A Simple Procedure for the Restoration of Shadows and Highlights in Digital Prints

Rodney Shaw
Hewlett Packard Research Laboratories
Palo Alto, California

Abstract

The optimum reproduction of images with extended latitude remains a challenging topic in both digital display and digital printing. Here a simple procedure is described for the restoration and enhancement of shadows and highlights in digital images where significant detail would otherwise be omitted in these regions. This procedure is based on a two-stage image-enhancement algorithm that is simple in concept yet designed according to rigorous visual-science criteria. An outline is presented of how these criteria influence the design of the two stages, and details are given of the practical nature of both these stages for the enhancement of typical extended-latitude digital images.

Introduction

Digital printing and display technologies suffer from similar limitations to optimization of perceived image quality as those for traditional (analog) technologies. In fact the problems posed by essential tone, color, resolution and noise requirements are augmented by the need to combat the artifacts introduced by the underlying grid-like structure of the image. In spite of this disadvantage, the newer digital technologies have an inherent advantage, in that they allow accessibility to the application of sophis-ticated image-processing techniques. Such techniques have been pioneered over recent decades within the image processing community, and according to a wide variety of performance criteria, from machine detection and recog-nition of complex signal-types, to the visual inspection of aerial imagery. It is well-known that the best of these image processing techniques, when appropriately applied, can have significant influence on most aspects of perceived pictorial image quality, from straightforward augmentation of tone reproduction to the more sophisticated nuances associated with sharpness filters and noise suppression techniques.

At recent conferences in this series the author has described the development of image quality metrics for digital photography and digital printing, and has provided details of absolute scales for digital sharpness and digital noise.[1,2] The approach taken in the development of both of these was based on a simple Fourier description of the visual process, and a spatial frequency integration of this visual function when combined with the relevant signal and noise spectra associated with the printed image. It has been an explicit thesis of these studies that basic knowledge of these essential spatial-frequency character-istics which are associated with the perceived impression of high-quality printed images can lead to significant clues as to where and how new digital enhancement techniques may be relevant.

Latterly, these same assumptions have been used to approach a slightly different problem,[3,4] namely the enhance-ment and optimization of digital images. Here the underlying thesis is the enhancement of the band of spatial frequencies that are crucial for the visual impression of quality at the expense of those which are only of marginal relevance. Whereas this technique is applicable to digital images in general, by nature it is ideally suited to those images which encompass a very wide dynamic range, since in these cases the problem of bringing the extremes into satisfactory visual range is also in effect one of deciding which spatial-frequency regions (*slow gradients)* of the image can, so to speak, be discarded without injurious effect on perceived image quality, and therefore accommodating more of the important spatial-frequency regions (*fast edges*). To do so in a systematic manner naturally involves a detailed quantitative analysis of these spatial frequency domains.

Spatial Frequency Analysis

Figure 1 shows the assumed typical visual transfer function (VTF) associated with normal print viewing. This has been used throughout the development of the absolute image quality scales, and has a long history as a surrogate for human vision in the design of imaging technologies and systems.[5] The surrogate for the signal, or scene, spectrum is obtained by assuming a flat spectrum over the band-pass of the observer, and the product of this with the VTF is shown in Figure 2. This product is for convenience referred to as the Visual-Detail Transfer Function (VDTF), and it is recalled that the curve of Figure 2 is identical to the radial integral of Figure 1, assumed circular symmetric (in a fixed image viewing-element the number of cycles of any given spatial frequency is directly proportional to the spatial frequency itself).

If the total area under VDTF is a measure of the ulti-mate sharpness when there are no other limitations other than those imposed by the observer, then the sharpness associated with any intervening imaging element (senor pixels, print resolution interval, etc) can readily be calculated in terms of its influence on this area. With this assumption the author derived an absolute digital sharpness scale from 0 to 10, and has subsequently used it to evaluate various digital imaging process, the number on this scale being designated the Sharpness Index.[2]

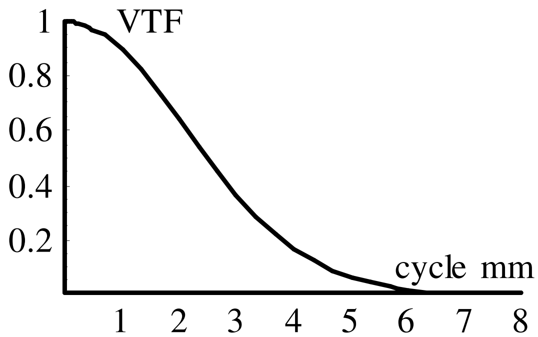

Figure 1. The Visual Transfer Function as a function of spatial frequency

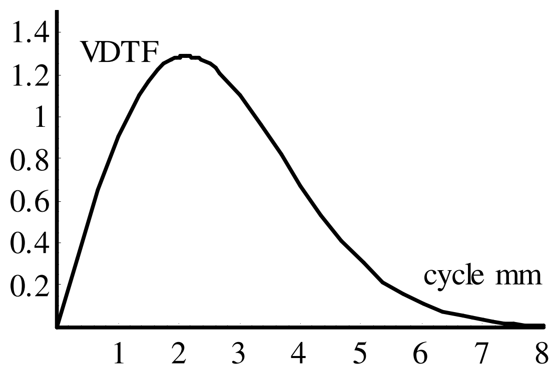

Figure 2. The Visual Detail Transfer Function

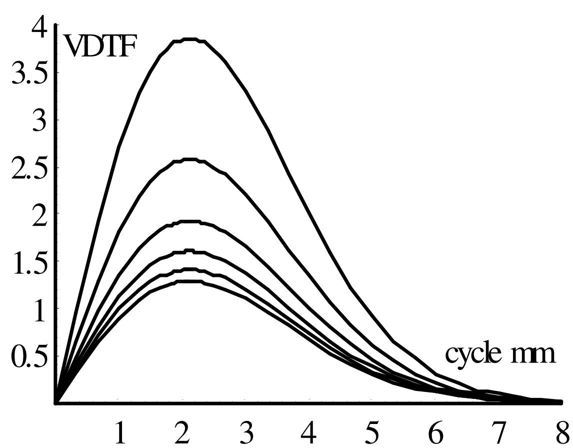

Figure 3. The Visual Detail Transfer Function, and as enhanced by a variable strength amplifier

These spatial frequency considerations give rise to the possibility of designing a simple variable sharpening algorithm for the enhancement of digital images. Figure 3 shows the goal of such a filter in terms of the desired influence on the Visual Detail Transfer Function, where is has been assumed that a conceptual variable-strength linear amplifier has operated on this function. The author[3] has previously described how this concept may be translated into a practical class of convolution filters. However here the practical problem is not necessarily one *per se* of remedying unsatisfactory image sharpness, but the closely related where selective enhancement is desired in the image shadows and highlights, and a slightly more complex approach is required, namely the use of a two-stage enhancement process.

Print/Display Latitude Extension

A common class of image-quality problem encountered in digital imaging relates to the accommodation of scenes having a high dynamic range of brightness levels, especially so with the increasing use of digital acquisition devices with technologies capable of extended-range scene-capture. The associated problem is often found in natural scenes where the preponderance of the natural print-detail falls within the highlight or shadow area of the print, and when the problem cannot be rectified by global mean-level amplification techniques (ie, the classical analog approach of tone-manipulation). One useful approach to this problem follows naturally as a corollary to that used above in constructing a sharpening filter. However the problem is now slightly more complex, and calls for a conditional type of spatial filtering and more specifically, selective local enhancement within defined spatial frequency bands.

It has been found useful to view this problem as a two-stage spatial-filtering problem, the first involving removal of very low spatial frequencies (*the slow gradients*) which have little on the perception of detail, and by doing so re-mapping the extreme regions towards the intermediate brightness regions, allowing a second-stage re-stretching of the important spatial frequencies that convey image detail (*the fast edges).* In this way all the existing boundary conditions across existing boundaries in the image can be maintained, while, in effect, '*providing image space*' for edges at the extremes of dynamic latitude which were otherwise absent or under-represented in the original perceived image. In this way the enhancement of the important higher spatial frequencies more than offsets discarding the very low frequencies, thus conveying the desirable visual impression of an overall enhancement.

In summary, the role of the first stage is thus essentially to define those low spatial frequencies which are of limited visual significance, and in effect to exempt them from the subsequent amplification of the higher frequencies which are crucial to the visual impression of image quality. In this way the highlights and shadows may be displayed within the normal perceptual range.

Figure 4 gives a practical illustration of the nature of this first stage, for simplicity assumed here to be a pseudo-circular digital convolution filter. The resulting Bessel function (denoted here as the Filter transfer Function, FTF) is expressed specifically with the assumption of effective viewing the digital image at 300ppi, and for different pseudo-diameters as expressed in terms of the number of image pixels.

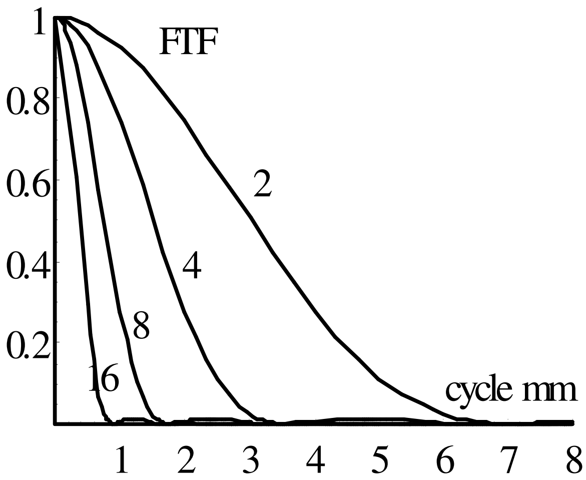

Figure 4. Filter Transfer Function of the first stage, in terms of the diameter expressed as number of image pixels.

Inspection of Figure 4 shows that even for relatively small pre-filter sizes, this first stage thus acts simply as a low-pass filter, and for example somewhere in the region of 8 to 16 pixels radius, only frequencies below around 1 cycle per mm are maintained in the ensuing blurred image. This first stage therefore may be used effectively to define those frequencies referred to above as *slow gradients*, and the influence on the VDTF can be expressed by the Filter Visual Detail Transfer Function (FVDTF).

Figure 5 shows the FVDTF corresponding to the pixel-sizes of the low-pass filter, and from this plot it may reasonably be concluded that for a 16-pixel diameter filter the contribution to the overall perception of image sharpness is only a small fraction of the contribution made over the entire visual range. In other words, Figure 5 provides a quantitative means of defining what has been previously been referred to as the contribution of the *slow gradients.*

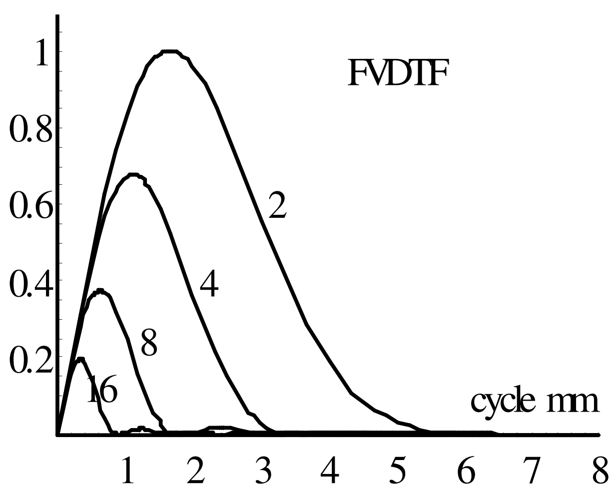

Figure 5. Filter Visual Detail Transfer Functions for the transfer functions shown in Figure 4.

The second stage then involves only a straightforward amplification of the visually-important spatial frequencies not implicitly filtered out by the first stage. This involves the decision of the degree of amplification to be made based only the local pixel value of the blurred function following the first stage, but then applied to the actual pixel value in the original image. Since this technique is based on enhancing local image brightness relationships across boundaries, the operation can be carried out in gray-scale, or in any proper color coordinate system, as appropriate.

Whereas any candidate function can be used for the second stage amplification and re-mapping, experience has shown that applying any of a reasonable set of smooth, single-valued amplification functions can lead to satisfactory results. Again, visual science offers clues to the optimum of such functions, but this is a lengthy and even controversial topic, and is left to a later date. But before leaving this topic we can readily illustrate the outcome when, say, the local pixel amplification is by a factor two. For this specific example, based on an 8-pixel diameter first stage filter, Figure 6 shows both the low-frequency region removed by the first stage, and the before- and-after VDTF relative to the second stage. It is readily seen the increased area under the curve due to the secondary amplification greatly outweighs any loss due to the first stage (the area under the FVDTF curve).

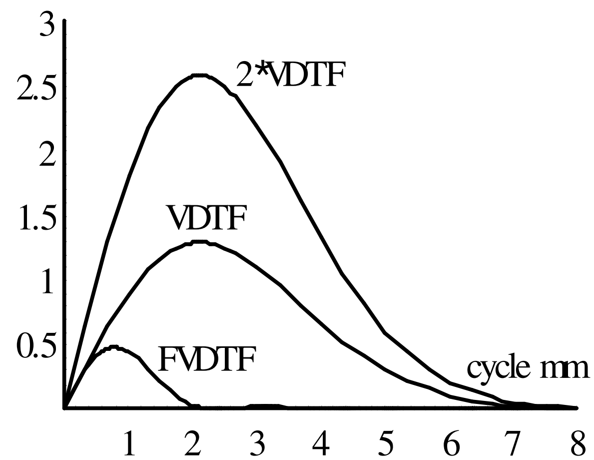

Figure 6. Visual Detail Transfer Functions for first-stage filter and second-stage amplification.

Figure 7 gives a pictorial example of typical results obtained by this approach. This shows a section of a digital scene wherein shadows and highlights dominate, and the main subject matter of interest contains both. Below this is an illustration of the results of conventional tone manipulation. The subject matter is now clearly visible overall, but at the expense of a washing out of the natural highlights. Finally, at bottom, the result is shown of a two-stage spatial-filtering technique, that as described above, selectively suppresses low spatial frequencies and enhances higher frequencies. It is seen the detail is now reproduced in a satisfactory matter everywhere, including both shadows and highlights. Again, this was made possible by effectively filtering out very low frequencies which themselves contribute little to the perception of detail.

Figure 7. Illustration of a section of a digital print involving strong shadows and highlights (top); as tone-manipulated (middle); and as adaptively spatial-filtered (below).

As in the case of the edge-enhancement filter, an attractive feature of this adaptive latitude-extension filter is its continuous nature, as illustrated in Figure 8. A common problem encountered in printing digital or analog photographs is posed by originals in which the use of flash has 'washed-out' the intended scene highlights, often in the form of facial features. Figure 8 shows an example where variable degrees of adaptive enhancement of the highlight details has followed conventional global tone-manipulation. This small but key element of the image may be enhanced to user preference, but not at the expense of detail in the general surrounding areas of the image. In Figure 8 the degree of enhancement has been taken beyond reasonable user choice in order to illustrate the continuity and extent of the operation made possible by the prudent choice of low-spatial-frequency band-with, and starting from same fundamental image-quality considerations.

Figure 8. Illustration of various degrees of restoration for scene highlights using an adaptive enhancement filter

References

1. R. Shaw, *IS&T Procs NIP 12*, 162 (1996).
2. R. Shaw, *IS&T Procs NIP 15*, 446 (1999).
3. R. Shaw, *IS&T Procs NIP 17*, 705 (2001).
4. R. Shaw, *IS&T Procs PICS*, Portland 221 (2002).
5. R .Dooley & R. Shaw, *J. Appl. Phot. Eng*., **5**, 190 (1979).

Biography

The author received a PhD in physics from Cambridge University. After several research and teaching positions in the UK and Europe he came to the USA in 1973, and following research appointments at Xerox and Eastman Kodak was Director of the Center for Imaging Science at RIT. He joined H-P Labs in 1994, and his current interests are in image processing and digital systems modeling. *rodshaw@hpl.hp.com*