

Displaying digitally archived images

Kate Devlin and Alan Chalmers
University of Bristol
Bristol, UK

Erik Reinhard
University of Central Florida
Orlando, USA

Abstract

In the area of cultural heritage, images — photographic, scanned, or computer-generated — are often used as virtual representations of real artefacts or scenes. For these images to be authoritative, they should be a faithful representation of the original object. To interpret these images, they must be displayed. The conditions under which an image is displayed can adversely affect its appearance, so care must be taken to ensure that the user sees the end product in the way that it was intended to look. However, in digital image archiving, perceptual fidelity between the stored image and the displayed image is desirable, regardless of the medium of display or the environment in which it is exhibited, but this requires careful consideration of such diverse factors as tone and color reproduction, display device specifications and physical viewing conditions, which all contribute towards the final displayed image that the user perceives. This paper summarises the issues concerning display quality control for digital archiving.

1. Introduction

Virtual heritage involves the application of information technology to the field of cultural heritage. This incorporates areas such as visualisation, virtual reality, data analysis tools and dissemination of information. With the decrease in cost of mass computer storage, digital archives have become an efficient and economical way of storing information that would previously have been documented in paper form. Anything that can be stored or viewed on a computer can be archived, allowing users to have fast and easy access to a wide range of information. Digital images often form an integral part of a digital archive. High quality digital images can be used as a representation of an original source, allowing study without the need for special access, preserving details that may be lost over time, or enabling manipulation without any damage to the original.

In terms of cultural heritage, computer graphics has enabled the capturing and creation of images with the aim of generating a perceptually equivalent representation of an

original¹. Virtual reality and visualisation techniques can provide a highly detailed model of a site or artefact. Improvements in scanning and digital photography have led to the widespread use of this technology to preserve original text and art. For digital archiving to be used as a technique for representation or preservation, the integrity of an image must be vouchsafed. The user must be confident that the image they are viewing is faithful to the original — they require *perceptual fidelity*.

In digital image archiving, perceptual fidelity is desirable between the image as it was created and the resultant image that is viewed by the end-user. However, a given digital image will not always be perceived in the same way. Problems may arise because the sequence of events from image capture to perception is open to adverse influence, which can result in an image that deviates from the way it was intended to look.

A digital image generally consists of a 2D array of picture elements (pixels). The color specified by each pixel is a blend of values for red, green and blue (*RGB*). Photography, scanning, and computer graphics all produce digital images by creating a set of digital values. A number of decisions need to be made during the process of image creation. These include: the file format (the type of coding used to structure the image); image resolution (the number of pixels); the bit depth (the amount of color information); the color space (the achievable color output); and file compression (lossless, where there is no degradation of quality, or lossy, where a bigger reduction in file size may be noticeable in the quality)².

Each of the decisions mentioned above that are made during the archiving process are open to subjectivity, resulting in the possibility of variability and inexactitude. Also, as images are often displayed in different ways and in different locations from where they were created, it is necessary to ensure that steps have been taken to ensure perceptual consistency, where any point in an image will look the same regardless of changes in viewing location and display device.

To ensure that the scene as it is displayed closely re-

sembles the scene that was created or captured, it is necessary to be aware of any factors that might adversely influence the display medium. The need to exert control over image presentation has given rise to standards and guidelines concerning digital display and ergonomics, such as the guidelines of the UK's Arts and Humanities Data Service (AHDS) and the International Organization for Standardization (ISO). In addition, museums and libraries using digital images are aware of the problems of inconsistencies in image display. Reilly and Frey's report to the American Library of Congress³ highlighted the differences between images when viewed on different systems or monitors, with Library staff finding it problematic "when discussing the quality of scans with vendors over the telephone, because the two parties did not see the same image".

This paper provides a survey of the display issues affecting digital image archiving, and proposes ways of working towards perceptual fidelity.

2. Display Devices

Once the image has been created in digital form it needs to be displayed. The two most commonly encountered visual display units (VDUs) are cathode ray tubes (CRTs) and liquid crystal displays (LCDs), although the use of plasma display panels (PDPs) in museums and art galleries is becoming more popular (for example, the Smithsonian and the Guggenheim Museum). This section examines and evaluates these devices.

2.1. Cathode Ray Tubes

A color CRT uses three electron guns (referred to as 'red', 'green' and 'blue' guns) which emit an electron beam that strikes the phosphors that make up the screen^{4,5}. When a digital image is created it is stored as an array of values that represent an intensity of a particular part of that image. These values that are used to express color actually specify the voltage that will be applied to each electron gun. The values are converted from digital to analog, and video signals are produced, exciting the phosphors of the display and emitting light, which results in an image on screen.

2.2. Liquid Crystal Displays

An LCD consists of two layers of polarizing material trapping a solution which has both liquid and crystal properties; that is, the liquid crystals may be fluid, but can also retain an ordered molecular structure. When an electrical field is applied to this solution, the crystals align so that light cannot pass through. Therefore, two states are possible: either light passes through a cell or light is blocked, with each cell representing a pixel⁴. Most LCD screens

are backlit with a fluorescent light which is evenly diffused to give a uniform display.

LCDs are often used as flat panel equivalents of a standard CRT monitor, and have grown in popularity in recent years because of the decrease in volume and weight when compared with the CRTs⁶.

2.3. Plasma Display Panels

Like CRTs, plasma displays use phosphor, and like LCDs they use a grid of electrodes as pixels. They work on the same principle as a neon sign, which emits light when an electrical current is passed through gas. Plasma is a gas which is electrically conductive, and as electrons move through it they ionize the individual gas molecules. The energy gained from ionization is emitted as light during the decay process. Although the process is simple, the implementation for mass production is costly and complex.

2.4. Evaluation

Price: CRTs are inexpensive and durable when compared to other VDUs and are still in widespread use, despite the growing interest in LCD screens. As LCD technology improves, LCDs are becoming less expensive. PDPs are still the most expensive of the three.

Size: CRTs are bulky due to the fact that the electron guns need to be situated some distance from the face of the CRT tube. An LCD takes up considerably less desk space, and is more lightweight, but is limited in size when a larger screen is required. Plasma screens come in large sizes (30 inches and upwards), but this is reflected in their price.

Resolution: A CRT can display multiple video resolutions, but their refresh rate is not always high enough to remove all traces of 'flicker'⁴. An LCD monitor is limited to a 'native' resolution — the highest resolution it can display best — but without any flicker. In addition, LCDs exhibit low temporal resolution, which can be problematic when displaying dynamic images. However, LCDs give a high spatial resolution, and as each pixel is separately addressable, it provides uniformity in space and color. Further to the above, CRT pixels have a Gaussian distribution of intensity for each pixel, whereas LCDs have perfectly square pixels. This can lead to differences in appearance between LCDs and CRTs⁷.

Viewing angle: CRTs and PDPs can be viewed from any angle without the image appearing altered, and this makes PDPs a good choice for large scale displays. However, the LCD image quality is affected by angular dependence due to the optical filtering properties. This is particularly

important if the user is outside of the optimum viewing position, or if there will be multiple users.

Contrast and color: Finally, contrast ratio claims vary from manufacturer to manufacturer, but these ratios may be calculated in too many different ways to provide any meaningful comparison. In terms of color, LCDs are improving, but some older LCD monitors are only capable of hundreds or thousands of colors, compared to the greater amount of colors displayable by CRTs and PDPs.

3. Image Display

In the previous section, the characteristics of various different display devices were discussed. CRTs, LCDs and plasma screens all have different limitations in terms of how and with what quality images are displayed. These limitations are not benign — display devices alter images in various perceptually significant ways. In this section, we discuss some of the corrections that may be applied to images before they are displayed, so that their visual quality is minimally affected by the chosen display device. In particular, in the following subsections we show the importance of gamma correction, tone reproduction and gamut mapping, all of which are specific image treatments that make an image suitable for display.

3.1. Gamma Correction

The mapping between the intensity values of the created image and the value subsequently emitted by each of the pixels would ideally be linear, so that the input to the display matches the output. However, for CRTs the mapping is normally not linear, but can be approximated with a power law

$$L_d = L_{max} L^\gamma \quad (1)$$

where L_d is the displayed intensity, L_{max} is the maximum displayable intensity, L is the input value between 0 and 1 and γ is a first-order approximation of the display's non-linearity (see for example ⁸). Different brands of computer deal with gamma correction in different ways, resulting in typical values for Macintosh computers of 1.8 and for Silicon Graphics machines of 1.5. PCs do not have gamma correction in hardware, and therefore the gamma for PCs depends on the monitor used, with typical values in the range of 2.0 to 2.6 ⁷.

For LCDs and plasma screens, the story is more complicated. Some LCD monitors have a built-in artificial non-linearity to mimic CRT devices. Others do not have such hardware added and may have other unknown non-linear responses to input signals.

It is possible to estimate the gamma of a computer system by displaying an image which consists of a set of grey

values next to an area of alternating black and white scan-lines. Seen from a distance, the black-and-white pattern fuses to appear grey and the grey patch which matches the fused pattern best, is selected. The intensity of this patch is used to find the gamma response of the CRT. Once the gamma for a particular set-up is known, images to be displayed may be corrected with the following transformation which is generally known as gamma correction:

$$L' = L^{1/\gamma} \quad (2)$$

3.2. Tone Reproduction

The range of intensities witnessed in the real world is vast, from the darkness of a night sky to a bright, sunny day. Current state-of-the-art image capturing techniques allow much of this range to be recorded in high dynamic range images by combining and amalgamating the various exposure levels ⁹. The range of intensities (the *dynamic range*) of such images is much larger than can be displayed on current display hardware. However, for virtual heritage applications, capturing imagery in high dynamic range format is desirable, because in the future high dynamic range display devices will become available allowing this data to be displayed directly ¹⁰. By capturing and storing as much of the real scene as possible, and only reducing the data to a displayable form just before display, the archive becomes future-proof.

As high dynamic range display devices are currently in the experimental stage, the display of high dynamic range data requires an extra step to reduce the range of intensities of the image to be within the range of intensities displayable by current display devices. This step is called *tone reproduction*, or *tone mapping*, and involves scaling the large intensity values in the input down to the displayable range.

A straightforward linear scaling between the original high dynamic range data and the display is not the best solution as many (if not all) details can be lost (Figure 1). The mapping must be tailored in some non-linear way.

A number of tone reproduction operators have been presented, each with its own visual characteristics ^{11,12}. Some of the operators are concerned with achieving perceptual fidelity with a real-world scene, and mimic aspects of the human visual system (HVS). Others concentrate on producing a subjective best image that is pleasing to the eye. With a large number of operators available, and validation of tone reproduction operators in its infancy, the choice of tone reproduction operator is currently a matter of deciding on the best tool for the job. Currently, there are no defined criteria for selecting the best tone reproduction operator for a specific task.




Figure 1: Linear scaling of HDR data will cause almost all details to be lost (top image). Here, the lightbulb is mapped to a few white pixels and the remainder of the image is black. Tone reproduction operators recover detail in both light and dark areas as well as all areas in between (bottom image).

3.3. Gamut Mapping

While the previous section deals with the range of image intensities that can be displayed, display devices are also limited in the range of colors that may be displayed. The term gamut is used to indicate the range of colors that the human visual system can detect, or display devices can display.

Even with 24-bit color, sometimes indicated as ‘millions of colors’ or ‘true color’, there are many colors within the visible spectrum that monitors cannot reproduce. To show the extent of this limitation for particular display devices, chromaticity diagrams are often used. Here, the Yxy color space is used, where Y is a luminance channel (which ranges from black to white via all greys), and x and y are two chromatic channels representing all col-

ors. Figure 2 shows a chromaticity diagram indicating the gamut of colors visible to humans (‘gamut of all colours’), and two restricted gamuts, one for a typical monitor and one for a printing device. Given that the triangle is completely contained within the shape of all perceptible colors, there are many visible colors that cannot be reproduced on a monitor. Assuming that some of the colors available in

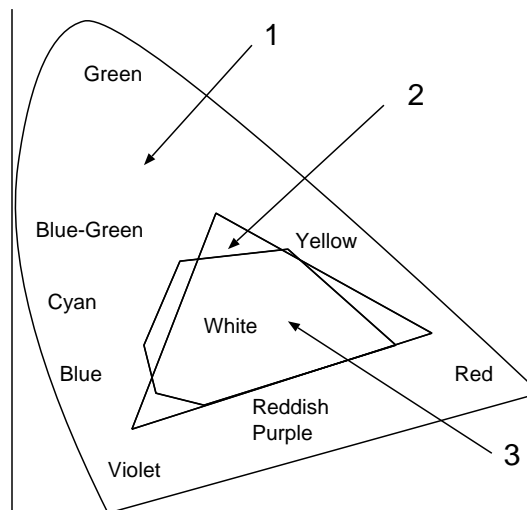


Figure 2: Example of a CIE Chromaticity diagram showing the range of colors that humans can detect (1), as well as the ranges of colors displayable on a monitor (2), and on a printer (3).

an image to be displayed are outside a monitor’s gamut, the image’s colors may be remapped to bring all its colors within displayable range. This process is referred to as gamut mapping^{13,14}. A simple mapping would only map out-of-range colors directly inward towards the monitor’s triangular gamut. Such a ‘colorimetric’ correction produces visible artefacts.

A better solution would be to re-map the whole gamut of an image to the monitor’s gamut, thus remapping all colors in an image. This ‘perceptual’ or ‘photometric’ correction may avoid the above artefacts, but on the other hand there are many different ways in which such remapping may be accomplished. As such, there is no standard way to map one gamut into another more constrained gamut.

4. Display environment

VDUs are open to influences from the environment in which they are located. The observers’ state of adaptation when they are viewing a displayed image may not match with their adaptation if they were viewing the real scene. In addition to this, the glass surface of display devices may reflect any light in the viewing environment.

4.1. Color appearance models

Unless large screen or immersive displays are being used, the VDU does not normally fill the whole field of vision. This means that the light present in the room, as well as the light emitted by the screen, will have an effect on the viewer's state of adaptation. Color appearance models can be used to try to predict the appearance of color in particular environments so that it is perceived as consistent across different environments¹⁵.

4.2. ICC Profiles

The International Color Consortium has specified a standard for the interchange of images between different display media¹⁶. It uses a device independent Profile Connection Space (PCS) for color management purposes. Transformations between specific native device color spaces and PCS is guided by ICC color profiles which effectively characterize display devices. Use of ICC color profiles is highly recommended if perceptual consistency is required between displayed material on different display devices.

4.3. Reflected ambient light

Ambient light can also cause a reduction in the perceived contrast of a displayed image. This is due to extraneous light in the viewing environment being reflected off the screen, causing an image viewed under bright ambient light to appear 'washed out'^{4,5,17}. There are two types of reflection, *specular* (or mirror-like) and *diffuse*. Specular reflections occur when light emitted or reflected by objects form images on the glass of a display screen. Diffuse reflections cause a uniform increase in luminance across all points of a display screen¹⁸. CRT technology is particularly prone to reflection off the screen. LCDs also suffer, but as they are more mobile they can often be easily moved, whereas a CRT's bulk means it cannot¹⁹.

Standards such as the ISO3664 Viewing conditions for Graphic technology and photography²⁰ outline a wide range of factors that should be addressed to achieve the best possible viewing environment, thus maintaining optimum perceptual fidelity. These guidelines consider the position of the screen, the illuminant and the observer; the spectral conditions for the reference illuminant; and the monitoring of apparatus for maintenance and degradation. Current approaches to this problem involve measuring the ambient illumination with specialised hardware, and altering the display device or changing the viewing conditions.

5. Discussion

Nowadays, digital archives for cultural heritage are used worldwide. Many of these archives are available online,

so image fidelity is of importance where Internet users are located worldwide, using different systems and in different environments. The wide availability of guidelines and standards pertaining to digital image creation are the best starting point in striving for perceptual fidelity. These guidelines provide a sound basis for the first part of the image archive process — the creation of the images and all the factors this entails — as described in Section 1. The actual delivery of the image should take into account the issues detailed in the previous sections.

Some institutions do address these factors. The Bodleian Library's online image catalogue at the University of Oxford states: "Note that the apparent quality of the images as viewed on-screen is in part dependent upon the quality of the monitor used to view them, and the apparent color-values are likewise dependent on whether the monitor has been correctly calibrated, and on the ambient lighting conditions of the room."²¹.

Suggestions as to the optimal display of archived images follow.

For the display of images in a controllable environment (for example, a museum display or art gallery):

- Choice of the most suitable display device based on space, cost, number of viewers.
- Correct display calibration.
- Choice of an appropriate tone and gamut mapping.
- Minimal impact from the surrounding viewing environment.

For the display of images over a network:

- Advice as to optimal display conditions.
- Information on image creation (for example, what gamma correction, if any, was applied). This can be addressed through metadata provision.

It must be remembered that future users of digital image databases may have technology and/or techniques beyond what we can hope to achieve today, but the information we have preserved must be available for them to use in the best format possible. By striving for perceptual fidelity, digitally archived images can become a useful part of cultural heritage, where an image can be used with more faith in its integrity.

References

- [1] J.A. Ferwerda. Three varieties of realism in computer graphics. In *Proceedings SPIE Human Vision and Electronic Imaging '03*, 2003.

- [2] C. Grout, P. Purdy, J. Rymer, K. Youngs, J. Williams, A. Lock, and D. Brickley. *Creating and Using Virtual Reality: a Guide for the Arts and Humanities*. Visual Arts Data Service. Oxbow Books, 2000.
- [3] J.S. Reilly and F.S. Frey. Recommendations for the evaluation of digital images produced from photographic, microphotographic, and various paper formats. Report to the Library of Congress, May 1996.
- [4] D. Travis. *Effective Color Displays*. Academic Press Ltd., 1991.
- [5] A.S. Glassner. *Principles of Digital Image Synthesis*, volume 1. Morgan Kaufman, San Francisco, CA, 1995.
- [6] M. Menozzi, U. Napflin, and H. Krueger. CRT versus LCD: A pilot study on visual performance and suitability of two display technologies for use in office work. *Displays*, 20:3–10, 1999.
- [7] C. Poynton. *Digital Video and HDTV: Algorithms and Interfaces*. Morgan Kaufmann Publishers, 2003.
- [8] P. Shirley. *Fundamentals of Computer Graphics*. A K Peters, Natick, Massachusetts, 2002.
- [9] P.E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 369–378, August 1997.
- [10] H. Seetzen, L. Whitehead, and G. Ward. High dynamic range display using low and high resolution modulators. In *The Society for Information Display International Symposium*, 2003.
- [11] K. Devlin. A review of tone reproduction techniques. Technical Report CSTR-02-005, Computer Science, University of Bristol, 2002.
- [12] J. DiCarlo and B. Wandell. Rendering high dynamic range images. In *Proceedings of the SPIE Electronic Imaging 2000 conference*, volume 3965, pages 392–401, 2000.
- [13] G.J. Braun and M.D. Fairchild. General-purpose gamut-mapping algorithms: Evaluation of contrast-preserving rescaling functions for color gamut mapping. In *IS&T/SID Seventh Color Imaging Conference: Color Science, Systems and Applications*, pages 167–172, 1999.
- [14] R.S. Gentile, E. Walowit, and J.P. Allebach. A comparison of techniques for color gamut mismatch compensation. *J. Imaging Technology*, 16:176–181, 1990.
- [15] M D Fairchild. *Color appearance models*. Addison-Wesley, Reading, MA, 1998.
- [16] ICC (International Color Consortium). *Specification ICC.1:2003-09, File Format for Color Profiles*, 2003.
- [17] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann, 2000.
- [18] D.C. Rogers, R.E. Johnston, and S.M. Pizer. Effect of ambient light on electronically displayed medical images as measured by luminance-discrimination thresholds. *Journal of the Optical Society of America A*, 4(5):976–983, 1987.
- [19] B. Schenkman, T. Fukuda, and B. Persson. Glare from monitors measured with subjective scales and eye movements. *Displays*, 20:11–21, 1999.
- [20] ISO (International Standards Organisation). *ISO3664 Viewing conditions – Graphic technology and photography*, second edition, 2000.
- [21] Bodleian Library, University of Oxford. <http://www.bodley.ox.ac.uk/>.

Biographies

Kate Devlin has a B.A. Hons. in Archaeology and an M.Sc. in Computer Science from Queen's University, Belfast. She is currently finishing her Ph.D. in computer graphics at the University of Bristol, and also works on digital image archiving for the archaeology department there. Her interest is in the perceptually accurate display of images, and its application to the field of archaeology.

Alan Chalmers is a Reader in the Department of Computer Science at the University of Bristol, UK. He has published over 80 papers in journals and international conferences on very realistic graphics. He was Vice President of ACM SIGGRAPH from 1999–2002. His research investigates the use of very realistic graphics in the accurate visualisation of archaeological site reconstructions and techniques which may be used to reduce computation times without affecting the perceptual quality of the images.

Erik Reinhard is an assistant professor at the University of Central Florida and has an interest in the fields of visual perception and parallel graphics. He has a B.Sc. and a TWAIO diploma in computer science from Delft University of Technology and a Ph.D. in computer science from the University of Bristol. He was a post-doctoral researcher at the University of Utah. He is founder and co-editor-in-chief of ACM Transactions on Applied Perception.