A Strategy for Pictorial Digital Image Processing (PDIP)

Jack Holm, Consultant
Rush, New York

Introduction

A large variety of approaches are used for processing digital images. These approaches tend to originate in different disciplines, and are usually developed for specific applications. For example, substantial work has gone into the reconstruction of scenes from spatially sampled data,[1,2,3,4] and into colorimetric reproduction models.[5] This work can be usefully applied to the processing of digital photographs (still pictorial images), but does not represent a complete solution to the pictorial processing problem. A number of factors must be considered, factors that may not have been significant to the application for which the original processing approach was developed. Some of these factors have been discussed elsewhere,[6] and will not be listed here. The purpose of this paper is to suggest a comprehensive strategy for the digital processing of still pictorial images.

Overview

The processing strategy developed seeks to accomplish four goals simultaneously:

- Process image data to produce the best possible result in terms of what is desired by the user.
- Minimize complexity whenever possible in order to reduce computational requirements and emphasize the basic function of the processing algorithms employed.
- Automate the processing to the greatest extent that is consistent with hardware capabilities and user quality expectations.
- Improve the efficiency of user adjustments by focusing capabilities on the more likely outcomes, and making the adjustment process as intuitive as possible.*

In accomplishing these goals, it is also helpful to clearly describe each step of the process. Many of the problems with current processing approaches stem from concerns about the *transparency* of the processing, and *proprietary* algorithms. It is difficult for anyone to determine if a particular processing strategy is accomplishing the desired functions, or even if it should be, if the nature of the strategy is kept secret and invisible.

The above goals force processing strategies in specific directions. In particular, it is desirable to consider the physics of imaging systems. Many operations are best performed with the image data in a particular physical representation. Also, physical measurements of the behavior of components in each system can be extremely useful in determining processing parameters. Since little of this information is obtained by the user, it is desirable to automate the transfer of this information, either as part of the image file, or between devices and the processing software. Several newer image file formats accommodate this transfer.

Another consideration in the development of the processing strategy is device independent performance optimization. Digital image data comes from a variety of sources, and may be used for a variety of purposes. For any strategy to be truly useful, it must be able to produce excellent results on a large variety of devices. Device independent performance optimization, however, should not be confused with most current manifestations of device independent color. Optimized performance occasionally results from reproducing colorimetric measurements, but frequently an optimized reproduction will be somewhat different from the original, particularly with photographs. Some of these differences are attributable to differences in human visual system adaptation. Development of a truly comprehensive appearance model, and the reproduction of appearance, would undoubtedly produce optimized reproductions in many cases. Such a model would by necessity be extremely complex, however, and might require difficult to obtain data about surround and viewing conditions, etc. Also, in some cases, even the reproduction of appearance might not produce the preferred result.

For many decades, photography has evolved an empirical set of preferred reproduction goals. These goals have been driven to some extent by materials considerations, but the relatively high quality ceiling of the photographic process prevents media limitations from greatly affecting the goals. A more significant problem is that the goals were not extensively documented. Also, the relative rigidity of chemical processes prevented the goals from being tweaked to take advantage of the flexibility of digital systems. Nevertheless, the implementation of these goals resulted in pictorial imaging systems capable of producing preferred reproductions of excellent quality, and relative insensitivity to changes in the viewing environment.

Recently, attempts have been made by the author to begin to document preferred photographic reproduction goals, and extend them for application to digital processing.[7] A result of this work is the emergence of several issues commonly considered to be of major importance in photography. In particular, the effects of flare, image key (high- or low-), scene dynamic range, viewing conditions, and veiling glare are addressed. Addressing these issues for device independent performance optimization in digital photography requires that the proposed processing strategy be scene and output viewing condition dependent. Photography deals with output viewing conditions through the use of specific media and standard recommendations[8] for specific applications. Scene dependent issues are dealt with by engineering materials for graceful failure, and by human intervention at both the image capture and processing stages. With digital systems, it is possible to shift the scene dependent intervention to smart processing algorithms.

Tone and color processing are of major importance in producing excellent images, but spatial processing can also have a significant effect on quality. The expense of manufacturing one-shot digital cameras with adequate numbers of pixels, and the use of color filter arrays has increased the importance of spatial issues even further. Over the past decade, research in spatial processing has been largely separate from color processing, but device independent performance optimization requires that the two be integrated. Optimized spatial processing is also becoming more output dependent. The spatial frequency capabilities of the output media and the viewing distance affect the degree of sharpening desired. Photographic artists and engineers have long known that mean-square-error (Wiener filter) based restoration is a start toward optimized processing, but that additional edge enhancement is needed. Recent work in information throughput based restoration is beginning to suggest new mathematical approaches to spatial processing.[3,4]

In discussing processing strategies, it is important to differentiate between pictorial processing to produce specific reproduction goals, and image editing or manipulation. In image editing, all or part of an image is intentionally altered to produce a specific effect or make some point. Moving objects around, changing peoples' faces, radically changing the colors of objects, and distortions are clearly image editing. Taking an image and processing it to produce a pleasing result is pictorial processing. The boundary between the two can blur when considering how much of an increase in contrast or saturation is preferred, as opposed to exaggerated. Most popular photographic image processing applications are well suited for image editing. The processing strategy outlined here is oriented toward pictorial processing.

Reproduction Goal Choices

The first step in the processing of pictorial images is to choose the desired reproduction goal. The goal chosen needs to be realizable with the intended capture and output equipment, as well as appropriate for the intended use of the image.

Exact Reproduction

Exact reproduction is where the reproduction and original are identical according to some objective, physical measurement criteria. Two common classes of measurements are used: *colorimetric measurements*—based on the CIE XYZ tristimulus values or some transformation thereof, and *densitometric measurements*—based on ISO status A, M, T, E, or I spectral products.[9] The measurement class chosen must be consistent with the capture device. For example, a non-colorimetric digital camera would not be appropriate to produce exact colorimetric reproduction, unless the colorants used in the scene or original to be captured were known and limited.[†] Different densitometric spectral products are intended for measuring specific materials. It is also important that all the conditions under which the measurements are obtained be well defined, and consistent with the image viewing conditions.

Appearance Reproduction

Appearance reproduction is where the reproduction and original have the same appearance when each is viewed under specific conditions. An exact match is an appearance match if the original and reproduction are on identical media and are viewed under the same conditions, and an exact colorimetric match is an appearance match if the viewing conditions remain constant. Currently, the only way to produce an appearance match under any condition is with manual, trial-and-error processing. Several appearance models have been developed that allow appearance matches to be produced under conditions that vary in specific ways, but the accuracy of the matches varies to some extent with the different models. Appearance models tend to be most successful in dealing with changes in illumination chromaticity. Unfortunately, many other changes are also important. In fact, one criteria for choosing photographic dyes is to minimize changes in appearance due to changes in illumination chromaticity, as long as the observer is adapted to the illumination. Table 1 lists a number of factors affecting the appearance of photographs.

Table 1. Factors Affecting Appearance

Human Visual System Factors (for viewing both the scene and the reproduction)		
Flare in the Eye	Adaptation State:	To the Overall Illumination Level
		To the Illumination Spectral Characteristics
		Spatial Variations in Adaptation
		Intermediate Adaptation to Multiple Conditions
Factors Relating to Characteristics of the Scene or Original (as viewed by the observer, as opposed to a camera or scanner)		
Overall Illumination Level	Illumination Spectral Characteristics	Colorants Used (if known)
Dynamic Range	Scene Key (high- or low-)	Scene Content
Factors Relating to Characteristics of the Reproduction		
Overall Illumination Level	Illumination Spectral Characteristics	Dynamic Range and Surround
Media Type - Surface or Illuminant Mode[10]		Surface Reflections and Veiling Glare
Media Color Synthesis Characteristics — Base Material and Colorant Gamut		

Preferred (Pictorial) Reproduction

In photography, the most common reproduction goal is preferred reproduction, where the aim is to produce the most pleasing image regardless of how well it matches the original. Since there is not an appearance model that deals extensively with dynamic range changes, and since preferred reproduction is highly dependent on dynamic range, it is difficult to say how much of preferred reproduction is an attempt to produce a dynamic range appearance match. If a truly comprehensive appearance model is developed, preferred reproduction may reduce to appearance matching with a slight s-shaped tone reproduction overlay and the corresponding saturation boost. For the time being, preferred reproduction processing frequently offers the best path to excellent photographic quality.

A minor consideration with preferred reproduction is that the exact nature of the reproduction desired depends slightly on the media type. In normal viewing contexts, illuminant mode color images, such as are displayed on a monitor, tend to require somewhat less s-shape than surface mode color images, such as prints, with equivalent dynamic ranges. However, quantification of preferred reproduction for digital systems is just beginning, and small effects are lost in the overall uncertainty.

Preferred Reproduction and Appearance Matching

Table 2 lists the most common reproduction goals for various applications of digital photography.

Table 2. Default Reproduction Goals for Digital Photography

Input Form/ Output Form	Scene	Transparency	Negative	Print
Transparency	Preferred	Exact	Preferred	Appearance
Print	Preferred	Appearance	Preferred	Exact

In Table 2, the default reproduction goal for producing a transparency from a print, or vice versa, is to produce an appearance match. Strictly speaking, the means for achieving this has not been developed because these media have significantly different dynamic ranges when viewed under typical conditions. However, a roundabout approach can be used to achieve the desired result. If it is assumed that the original exhibits preferred reproduction, then it is possible to undo this reproduction back to a linear space, and then implement preferred reproduction on the new media. The result will be very close to an appearance match. This type of processing can be done using a simple look-up-table (LUT) if the output devices utilize a common RGB data space.

sRGB Color Space Processing

Device performance optimization requires that pictorial processing algorithms know the meaning of the digital image data they are presented. If the processing algorithms are specific to particular devices, there are various ways in which this information can be communicated. Device *independent* performance optimization requires that the meaning of the data be understandable regardless of the device. The only way to accomplish this is to establish some sort of standard data space.

Most current color management paradigms make use of a perceptual device independent color space, such as CIE XYZ or CIE L*a*b*. The motivation for this color space type is that images are meant to be viewed, and color descriptions based on psychophysical measurements best predict appearance. This approach is theoretically indisputable if it is possible to construct transforms that accurately convert image data to a psychophysical space, and if the psychophysical space accurately describes appearance. Unfortunately, it is not always possible to construct such transforms, and the lack of a totally comprehensive appearance model may prevent current psychophysical descriptions from predicting appearance. The use of strict perceptual color spaces can also result in fairly intensive processing requiring high precision, since the image data may be transformed through a non-native state.

Alternatives to perceptual color spaces are physically standardized, but more device native “color” spaces. Such spaces describe the physical meaning of the data. It may also be possible to correlate these physical descriptions to current appearance model descriptions for limited sets of viewing conditions. Obvious candidate spaces of this type are the standard RGB spaces. Of those available, the most appropriate are the monitor spaces. Monitor spaces have wide gamuts, most images are viewed on a monitor at some point, a great deal of manual processing is accomplished using monitor feedback, and an internationally standardized monitor space already exists, as well as correlations to appearance under specific viewing conditions. Standard monitor data, when printed on photographic media using devices with independent channels, also tends to correlate reasonably well with Status A densitometry. This means that photographically derived preferred reproduction models can be applied. It is also interesting to note that recent work in the color appearance area is indicating that the use of spectrally sharpened visual response functions is advantageous.[11,12] These functions are much closer to RGB than the unsharpened visual response functions. Table 3 summarizes the advantages and disadvantages of the two types of standard color data spaces.

The standard monitor data approach provides a common ground to link perception to the native physical behavior of devices based on RGB and CMY colorants. Most devices that produce outputs of varying dynamic range use colorants of these types. Output devices that use other colorants, but have dynamic ranges of around 100:1 can also be accommodated since the monitor data can be correlated with perceptual metrics at the fixed dynamic range. Various forms of standard monitor RGB have been successfully used by practitioners of digital photography for several years.[13] Recently, a few major corporations have formalized this approach by proposing a specific standard monitor color space, *sRGB*.[14] This proposal paves the way for the use of a standard RGB space for PDIP.

Table 3. Advantages and Disadvantages of Perceptual and sRGB Color Spaces

CIE XYZ and L*a*b* Color Spaces

Advantages

- Excellent color appearance reproduction if the capture is colorimetric or the colorants used in the original are known, and the viewing conditions and media dynamic range are appropriate.
- Can reproduce color using unusual colorants as long as the viewing conditions and media dynamic range are appropriate.
- L*a*b* is reasonably uniform perceptually.

Disadvantages

- The color reproduction accuracy advantage is lost if the capture is not colorimetric or the colorants used in the original are not known, as is usually the case with digital cameras.
- Color appearance prediction may be poor if the output media dynamic range and/or viewing conditions are significantly different from the original.
- Processing may be more extensive and require higher precision.
- No model is available for preferred reproduction.
- If all the gamut benefits are to be realized, the image data may need to be stored at high precision, or the raw data stored with a transform.

sRGB Color Space

Advantages

- Similar to many device native color spaces. Transformations to sRGB tend to be simpler, more accurate, and require less precision for storage. It is less necessary to save the raw data with a transform. Transformations from sRGB to output device spaces also tend to be simpler.
- Since sRGB image data can also be described perceptually, the advantages of the perceptual color spaces can be applied.
- Photographic preferred reproduction models can be applied.
- Reasonably uniform perceptually.
- Relatively independent channels help with signal-to-noise issues in capture.
- May be similar to the spectrally sharpened tristimulus metrics to be used in future appearance models.

Disadvantages

- Colors that are out of the monitor gamut are expressed using negative values, requiring larger data volumes.

Pictorial Digital Camera Processing Pipeline

The following is a proposed optimized pipeline for PDIP.

Preliminary Measurements

1. Set the camera gain(s) and offset(s) as they will be set during use, hopefully the optimum settings. Ideally, the offset should be set so that a focal plane exposure of zero produces a digital level of zero after bias and dark current subtraction. The spacing of the levels should be chosen so that virtually all the information the camera can capture is recorded with minimal quantization error. A rule of thumb is that the standard deviation, expressed in digital levels, of any even focal plane exposure, should be at least 0.5 after all fixed pattern noise removal processing. Higher standard deviations are also acceptable, but will require more bits for image data storage.
2. Determine the camera fixed pattern noise characteristics, such as dark current and pixel sensitivity non-uniformity.
3. Measure the camera response characteristics to determine the focal plane OECF for each channel for each illumination type of interest.[15] Variations due to illumination type can usually be dealt with using a single OECF curve shape and channel multipliers. For example, a particular camera may have relative sensitivities of 0.6, 0.9, and 1 with 5500K daylight illumination and 0.2, 0.7, and 1 with 3200K tungsten illumination. If channel multipliers are used, it is possible to determine a mathematical model to predict multipliers for intermediate illumination types.
4. Determine the camera OECF's for a variety of simulated scenes.[15] Use this information in conjunction with the focal plane OECF measurements to devise a model that predicts flare based on focal plane image statistics.
5. Measure the camera channel spectral responses and determine a linear radiance space spectral sharpening matrix.
6. Measure the linearized camera spatial frequency responses and noise power spectra for each channel in the horizontal and vertical directions. Determine a reasonable processing kernel size and construct a maximum information throughput spatial reconstruction kernel. Note that each channel may require a different reconstruction kernel, and that the ideal reconstruction kernel may vary with focal plane exposure.
7. Measure the neutral EOCF (electro-optical conversion function), the spatial frequency response, and the noise power spectrum of the output devices on which the image data may be rendered. If an output device has minimal noise, it mat not be necessary to measure the noise power spectrum. Spatial frequency response and noise power spectrum measurements are also not necessary with halftoning output devices that incorporate perceptual noise suppression and sharpening into the halftoning algorithm. If the output device is not known, the EOCF for sRGB can be used.

Processing Step 1: Determination of Flare and Scene Key

1. Divide up the image data into the color channels (if not already done).
2. Pixel average (boxcar filter and sub-sample) each channel to obtain meter arrays of 10,000 to 20,000 digital values, and store the meter arrays. Previous work has indicated that meter arrays with a short dimension of about 100 pixels are optimal for most subject matter.

The pixel averaging is done in what is assumed to be an approximately logarithmic gamma type camera data space because geometric means are preferable to arithmetic means.

Note: If the camera is linear, it would be best to convert the digital values to log space prior to the pixel averaging, and then back again, although the effect of this rather intensive additional processing should be minimal.

3. Transform the meter array values to focal plane exposures using the inverse OECF, and take the log.
4. Determine the minimum, maximum, and mean log exposures. Estimate the flare factors based on these values.
5. Determine the image specific camera OECF's for each channel based on the focal plane OECF's for the illumination type used, and the flare factors.
6. The scene key is determined by subtracting the average of the minimum and maximum log exposures from the mean log exposure.

Processing Step 2: Determination of Scene Zone 1, Zone 5, and Zone 9 Luminances

1. Transform the meter array values into spectrally sharpened linear scene values using the processing sequence as outlined in steps 4 and 6.
2. Combine the channel values into a luminance channel using an equation appropriate for the channels used. If the spectral sharpening is intended to produce sRGB, the luminance conversion equation from ITU-R BT.709 can be used.
3. The minimum meter array luminance is assumed to be the scene Zone 1 luminance, the arithmetic mean luminance the Zone 5 luminance, and the maximum luminance the Zone 9 luminance.

Processing Step 3: Determination of Output Table

1. Select the output device and pixel pitch for the desired rendering. If the output device is not known, assume a standard monitor as represented by sRGB.

Linear Reproduction

2. Determine the digital level that produces 18% reflectance, or 18% transmittance relative to the base transmittance, on the output device.‡ This is designated as the midtone reflectance level.
3. Determine a LUT that will produce an output with reflectances that are a constant multiplier of the scene luminances, with the constant chosen so that the Zone 5 luminance reproduces at the midtone reflectance.

Preferred Pictorial Reproduction

2. Determine the output device Zone 1 and Zone 9 densities, typically 0.04 above base plus fog and 90% of the maximum density.
3. Calculate the preferred reproduction relationship between the scene luminances and output densities. The details of this calculation are provided elsewhere.[7]
4. Determine a LUT that will produce preferred reproduction on the selected output device, or a standard monitor.

Processing Step 4: Scene Linearization

1. Subtract and divide out the fixed pattern noise (if not already done).
2. Construct scene linearization tables by taking each possible digital value through the image specific inverse camera OECF's.
3. Convert the pixel digital values to linear scene channel radiance.

Processing Step 5: Spatial Restoration

Note: This step is placed here because most spatial restoration techniques assume the data is in a linear radiance space. There may be advantages to placing this processing later in the pipeline if the reconstruction algorithm is designed to deal with spectrally sharpened and/ or nonlinear data.

1. Apply the maximum information throughput spatial reconstruction kernel to each channel.
2. Apply any morphological or other nonlinear processing to the image to reduce artifacts (most common with CFA camera data).

Note: Memory requirements can be reduced by performing the spatial restoration operation on sections of the image at a time.

Processing Step 6: Spectral Sharpening

Note: this step is not necessary if the spectral capture bands are sufficiently narrow, such as if color separation filters are used for a color sequential exposure.

1. Apply the linear radiance space spectral sharpening matrix.

Processing Step 7: Output Processing

1. Apply the desired output LUT.
2. Apply any subsequent output processing, such as sharpening, noise reduction, halftoning, etc.

Subsequent Processing

Image data that has been processed to a particular reproduction goal on one output device can be processed for another output device by undoing the processing back to the point where the processed data is common to both output devices. For example, if the original and new output devices assume sRGB data, but have different EOCF's, it is possible to transform the data using one simple LUT. If the color spaces are different, it will be necessary to go back before the spectral sharpening matrix, and then redo the matrix using one that is appropriate for the new output device. Changes in reproduction goal are similar; the new reproduction goal on the same output device can be viewed as a different output device by the processing.

Processing Strategies for Other Applications

Colorimetric Digital Camera Processing Pipeline

The prime requirement for colorimetric capture is that the camera channel spectral sensitivities, or some linear combination thereof, be able to simulate the human visual system spectral sensitivities. If this is the case for a particular digital camera, it is possible to use the pipeline described above to produce colorimetric measurements. Transformations to the desired colorimetric space, instead of sRGB, can be accomplished after linearization. This pipeline has an advantage over most current approaches in that flare effects are removed in an explicit, image dependent manner.§ Preferred reproduction can also be produced by using a spectral matrix that produces sRGB data.

Transparency Film Scanner Processing Pipeline

Colorimetric Scanner or Known Colorants, and Exact Reproduction

If a scanner has colorimetric channel sensitivities, it can be equated to a colorimetric digital camera, and the same processing approach used. In many cases, however, non-colorimetric scanners are used to obtain colorimetric data through the use of a model that maps the scanner measurements to colorimetric measurements for a particular set of known colorants. The IT8.7/1 scanner characterization targets can be used to construct such models. These models frequently take the form of multiple LUT's or matrices of equations. While it is possible to use a matrix of equations for the spectral matrix in the pictorial image processing strategy described above, doing so is a significant extension of the approach. It is probably more appropriate in such cases to consider the transformation to colorimetric data as part of the capture processing. The data then becomes analogous to colorimetric digital camera data and can be dealt with similarly.

Unknown Colorants or Preferred Reproduction

If the colorants in the film scanned are unknown, or if preferred reproduction is the goal, the processing pipeline described in the previous section is appropriate. The accuracy of color reproduction with this approach will usually not be as good as with exact colorimetric reproduction, but the advantage of being able to produce preferred reproduction, particularly on media with different dynamic ranges, may outweigh the drawbacks. If the colorants used in the transparency are known, the color reproduction accuracy can be improved substantially by customizing the spectral matrix for each colorant set. The spectral responses of the channels determine how sensitive a particular scanner will be to this effect. Status A responses are designed to produce good results when used to measure any of the dye sets commonly found in photographic transparencies.

RGB processing of scanned transparency data has other advantages. Frequently, little or no matrixing is required and the desired result can be achieved using three LUT's. Also, models can be used to correct exposure errors on original transparencies. Measurements of the transparency are used to estimate linear scene RGB radiances, which are then processed to preferred reproduction. However, it is important to remember that the variability of film processing reduces the accuracy with which film scans can predict scene radiances. The photographic process relies on the attractiveness of preferred reproduction to compensate for processing errors. Also, if a transparency is assumed to already exhibit preferred reproduction, it is not necessary for the scene radiance estimates be accurate. Undoing the preferred reproduction on the original, and then redoing it on the reproduction will produce preferred reproduction even if the scene estimates were incorrect.

Negative Film Scanner Processing Pipeline

Exact colorimetric reproduction is not a viable option in scanning negatives to produce positive images. The film did not capture the scene information colorimetrically, and the colors in the negative are not reproduced in the positive image. However, in some respects negatives are better suited for film capture than positives. They have a larger input dynamic range, more exposure latitude, and the unwanted absorptions of the film dyes are corrected using masks, making the RGB channels in a negative more orthogonal. Also, the output dynamic range of negatives is lower than with transparencies, reducing flare and the dynamic range requirements for the scanner. Unfortunately, many scanners have difficulty producing good image data from negatives because they are designed to scan transparencies. The level spacing is too large for the lower dynamic range, and the processing software does not know what to do with the reversed and offset RGB color data. Another complication is that negative dyes are designed to modulate the red channel information at longer wavelengths. Negatives should be scanned using channels with spectral sensitivities that allow Status M densities to be obtained, as opposed to the Status A densities appropriate for transparencies. The best way to process negative data is to use models that estimate linear scene RGB radiances, which can then be processed to preferred reproduction.

Print Scanner Processing Pipeline

Print scanning is complicated by surface reflections which are highly dependent on the scanner optical geometry. The two approaches listed for transparency scanning can be applied to print scanning, but if preferred reproduction is desired it is necessary to scan the print in such a way that the surface reflections are minimized. The geometry specified for reflection densitometry measurements[16] is best for this purpose.

If exact colorimetric reproduction is desired, it becomes necessary to simulate the viewing geometry in the scanner. Ideally, different scanner geometries should be used for different viewing conditions. This is rarely practical, however, and the most common approach to colorimetric scanning is to use an integrating sphere, frequently with a gloss trap. However, measurements taken using this type of geometry will be inferior for producing preferred pictorial reproduction. Another approach is to use the geometry for reflection densitometry and simulate veiling glare. This approach is rarely used because a model that can simulate the veiling glare characteristics of the material scanned is required.

User Adjustments

A comprehensive strategy for PDIP must allow for user input to the final result. No automated routine can account for individual taste. Also, the strategy presented in this paper does not deal with the nature of objects in images. These objects (such as people) can have a significant effect on the desired reproduction, and object recognition and classification is probably the next frontier in automated image processing.

Another consideration is that current processing software does not provide for the quick and easy adjustment of reproduction. It is necessary to have a processing approach in place before user adjustments can be added. Automated image processing to preferred reproduction will be a major advance in digital photography, but quick, easy, and intuitive tweaking of the result is almost equally important. Table 4 lists a variety of user adjustment options. These options are divided into levels, so that the novice user will see only the simple controls, while more advanced users can choose to view more options. Structures of this type are common in many software applications. In all cases it is assumed that visual feedback for all choices is provided in real time, that user preferences can be stored, that all technical and device information that can be transferred automatically is, and that the user can request that default values be selected where not provided and used for all values, for some values, or that the interface ask each time. Complete descriptions of how each adjustment affects the processing of the image data should be available for advanced users.

Table 4. Manual Adjustment Options

Level O Choices:

Illumination source - accept default or device estimation, specify type, or use a neutral balance feature.

Reproduction goal - exact or preferred.

Output device (default to sRGB, or to a device specified in the image file).

Level 1 Adjustments:

Brightness Slider

Exact reproduction - adjusts the output density of the scene arithmetic mean luminance.

Preferred reproduction (fine) - adjusts the amount of high- or low-key shift.

Preferred reproduction (coarse) - implements a gamma type brightness shift.

Contrast Slider

Exact matching - disabled.

Preferred reproduction (fine) - adjusts the amount of flex.

Preferred reproduction (coarse) - adjusts the meter array cell size.

Color Balance Slider (3 options)

Adjust using color temperature slider.

Adjust using trilinear coordinate (RGB) joystick.

Adjust using rectangular coordinate (opponent color) joystick.

Sharpness Slider

Level 2 Adjustments:

Exact reproduction

Choose to base the midtone on the geometric mean.

Preferred reproduction

For each channel, allow adjustment of the meter array cell size, Zone 1 and Zone 9 output density, and flare factor.

Choose to base the key on the arithmetic mean.

Level 3 Adjustments:

Apply user specified look-up-tables and matrices

Specify look-up-table to convert each channel to scene radiance.

Specify matrix to convert to different channels for output.

Specify look-up-table to produce desired reproduction on output.

Conclusions

The tools and techniques required for the improved, efficient, and automated processing of pictorial images are becoming available. If processing of this type is implemented, the quality of digital photographs should surpass that of conventional photographs in most areas, resulting in rapid acceleration in the acceptance of digital photography.

References

1. W.K. Pratt, *Digital Image Processing*, Wiley, 1978.
2. C. R. Giardina and E. R. Dougherty, *Morphological Methods in Image and Signal Processing*, Prentice-Hall, 1988.
3. C.L. Fales, , F.O. Huck, R. Alter-Gartenberg and Z. Rahman, “Image gathering and digital restoration,” *Phil. Trans. R. Soc. Lond.*, in press.
4. F.O. Huck, C.L. Fales and Z. Rahman, “An information theory of visual communication,” *Phil. Trans. R. Soc. Lond.*, in press.
5. R.W.G.Hunt, *The Reproduction of Colour*, Fifth Edition, Fountain Press, 1995.
6 Jack Holm, “Factors to Consider in Pictorial Digital Image Processing,” *Proceedings, 49th Annual Conference of the Society for Imaging Science & Technology*, p. 298-304, 1996.
7. Jack Holm, “Photographic Tone and Colour Reproduction Goals,” *Proceedings, CIE Expert Symposium ‘96: Colour Standards for Image Technology*, p. 51-56, 1996.
8. ISO/CD 3664, *Viewing conditions - Prints, transparencies, and substrates for graphic technology and photography*, June 15, 1996.
9. ISO 5-3, *Photography - Density measurements - Part 3: Spectral conditions*, 1995.
10. L. Stroebel, J. Compton, I. Current and R. Zakia, *Photographic Materials and Processes*, Focal Press, 1986.
11. Finlayson, G.D., M.S. Drew and B.V. Funt, “Spectral sharpening: sensor transformations for improved color constancy,” *J. Opt. Soc. Am. A*, **11**:5, p. 1553-1563, 1994.
12. Finlayson, G.D., M.S. Drew and B.V. Funt, “Color constancy: generalized diagonal transforms suffice,” *J. Opt. Soc. Am. A*, **11**:11, p. 3011-3019, 1994.
13. J. Holm, and N. Judge, “Electronic Photography at the NASA Langley Research Center,” *Proceedings, 48th Annual Conference of the Society for Imaging Science & Technology*, p. 436-441, 1995.
14. M. S. Anderson, R. Chandrasekar, R. Motta and M. Stokes, “Proposal for a Standard Color Space for the Internet - sRGB,” included in these proceedings.
15. ISO/CD 14524, *Photography - Electronic still picture cameras - Methods for measuring opto-electronic conversion functions (OECF’s)*, June 15, 1996.
16. ISO 5-4, *Photography - Density measurements - Part 4: Geometric conditions for reflection density*, 1995.

[*] It is important to remember that intuition is relative. For example, color scientists used to working with CIE L*a*b* values may find adjustments in L*, a* and b* to be the most intuitive, while photographers used to working with R, G and B may find adjustments of this type to be more intuitive. A novice might prefer the most computationally simple approach, as opposed to any particular color space, so that feedback is quicker (as long as the increments are perceptually reasonable).

[†] If it is possible to determine colorant amounts through any analysis technique, it is usually possible to convert these amounts to colorimetric (or densitometric) measurements.

[‡] Print materials are viewed using reflected light, and the visual reference white is a true 100% diffuse reflector. Transparencies are viewed with transmitted light, and the appropriate reference white is the transmission of the transparency base.

[§] When using a camera for exact colorimetric measurements, the removal of flare tends to be less important than when producing preferred pictorial reproduction. Colorimetry is based largely on the measurement of surface colors, which are limited in dynamic range. The effect of flare is minimal with scenes of low dynamic range.