4000ED film scanner. HDR radiance map was constructed using estimated relative EV values of +5, +4, 0 and -2 respectively for images 1 through 4. Tone mapping and exposure fusion images were generated in Photomatix Pro 3.0 software using ‘default’ settings. Histograms are shown to illustrate the general distribution of values, and are not necessarily to the same scale for each image.