Image Quality, Spatial Bandwidth, and Design Criteria for Optimum Digital Enhancement Techniques

Rodney Shaw
Hewlett Packard Research Laboratories
Palo Alto, California

Abstract

While image evaluation has evolved to a mature science in terms of both machine-measured and visually based metrics, its application to the systematic design and development of digital-image enhancement techniques has lagged. This is unfortunate, since the explicit end goal of the majority of these techniques is the optimization of the final image quality, and especially in consumer pictorial applications the quest for increased image remains an important one.

The spatial-frequency-domain analysis of the most important image quality attributes is also now widely accepted, especially since those human visual attributes that are important in engineering sophisticated image-capture and printing devices can also be most conveniently be expressed in this domain. Therefore the spatial-frequency domain, and the associated concept of the selective manipulation of different frequency bands, are tailor-made for the design of optimum image enhancement techniques.

This study investigates the spatial-bandwidth concepts which are the basis of vision-based image quality metrics, and illustrates their application to robust but computationally-simple image enhancement techniques.

Introduction

Digital printing technologies offer an obvious advantage when compared with the traditional (analog) technologies, in that they allow accessibility to the application of sophisticated image-processing techniques for the goal of achieving of optimum image performance, according to whatever criteria are most critical. Such techniques have been pioneered over the past few decades within the image processing community, and cover a wide variety of performance criteria, from machine detection and recognition of complex signal-types to the visual inspection of pictorial imagery. It is well-known that the appropriate sub-set of these techniques, when properly applied, can influence most aspects of perceived pictorial image performance, from straightforward augmentation of tone reproduction to the more sophisticated requirements of sharpness filters and noise suppression techniques.

In spite of having this well-established image-processing background to draw from, the consumer digital-photography/digital print community has been somewhat slow to capitalize on this potential competitive advantage over analog imaging systems, and in the main, only ad hoc techniques have actually been incorporated, and on a fairly scattered basis. The reasons for this are many, and range from unfamiliarity of what has already been achieved in other imaging technologies, to more esoteric questions concerning the stage at which these techniques should reside within an end-to-end digital imaging system (including acquisition and print stages).

At recent conferences in this series the author has described image quality metrics for digital photography and digital printing, and has detailed 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 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 characteristics 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.

In this present study the nature of these crucial spatial frequency domains is explored in more detail, with especial concentration on the visual-descriptor basis, and the concept of simultaneously enhancing and suppressing different spatial frequency regions for the purpose of optimizing the perceived impression of image quality.

The Spatial Frequency Domain

Figure1 shows the assumed typical visual transfer function (VTF) associated with normal 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 systems.

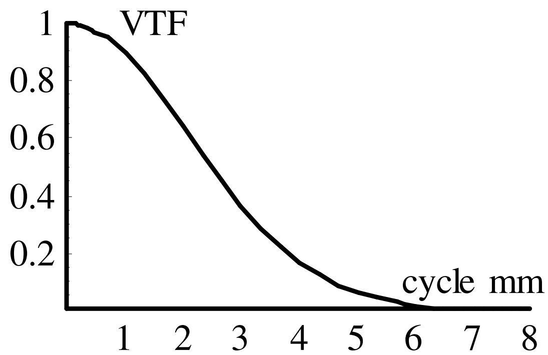

Figure 1. The Visual Transfer Function as a function of spatial frequency

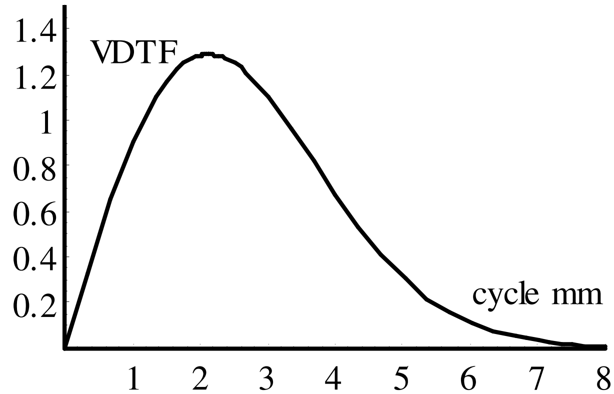

Figure 2. The Visual Detail Transfer Function

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 ultimate 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.

In digital photography the influence of the independent resolution element of the sensor plays an important role in defining the image sharpness, and its influence can readily be calculated in these terms.

Figure 3 shows a calculation of the combined influence of the number of sensor pixels on a side and the size of the resulting digital output (display or print). On this scale the range of around 8 to 10 corresponds to the practical range of acceptable analog photography. Systems design data in this form allows an at-a-glance understanding of the sensor pixel requirements in terms of the desired image sharpness. For example, in the absence of sharpening techniques, a sharpness level of 9 associated with a 10" print requires a nominal sensor array of 300 pixels on a side.

Similarly, in digital printing the resolution element of the device, or *ppi*, imposes fundamental limits as summarized in Figure 4. It is seen that there are effectively three regions of this curve. Up to around 200ppi there is an approximate linear increase in sharpness with print ppi, while from around 200 to 400ppi the increase is at a slower rate, while beyond 400ppi there is only a very gradual increase.

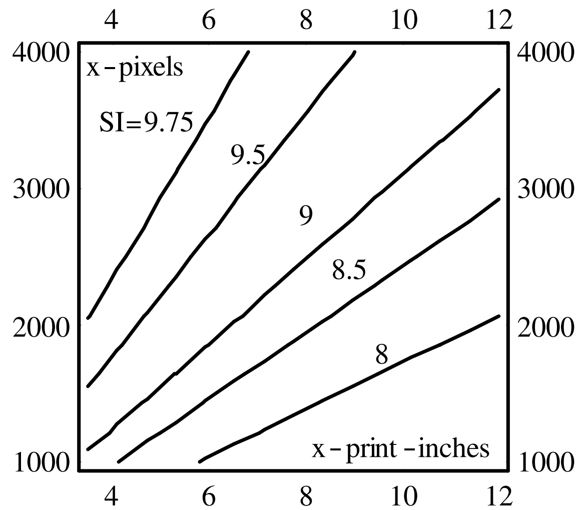

Figure 3. The Sharpness Index for a digital acquisition device in terms of number of sensor elements and the final print size.

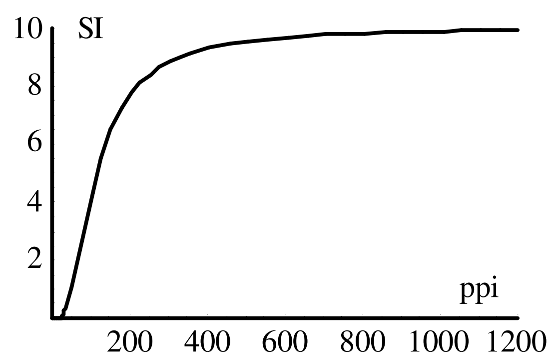

Figure 4. The Sharpness Index for a digital printing device in terms of number of sensor elements and the final print size.

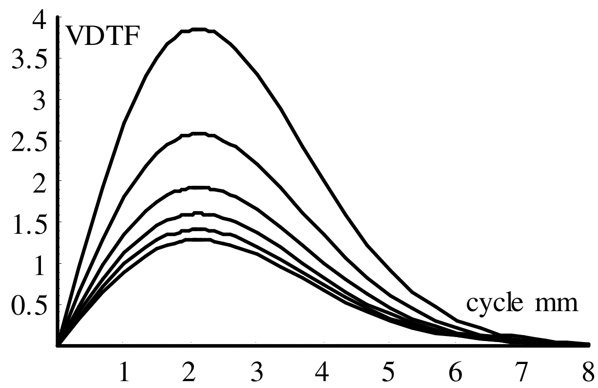

Figure 5. The Visual Detail Transfer Function, and as enhanced by a sharpening filter of variable strength

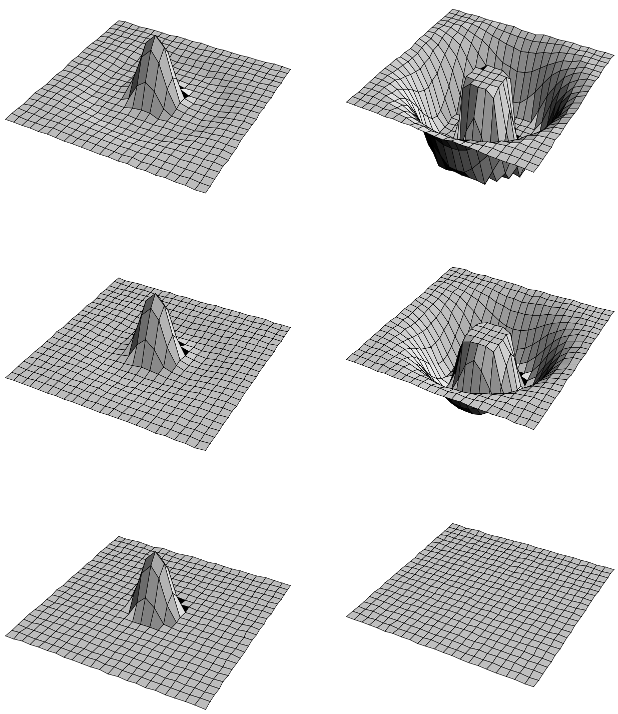

Figure 6. Examples (left) of sharpening (convolution) filters of different strength: The negative lobe (right) is also shown, as amplified for illustration.

Sharpness Enhancement Techniques

These spatial frequency considerations give rise to the possibility of designing a simple variable sharpening algorithm for the enhancement of digital images. Figure 5 shows the goal of such a filter in terms of the desired influence on the Visual Detail Transfer Function, leading to as set of convolution filters as illustrated in Figure 6.

Figure 7 shows a practical example of the successful application of this sharpening techniques. This is an enlarged section of a digital print considered to be of satisfactory quality other than in the attribute of sharpness. Starting with the original un-enhanced version (top), the results of successive increases in the amplitude of the filter are shown below. In this particular case the picture on the scale as printed was originally judged to be outside the generally acceptable photographic range, but after enhancement was judged to be well within this range, and this was confirmed by calculation of the changes on the associated digital sharpness scale.

Of course, in general the limits to which such enhancement may be made will be governed by personal preference, as well as by the inherent noise level and its own spatial spectrum, since this will also be enhanced unless other noise-suppression techniques are used in unison. However such considerations are beyond the present scope.

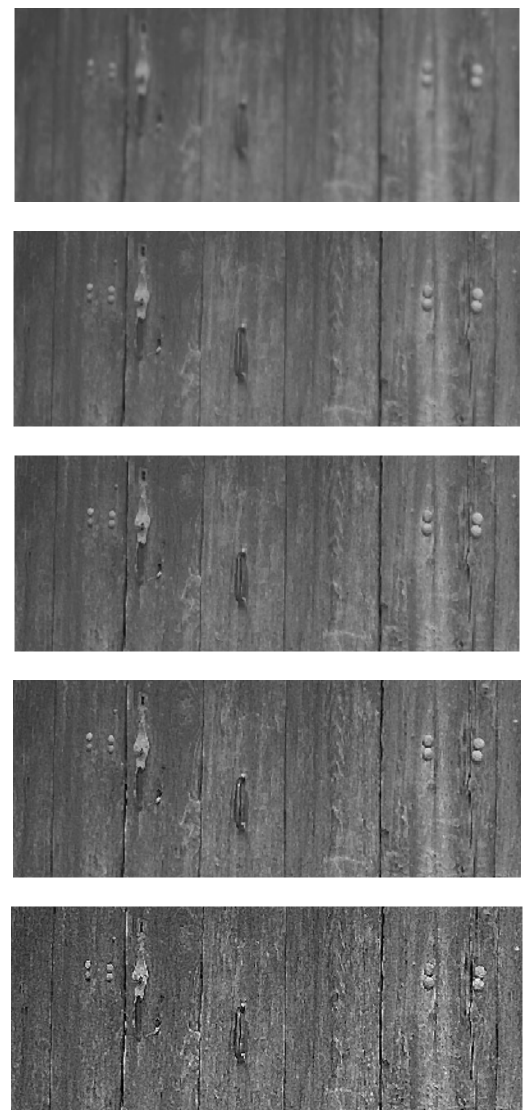

Figure 7. Enlarged section of digital print (top) and as edge - enhanced by a variable digital sharpening filter.

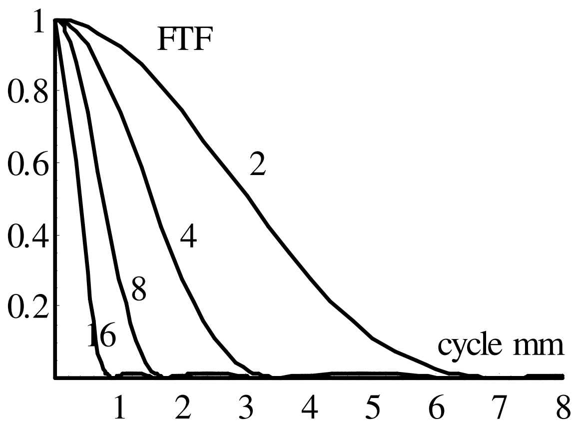

Figure 8. Transfer Function of first stage filter for a series of diameters in terms of number of image pixel.

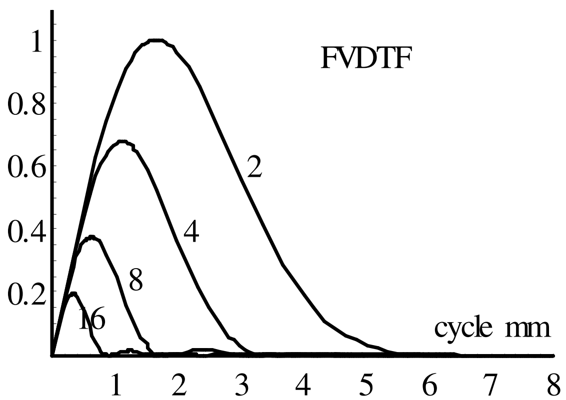

Figure 9. Filter Visual Detail Transfer Functions for the transfer functions shown in Figure 8.

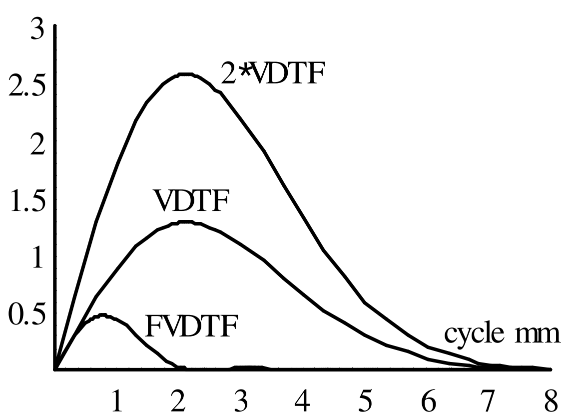

Figure 10. Filter Visual Detail Transfer Functions for the transfer functions shown in Figure 8.

Print/Display Latitude Extension

Another common class of image-quality problem that is 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 more complex, and calls for an adaptive form 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 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. In this way all the existing boundary conditions across existing edges in the image can be maintained, while, in effect, 'making room' for edges at the latitude extremes which were otherwise absent or under-represented in the perceived image. In this way the enhancement of the important higher spatial frequencies more than offsets discarding the very low frequencies, thus conveying the visual impression of increased sharpening overall, as well as the obvious advantage of viewing detail in the latitude extremes not otherwise above the visual threshold.

Figure 8 shows the (Bessel function) form of the low-spatial-frequency filter used at the first stage, designed specifically for the assumption of effective viewing of the digital image at 300ppi, and for different diameters as expressed in terms of image pixels. Figure 9 shows the equivalent FVDTF for each of these filters. Figure 10 shows the 16-pixel filter, along with the higher spatial frequencies VDTF that are preferred for enhancement, say by a factor 2, as also shown in Figure 10. In other words, conditional pixel re-mapping which effectively enhances the spatial frequencies which are most important for image detail is at the potential expense of these lower and less important spatial frequencies.

Figure 11 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 11. 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 12. 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 12 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 12 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, starting from fundamental image-quality considerations.

Figure 12. Illustration of various degrees of restoration for scene highlights using an adaptive enhancement filter

Conclusion

It has been demonstrate how the spatial frequency range associated with optimum image quality may be used as a basis for the design of visual image processing techniques for the purpose of the enhancement of desirable image features. Examples have been given in terms of a variable edge enhancement filter matched to the spatial frequency bandwidth of the print as viewed, and to the extension of perceived print latitude when otherwise detail would be lost in the shadows or the highlights. In the latitude-extension case these important considerations lead to a technique which adapts naturally to the local image detail.

References

1. R. Shaw, IS&T Procs. PICS 2000, pg 71.
2. R. Shaw, IS&T Procs. PICS 2001, pg 124.

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