Image-processing for the digital consumer market

Rodney Shaw & Paul Johnson; White Rose Digital; Aptos, California

Abstract

Billions of pictures go unprinted or cause consumer dissatisfaction due to their poor image quality and the inability of users to make simple enhancements that would render them satisfactory. This conclusion is true for applications such as digital photography, where high quality prints are desirable, and for digital documents where it would often be convenient to carry out in situ image processing of embedded pictures. By combining both traditional and non-traditional imaging tools, a simple yet highly effective overall image-enhancement methodology has been developed that is capable of placement everywhere a user interacts with a digital image, including a desk-top printing menu. The authors will describe the basic imaging principles used in the development of this practical methodology, from initial concept through to end-solution, and a demonstration will be given of typical user operation.

Introduction

The dramatic growth and pervasive influence of image-wise communication is changing society on a daily basis, and the digital image is now an adjunct of everyday consumer life. While the internet and the many digital services it now provides have been the main drivers of this revolution, the vast and growing numbers of sophisticated digital cameras, printers, video-phones, etc, that are currently in the hands of ordinary customers have also been key contributors in this democratization of the entire consumer field.

These rapid and profound changes associated with digital imaging call for new digital solutions to decades-old imaging problems, as billions of digital images of all descriptions now circulate freely at the frontiers of modern communications. These digital solutions should empower the consumer to have control over all aspects of their digital images throughout the complete usage cycle, from acquisition to viewing, sharing, printing and storing, and this control should especially include the ability to render each image to the maximum of the user preferred quality. The latter implies comprehensive and user-friendly digital image-enhancement facilities not currently available to the consumer.

The general field of digital mage processing has in fact a long and well-documented history, with many advanced problems addressed and solved in fields ranging from medical diagnostics to aerial reconnaissance. These advances have not however been accompanied by parallel solutions in the consumer field, and typically remain the province of the advanced specialist. For the ordinary user the current image processing tools tend to fall into two main categories. The first consists of the comprehensive menu-approach, as used in the more substantial imaging-software packages that are available. These typically collect together a number of statistical-functions, digital-descriptors, analog-photography concepts, etc, and then attempt to provide some degree of control to the user for some or all of these. Such menu-based methodologies inevitably have their inherent advantages and drawbacks. The controls provided are usually non-independent, often have no implicit preferential order of use (even though they are generally non-commutable), and while an experienced and knowledgeable user may apply these powerful software programs to great benefit, the necessary technical background is generally prohibitive to the average consumer. In fact user statistics show that only a very small percentage of consumers make use of such comprehensive facilities.

An alternative approach, growing in commercial popularity but often of very limited effect, is the use of fully- or semi-automatic image-enhancement tools, usually based on specific algorithms that operate on anticipated defects existing in certain classes of digital imagery. While in some limited cases and for certain image-types, these may provide remarkable and satisfying results, all tend to contravene to some degree the important principle of *'first do no harm'*, and are thus usually and necessarily provided with the means to undo their influence and revert to the original image.

As a result of this present status of the consumer field, the overwhelming proportion of digital-images remain in their primitive image-quality state, even though a large fraction of them would naturally lend themselves to immediate perceived improvement and appreciation by the consumer. The trickle-down implications are such that billions of images are considered by users to be unsatisfactory either to print, store, or share, and the negative economic implications to the digital printing imaging industry as a whole are considerable.

The Image Enhancement Solution

Due to the above considerations, there is an obvious and urgent need to provide a simple and effective image-enhancement methodology for the average non-technical consumer. This methodology should embrace all the enhancement benefits of the more sophisticated photo-software programs whereby users retain control and have the ability to choose the enhanced version of the image entirely to their own personal criteria. At the same time, the procedure should have the simplicity of use associated with the typical existing algorithmic *one-click fixes*. The image-processing problem then reduces to that of the feasibility of combining the benefits of these approaches while introducing none of the disadvantages.

To meet this need we have developed a novel enhancement methodology [1] for digital images that takes into account these and other important practical considerations. This resulting enhancement methodology, and the associated consumer software, which by virtue of its simple user-interface, real-time computation, and lack of any appreciable user learning-curve, naturally lends itself to many practical imaging applications in addition to that of stand-alone software. These practical applications include digital printers, cameras, and photo-kiosks, or provision as an image-processing web-service. The development of this methodology has involved a new general approach to image enhancement, as will be described here.

Technical Approach

The general field of image-enhancement covers a wide range of loosely defined terms, and so it is appropriate to give a more precise definition of our own usage. Under our present definition we include all those image attributes that may be thought of as the digital surrogates in the translation from classical analog tone- and color- reproduction theory. These represent all aspects of the image relationship to the original scene in terms of its perceived brightness across all regions of the image, likewise the color reproduction, and the tone or contrast associated with each brightness region of the image. We further make the fundamental assumption that all image manipulations within this domain are obtained within the rule of *determinate pixel mapping*. In other words, only enhancements are assumed permissible which operate in a predetermined manner on each pixel, independent of the state of any adjoining pixel, or groups of pixels.

In this context we note that those techniques that operate conditionally on pixels depending on the state of defined regions of adjoining pixels may be thought of in this present context as advanced image enhancement. In addition to many of the 'single-click-fix' consumer facilities, this class of techniques includes such well-established image processing methodologies as those used for increasing sharpness or reducing noise. In imaging terms, these are often defined in terms of spatial-frequency-dependent operations. In practice the use of advanced enhancement may become a balancing act between desirable image improvements and the addition of new undesirable image artifacts. Examples of these artifacts include image-contouring, haloes and ringing effects, and color spills into adjacent image regions.

The practical reason why the basic set of image enhancement methodologies are not attempted first, and only augmented later by advanced methodologies if and when necessary, lies in the simple reality that there is no obvious systematic way of doing so. However, by consideration of the determinate pixel-mapping basis, so long as the pixel-mapping procedures obey certain obvious rules (continuous, single-valued, finite differentials, well-behaved at the extremities of the pixel range) they are relatively free from the introduction of unwanted image-defects of their own, and the methodology described here is based on this premise.

The key element of the procedure lies in a systematic exploration of the entire basic image-space of brightness/darkness, contrast/tone and color-reproduction. In our own practical experience, when this basic image-space has been fully explored, then the need for augmentation by advanced enhancements is reduced by such a significant amount that typically the consumer is entirely satisfied with the image quality in the absence of any additional advanced image-quality enhancements. Only a much smaller fraction of all consumer images are then deemed to need advanced techniques for, say, image-sharpening or extended-latitude imagery.

In view of the very large number of combinatorial pixel-mapping functions that might be chosen as surrogates for the basic imaging concepts of brightness/darkness, tone/contrast and color-reproduction and balance, this may seem a formidable challenge. Yet any digital photograph has only around five or six independent variables, as seen from a strictly physical viewpoint. If correctly defined, the states of these variables can form the basis of a robust image-quality choice hierarchy.

Technical Details

The essential steps towards a practical solution that includes all the above observations may thus be summarized as follow:

- Define a basic set of individual physical variables representing any digital image.
- Order these independent variables within an overall logical hierarchy.
- Define the practical range of these variables for a comprehensive consumer image set.
- Set the interval scales within these ranges in terms of linear visual effect.
- Determine the just-appreciable visual differences within these scales for the same typical consumer image-set.
- Ensure that the full operation of these variables introduces a negligible degree of associated image artifacts.
- Calculate the total number of combinatorial image states in the image.
- Provide consumer-access to each of these image states using a critical choice hierarchy.

We make the fundamental assumption that the essential physical variables to be used in this basic enhancement methodology can all be determined by unique operations on the basic pixel-map representing the digital image, and that such operations are 'well-behaved' (continuous, single-valued, cover the entire pixel-range, are rational at the pixel-extremities etc).

Figure 1 illustrates the first mode of pixel mapping as a systematic change of the image brightness level. As shown, this yields a systematic and defined enhancement of brightness in the image, while the mirror image of this function naturally represents a corresponding systematic enhancement of image darkness.

Figure 2 illustrates the second mode of pixel mapping, representing a tone-manipulation of the image, whereby mid-tones (mid pixel regions) are associated with increased gain (contrast), at the expense of decreased gain in the shadows and highlights (low and high pixel regions). In this case the mirror image of the curve represents the inverse effect on the contrast associated with these pixel regions.

Finally, as in Figure 3, the individual pixel color (RGB) components are themselves operated on in the above brightness/darkness sense, and using a similar pixel-mapping function.

Having classified the pixel-mapping variables into three basic modes in this elementary manner, we then make the further assumption that these modes may be optimally combined as a linear sequence of ordered operations, starting with brightness/darkness, continuing to contrast/tone-reproduction, and finally addressing color-balance/reproduction, and in this sense constitute a sequence from the largest to the smallest changes in typical images. Hence this sequence comes as close as is practically feasible to representing an independent set of enhancement variables. Our experience with a very large number of consumer images shows that assumption is entirely satisfactory from a practical viewpoint. But in those less frequent cases where large changes in contrast or color balance are the prime need, the sequence may be readily reordered, and in a future version already under development, the use of artificial intelligence will allow the sequence to be automatically selected based on the incoming image characteristics.

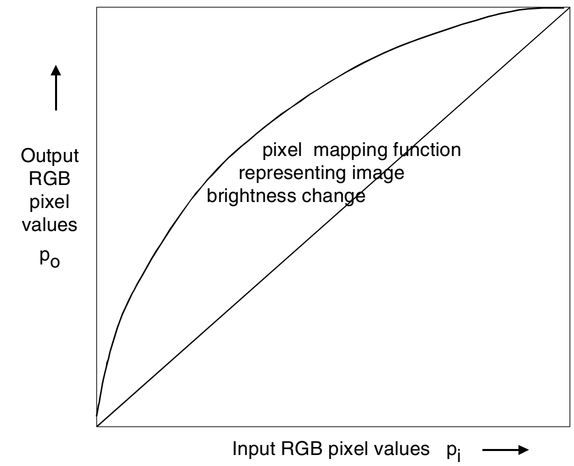

***Figure 1**. Mapping function illustrating overall change in image brightness*

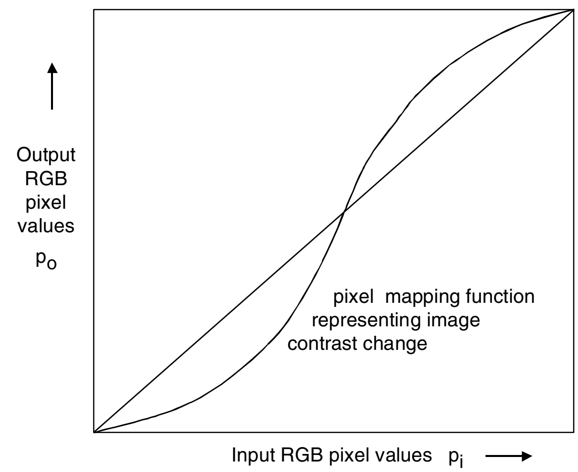

***Figure 2**. Mapping function illustrating change in tone reproduction*

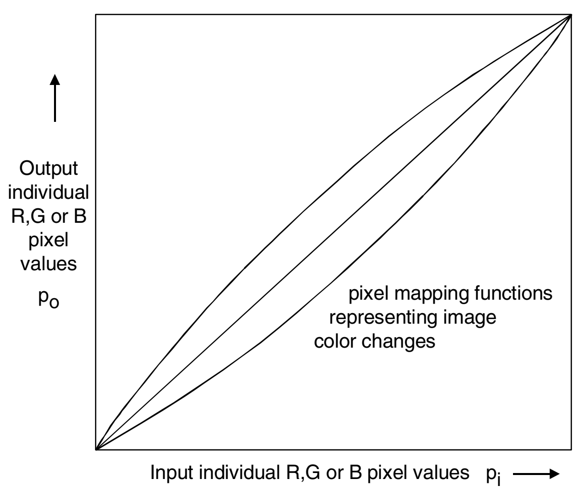

***Figure 3**. Mapping functions illustrating changes in RGB components*

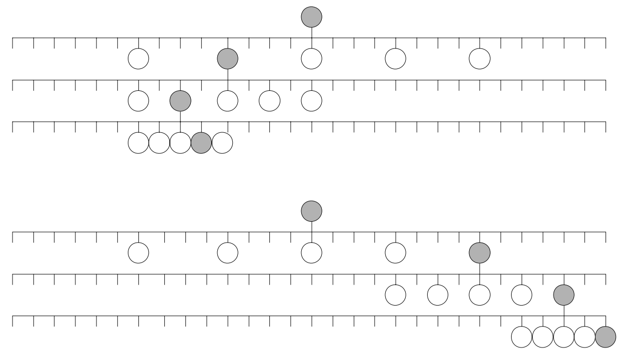

***Figure 4**. Illustrations of the sequential image-choice methodology*

With the pixel-mapping variables established in terms of these corresponding practical image properties, the problem remains of determining the practical range of each variable, and defining the minimum interval within this range that corresponds to a significant difference for each specific image-variable. It is essential that these relate directly and linearly to visual effect, and especially that the intervals within the practical ranges are defined in fixed intervals of visual response. They were in fact determined from several thousand typical images, which included samples ranging from lowest to highest image quality.

The somewhat surprising conclusion was that, within the assumptions of visual linearity, the number of discrete image-states for each attribute could typically be reduced to around thirty. In other words, thirty distinct image-states, correlated with thirty distinct pixel-mapping variations of the assumed basic brightness/lightness surrogate, covered the entire practically established-range, and thereby accommodated more than ninety-nine percent of all consumer images. A similar conclusion was reached for the pixel-mapping function associated with contrast/tone-reproduction, and likewise for color balance and reproduction, although in the latter case, for example for badly faded scanned negatives whose color balance has shifted drastically over the decades, an increased gamut of accessible image color states may be appropriate.

Based on these approximate numbers, the problem associated with a basic enhancement choice-hierarchy becomes clear. Assuming that a number around thirty is appropriate for each of the sequential choices, and that the image-states are independent and combinatorial, the conclusion is reached that there are around 30^3, or 27,000 overall distinct image-states. More exact calculations have shown that while in fact a much greater number of images states (around 10^6) would be necessary to cover the complete gamut of image quality, nevertheless, this lesser number would serve to satisfy the overwhelming majority of user choices. The practical problem of presenting each image to the consumer for individual choice would at first sight seem insurmountable. This would seem especially to be the case in light of the further practical aim of making the choice simpler and more intuitive than in existing comprehensive software packages, to do so without the necessity of setting variables, adjusting curves or moving sliders, etc, and to develop a procedure such that with the minimum of training an inexperienced and non-technical consumer could make this selection without ambiguity within several seconds.

Figure 4 shows a schematic illustration of the principles used in devising an optimum image choice hierarchy. The practical number of distinct brightness states is represented by equally spaced increments. Initially five states are presented to the user, including the original. Choice of one of these (solid-gray) leads to a further set of states being presented to the user. As illustrated, three such choices allow for total exploration of the entire range, even though in the top example only eight image brightness states have been presented to the user along with the original. In the lower example of Figure 4, a critical choice path is shown that allows the user to navigate to the extreme of the practical range of states for this particular image variable.

Having established the basic physical variables, their surrogates in imaging-space, the practical ranges and intervals in linear visual-space, and having then placed the entire image-quality space within a critical choice hierarchy, it remains to translate all these operations into a user-friendly software package. The main ingredient making this possible is the assumption of the simplest mathematical form of the underlying pixel-mapping functions. In this way real-time calculations can readily be made corresponding to each image variable, and recalculation and representation of the next image choices can be made almost simultaneously.

Field Studies

To date several thousand consumer digital images have been enhanced using this new simple procedure, and a small control group of typical users have sampled the software and applied it to their own collection of digital images, mainly but not limited to those images acquired using modern digital cameras.

Some of the initial observations have at first sight been surprising. The original premise for the design and development of this software was primarily that of a rescue operation for the significant number of digital images suffering from any number of a common set of image-acquisition problems, with the implicit assumption that the there would be a decreasing need for the newer generations of sophisticated imaging devices, for example as high-mega-pixel digital cameras proliferate. However this premise has proved wrong, and on reflection this should be no surprise, since using the very logic and principles described above, these more sophisticated image-capture devices will acquire images having the highest signal-to-noise ratio, degrees of freedom, etc, which in our terms we think of as potential independent image-quality states available for exploration. At the same time, the probability that the collective acquisition technology associated with these cameras will place the image in the optimum available image quality state is increasingly remote. Thus the implications are that these high-quality devices present a capability for image enhancement not present with less sophisticated devices.

Our own test enhancement of very high quality digital images has provided a satisfying learning experience in this respect. In effect, the freedom to explore all the alternative available image-quality states leads to a personal selection for each image that may be well displaced from the original, and can transform an already outstanding image into one of ultimate satisfaction.

A large set of samples of typical before-and-after images has been collected from these initial applications of the software. By definition none of these can be reproduced here in any meaningful way, and thus no attempt will be made to do so. Selective sets of these images, representing a wide category of consumer interests and imaging-capabilities and spanning all quality levels, can readily be seen elsewhere [2]. But even these miss the point of the central thesis presented here, namely that the technology, methodology and associated software were developed to exist at every convenient point of consumer access (printer, camera, scanner, computer, photo-kiosk, etc), and that the only meaningful before-and-after comparisons are those made on images enhanced by the actual consumer in the context of the specific viewing conditions at their own particular point-of-access to the image.

Summary and Conclusions

We have described the concepts and operating principles of a practical image-enhancement methodology designed specifically for technically unsophisticated consumers. The associated software is intended for distributed use at any point where the consumer interacts with a digital image, whether in camera, printer, scanner, computer screen, photo-kiosk, or embedded in graphical word-processing software, etc. However it is also ideally suited for central point-of-service applications, exemplified by larger-scale digital printing facilities or central web-based image-enhancement services. The extreme simplicity of use enables instant consumer familiarity without the usual technical complexity of operation.

The imaging theory used as the basis of this methodology has been translated into a primitive set of pixel-mapping equations representing brightness/darkness, contrast/tone, plus color reproduction and balance. The resulting methodology allows the consumer to choose between many thousands of potential image-quality states based entirely on personal preference, and to do so without ambiguity in a matter of seconds. These image quality states are pre-determined by the establishment of calibrated visual ranges and linear visual intervals. Due to the nature of the pixel-mapping equations, minimal image artifacts are introduced during the process. The comprehensive nature of the image-quality space available for exploration means that many previous separate image-enhancement algorithms are implicitly folded into this new overall enhancement methodology.

References

[1] Rodney Shaw and Paul Johnson, Ubiquitous Image processing: A Novel Image-Enhancement Facility for Consumers, Procs IS&T/SPIE Electronic Imaging, Color Imaging XI, SPIE Vol. 6058 (2006).

[2] See practical examples at: www.whiterosedigital.com

Author Biography

Rodney Shaw has a PhD from Cambridge University. His career in imaging research and education includes over a hundred publications, and recognition by IS&T Honorary Membership in 2002. Most recently he was Vice President for Imaging at OurPictures of Palo Alto, and in 2006 co-founded White Rose Digital, an internet provider of digital-imaging solutions, where he currently serves as President.