

Integrating New Color Image Processing Techniques with Color Management

Jack Holm
Hewlett-Packard Laboratories
Palo Alto, California, USA

Introduction

The advent of digital photography as a serious alternative to silver halide processes has led to the invention and investigation of a number of new color image processing techniques. The performance of these techniques has far reaching implications. It is clear that a number of visual phenomenon associated with the viewing of complex scenes and images are becoming increasingly important. These phenomenon should eventually result in a fundamental generalization of vision models from single stimulus to complex image models, but the difficulty of this task is hard to overstate. The benefits are also large, both scientifically and commercially. It is beginning to appear that a robust model for complex images will require much closer adherence to human visual system processing.[1] Current single stimulus empirical models such as CIECAM97s are quite complicated, and the number of degrees of freedom expands enormously with complex images. It is simply no longer practical to pile complexity on top of complexity in empirical models.

The commercial implications of this situation are that it is becoming more necessary to use a reproduction model approach for color management. With this type of approach, color descriptions are spectral and/or colorimetric, with associated information such as media gamut and dynamic range, and viewing conditions. There is also a distinction between scene descriptions and reproduction descriptions. The reproduction model approach allows proprietary solutions to be utilized to go from the scene description to the reproduction description. Such models have been used in practice for many years in the photography, graphic arts, motion picture, and television industries.[2,3] However, since these industries generally involved closed loop systems, it was not necessary to recognize the existence of reproduction models per se, even though they were in effect being used. The cross-application nature of digital imaging makes it necessary to recognize what is going on for effective communication of color information.[4,5]

Another interesting observation is the fact that before digital scene capture (and computer scene rendering algorithms), all image data comprised a reproduction description. The only way to get an image was to photograph it or create it by some artistic means. In each of these cases, a reproduction description was produced. For pictorial imaging, these reproduction descriptions tended to be limited to the following classes:

1. Reflection hardcopy, with density ranges typically between 1.4 and 2.4. There are many types of reflection hardcopy media, but the most important are photographic, graphic arts, and artistic materials. These materials tend to have somewhat similar gamuts and density ranges as noted above, as is required for producing pictorial images of reasonable quality. Viewing conditions are variable, but since these images are transportable, standard viewing conditions are typically used for evaluation. It is also desirable to link the evaluation viewing conditions to the optimal ones for the medium used.
2. Transparencies, with density ranges up to 3.0. Transparencies are perceived differently than reflection hardcopy, but there are not too many types of transparencies, and not too many different viewing conditions are used.
3. Color negatives, which are at an intermediate stage. It is possible to obtain scene information from a color negative, but most approaches to date have "printed" the negative on "digital photographic paper" or used some other proprietary algorithm to effectively produce a reflection hardcopy reproduction description.

The result of this situation is that, prior to digital capture, color management only had to deal with reproduction descriptions with a limited number of characteristics. Another circumstance that supported early approaches was the generally abysmal nature of digital images before color management. It was possible to effect enormous improvement without using conceptually complete approaches. It was also a fact that with many of the tasks color management had to perform it was not that far from conceptually complete, since it was starting with a reproduction description and just had to produce another reproduction description that was a pleasing and reasonable appearance match.

Digital scene capture "changes everything". It is no longer possible to assume that the original has a well defined white (or even a choice of two),[6] or a dynamic range or gamut reasonably close to that of the desired output medium. This situation forces the differentiation of scene and reproduction descriptions and the explicit use of reproduction models for cross-application imaging as noted above. The flexibility of digital systems also makes the use of reproduction models desirable for single-application processing. The question then becomes, how can the new image processing techniques encompassed by the reproduction model approach be integrated into color management? This paper will discuss work that is being done to facilitate this integration for digital photography, and propose solutions as to how such integration can be accomplished in general.

Discussion

The value of digital color management as the means for communication of color in open systems is unquestionable. Closed system approaches are no longer viable for many applications. Also, the traditional approaches embodied in conventional graphic arts and photography systems are not necessarily the best approaches, even if it is possible to digitally replicate the chemical and mechanical processes. However, the honeymoon for color management is ending. Customer image quality objectives remain the same regardless of the technology used, and these desires can be fulfilled by available solutions. In some applications, the novelty or utility of a digital imaging solution is of sufficient value that image quality can be sacrificed, or the demands may not have been too high in the first place. In other applications, the use of open systems, digital color management is no justification for a quality decrease, and need not result in one.

Until recently, the capabilities of color management have been limited by a number of misconceptions, or assumptions that were reasonably accurate only in the protected environment of multimedia and computer imaging. With these applications, the question was not whether to use a conventional or digital approach, but how good were the results produced by the digital approach. Many problems have been identified and dealt with and quality has improved. However, there has been a reluctance to abandon the initial assumptions because of the potential broad impact on systems in place, and because when one gets used to thinking along certain lines it is disorienting to change the foundation.

Another major contributing factor is that many of the color reproduction approaches used in traditional graphic arts and photography are difficult to translate for open systems. An overly simplistic summary is as follows:

Traditional Color Reproduction Practitioners - know what a good reproduction is, how it relates to the scene or original visually, and how to produce it using conventional materials and processes, but may not remember how the conventional methods were arrived at in the first place from basic imaging concepts, and can't communicate knowledge of how to develop analogous methods for digital systems, particularly open systems.

Digital Color Reproduction Practitioners - know how to work in an open systems environment and develop digital methods for color reproduction; but have difficulty understanding the motivations for various aspects of closed system methods, and in particular discerning which aspects result from the closed system approach and materials and process limitations, and which are necessary prerequisites to achieving the required image quality.

The resulting communication gap has proven to be formidable. The following discussion is an attempt to provide some parts of the bridge.

Common Misconceptions

Device Independent Color

Originally, the idea of device independent color related to metrics for measuring color. The Status A, M, and T spectral products used in graphic arts and photography were designed for measuring the particular colorants used by these processes, and it was observed that when other colorants were used the same visual appearance did not necessarily produce the same measured values on patches. Other factors relating to the geometric measurement conditions also affected results, but it was clear that the use of CIE standard observer functions (or derivatives) for spectral integration was desirable for quantifying the appearance of patches produced using different colorants. This had long been the practice in the paint and textile industries.

The misconception about what the device independent color really means began when people made two leaps that are not necessarily appropriate for imaging. The first was that the appearance of a patch is independent of the other colors in the image. The second was that it is possible to come up with a representation of the colors in an image in a way that is independent of the medium. The idea of using CIE observers to produce colorant independent measurements of color patches is valid. The extension of the device independent color concept to imply that a single, media independent color representation for an image is possible is not.

Gamut Mapping

One of the first problems noticed in attempts to produce device independent color descriptions of images was that different media have different color gamut and dynamic range capabilities. This meant that when a device independent color description was produced for an image, it might not be possible to reproduce all the colors in the image on a particular medium. It also became apparent that if the description was produced for a large gamut medium, that simply clipping the out of gamut colors for a small gamut medium did not produce acceptable results.

Producing colorimetric descriptions with sufficiently limited gamut to be within the capabilities of most media also did not produce acceptable results on large gamut media.

This problem brought about considerable research in gamut mapping, or the art of transforming an image description produced for one gamut to one appropriate for a different gamut. However, this research was arguably not focused on exactly the right objective. It was assumed that the problem was the gamut mismatch, and that the transforms required were independent of the image content. This was in contradiction to the experience of practitioners in graphic arts and photography.

Appearance Matching

Another observation was that a colorimetric description of an image produced for one viewing condition is not necessarily appropriate for another viewing condition. Generally, the darker the surround and the brighter the image, the darker a particular tone in an image should be reproduced (effectively making the image darker and contrastier). A striking example is the difference in the luminance factor of a midtone relative to media white for a reflection print viewed at a typical indoor illumination level, as opposed to a slide or motion picture projected in a dark room. In this case, most of the effect is due to the surround, because the absolute luminance levels are not too different. Another example is that in the graphic arts, the high illumination levels used for comparing proof to print generally resulted in prints that were considered to have good tone reproduction under typical viewing conditions being perceived as too light and flat.

This seemed to be a breakthrough in that it provided an explanation of why reproducing colorimetry did not always produce a desirable result. The stated objective all along had been to reproduce the "appearance" of an image on various media. The initial assumption (by some) in the color management community had been that the colorimetry described the appearance. The color science community had known for some time that viewing conditions also affect appearance, even of single stimulus patches, and a lot of work had gone into modeling these effects. It began to seem that there was no problem with the basic concept of a media independent representation of a color image, but the problem was that colorimetry did not completely describe appearance.

Unfortunately, this is not entirely correct because frequently, media limitations as opposed to viewing conditions, are the driving factor in determining how to render an image. Also, to the extent that describing appearance does provide a media independent description of a scene, there are implementation problems because most of the appearance models developed by color scientists are intended for use with single stimuli, and can be demonstrated to break down when applied to complex images in a global manner. Some research has been done to develop appearance models for complex images,[7,8,9] and to apply single stimulus appearance models locally to complex images, but this research is not yet conclusive.

Example - Transportable Appearance

Say that one has a reflection device capable of actively adjusting image relative colorimetry to compensate for changes in viewing conditions using an appearance model. One might consider this to be an ideal device, but would it be? Say that we photograph an outdoor scene and view it using this device. As the viewing conditions get dimmer (average surround), the device would boost colorfulness to try to maintain appearance. At first this might be pleasing, but as conditions got dimmer the increasing boost would begin to look odd. We are not accustomed to seeing daylight levels of colorfulness at very dim illumination levels. A standard reflection print, on the other hand, would have some amount of boost built into the print by default, but would then behave as a "normal" object when viewed in increasingly dimmer conditions.

Conversely, if the same device is used to reproduce a dim indoor scene in a bright viewing environment, it will tend to prevent the reproduction from being too colorful. The reproduction may therefore accurately represent the appearance of the original scene, but most viewers would choose an "incorrect" reproduction with more colorfulness. An extreme example is astronomical photography where grayish wisps as viewed through the telescope become a brilliant colorful universe when colorimetry, as opposed to appearance, is reproduced.

If the performance of the "ideal" device is compared to that of a typical photographic print, the photographic print will frequently be preferred. However, there is a caveat. Changing tone reproduction characteristics as the viewing illumination level changes is known to be desirable. Perhaps we need an "active appearance" device that operates primarily on luminance.

Color Space Transformations

Another area of confusion that grew out of the above misconceptions was the idea of a color space transform as a model for color image processing. Given an image description in a color space suitable for a particular application, many people assumed that if the image data were transformed to another color space appropriate for another application, that the color space transformation would include the appropriate rendering transform. While this may occur in some situations, it cannot be assumed to occur. A number of color spaces, such as the CIE spaces XYZ, LAB, and LUV, and the newly introduced sRGB-64,[10] can be used to contain a variety of color renderings. Their strength is their flexibility. Their weakness is that knowing image data is in one of these spaces does not in itself provide definitive information about how the data has been rendered.

There are some caveats. For example, if no additional information is provided identifying some other rendering, image data encoded in sRGB-64 should be assumed to be rendered to the 8-bit sRGB dynamic range, gamut, and viewing conditions. Also, since CIELAB approximates the Munsell color space, CIELAB data most appropriately describes a reflectance image viewed using conditions similar to those under which the Munsell space was constructed. Such an image could be an original painting or other two dimensional artwork where the CIELAB values describe the colorimetry of the original, or a reproduction description intended for a particular media. There is no way to tell the exact nature of the rendering from the fact that the image is encoded as CIELAB, although one can reasonably assume a two dimensional reflection medium with no specular highlights.

Multi-purpose color spaces are useful in that rendering processing can occur without the additional need of a color space transformation. Single purpose color spaces are useful in that the designation of the color space also designates the rendering. Complimentary color spaces where one is rendered into the other offer both advantages.

Example - sRGB Gamut Limitations

A particularly noteworthy example of this confusion relates to debate that has occurred concerning the gamut of sRGB reproduction descriptions. Many people have complained that the gamut limitations of sRGB cause problems with encoding images into sRGB, or repurposing sRGB image data to some other medium. In extreme cases this criticism could be valid, and it would be inappropriate to encode image data for such an application into sRGB. However, in the vast majority of cases, the complaints about gamut limitations relate more to the lack of tools for rendering into and out of sRGB than to the characteristics of the color space itself.

Consider a scene that is captured using a digital camera into ISO RGB,[11] and also using Brand X slide film which is then scanned colorimetrically into sRGB-64. Say that the ISO RGB image data is also encoded into sRGB-64 for convenience. In both cases the image data is encoded in sRGB-64, but the data is different because in one case it represents scene colorimetry and in the other the colorimetry of the transparency reproduction. The scene representation should be tagged "ISO RGB" and the slide film representation "Brand X slide film". Now say that a proprietary rendering algorithm takes these image files and renders them for sRGB, with the rendering intent being to produce a preferred reproduction of the scene. (The only other reasonable rendering intent for sRGB would be to produce an appearance match the transparency, but this would only be appropriate if the transparency embodied some artistic preference, and a good scene rendering algorithm could embody a similar artistic preference.)

Now say that the sRGB rendering is encoded in sRGB-64. We now have three colorimetrically different versions of the same scene encoded in the same color space. This clearly demonstrates that a color space need not contain a unique rendering. It also demonstrates that a color space transform does not necessarily perform rendering, because in transforming the scene description from ISO RGB, or the CRT display reproduction description from sRGB, no rendering occurred. The problem with the sRGB gamut being viewed as inadequate for encoding data results from the rendering of the image to sRGB not being performed adequately before encoding into sRGB. This would happen if either the scene or slide film reproduction were encoded as sRGB with no rendering (other than clipping).

The other criticism of the sRGB gamut concerned repurposing sRGB image data. There is also no need for this to be a problem if the data was well rendered into sRGB in the first place. If it was, specular highlights, shadows darker than the sRGB maximum density, and out-of-gamut colors will have been compressed (as opposed to clipped), and the only limitation on re-expansion for a wider dynamic range medium is the precision of the data. In extreme situations, or when a significant re-rendering to change the artistic intent is performed, lack of precision can be an issue. In most cases it is not, and sRGB image data can be repurposed to virtually any reflection media (and even transparencies) with excellent results, if it was well rendered into sRGB in the first place, and the nature of the rendering transformation is known. The latter is necessary so that the physical meaning of the data can be understood to facilitate intelligent re-rendering.

It may be worth mentioning that in the example described above, there is no reason to assume that if the sRGB-64 colorimetry was transformed into CIECAM97s appearance correlates, that the appearance descriptions would be equivalent. The goal of the slide film and sRGB renderings may not be to reproduce the appearance of the scene, even assuming that the CIECAM97s model is adequate for describing complex image appearance.

Too Many of the Wrong Tools

From this review, one may conclude that problems with color management do not stem from color space limitations, but from a lack of information about the physical meaning of the image data the various color spaces contain, and from a lack of tools for performing rendering to various media. If a transparency is scanned into sRGB by clipping the highlights and out-of gamut values, the problem is not the color space, but the adequacy of the rendering. The same issue would still exist if the transparency was scanned into sRGB-64 and then transformed to sRGB using the default transforms (which clip, because the default assumption is that the image data has already been rendered into sRGB).

The other way in which tools are inadequate relates to what is commonly known as "gamut mapping". Whether in relation to a colorimetric gamut or a gamut of appearance correlates, it is incorrect to assume that a single transform can be determined that gives optimal results with all images. For example, say that we have two scenes, one with mostly low chroma colors, and another with mostly high chroma colors. The optimal gamut compression algorithm cannot be the same for these two images if the difference in the source and destination gamuts is large. However, the optimal transformations will become more similar as the gamuts become more similar.

Proprietary Advantage vs. Open Systems

It is clear that a lot of work remains to be done in color management. Whenever this is the case, the result is a fundamental conflict between open systems and proprietary advantage. In a free market system, it is difficult to argue against maintaining proprietary advantage. One argument does hold -interoperability. Proprietary advantage must be sacrificed to the extent necessary to achieve interoperability if progress is to be made and the advantage realized. The challenge is to maximize both interoperability and proprietary advantage simultaneously. This can be achieved using current color management strategies and the reproduction model concept.

The Reproduction Model Concept

The basic idea behind the reproduction model concept is to distinguish the problem of determining appearance from that of producing a reproduction. This accomplishes two things: it is no longer necessary to have a fully functioning and robust complex scene appearance model, and, with a reproduction model, considerations relating to the output medium are factored directly in to the model. This means it is not necessary to develop "gamut mapping" algorithms, because the algorithms are built into the model, as opposed to being used in conjunction with an appearance model.

There are several advantages to combining media and appearance considerations together. The first is that it is no longer necessary to worry about the ability of the medium to produce a particular aim appearance when applying an appearance model. The second is that in many cases it is possible to construct a robust reproduction model that is much simpler than either an appearance model or the associated gamut mapping algorithms because some of the appearance and gamut mapping effects cancel each other out. The final advantage is that the development of competitive, robust reproduction models can provide an enormous amount of valuable information about human visual system processing. As proprietary models compete, the whole world becomes a giant subjective study. The enduring approaches should provide valuable clues about vision, assuming the vision clues can be unraveled from the media considerations.

For clarity's sake, the following is a brief outline of the differences between various processing models:

A Colorimetric Model

Input: patch colorimetry
reference white colorimetry
conversion to destination color space

Output: relative colorimetry in destination color space

A Single-Stimulus Appearance Model

Input: patch colorimetry
viewing conditions: proximal field, background, surround, luminance of adapting field, adopted white, reference white, D-factor

Output: patch appearance correlates

A Complex Scene Appearance Model

Input: scene colorimetry
adopted white for all scene points
viewing conditions for all scene points
human visual system spectral characteristics
human visual system spatial characteristics
human visual system processing model

Output: scene appearance correlates

A Complex Image Reproduction Model

Input: scene colorimetry
adopted white for all scene points
viewing conditions for all scene points
media characteristics
color rendering model

Output: a colorimetric description of a reproduction optimized for a particular medium and viewing condition.

It is interesting to note that a color rendering model could be an appearance model followed by gamut mapping. In this case a single stimulus model may be appropriate because the objective is no longer to correctly describe the appearance of a complex scene, but rather to use the appearance transforms as part of the means for producing the reproduction description.

Reproduction Model Examples

Eight color plates are provided at the end of this paper to illustrate some of the advantages afforded by proprietary reproduction models. These plates were produced on a short run press that was approximately, but not exactly, calibrated to print sRGB image data. The exact image data is available on the CD-ROM for printing in a more carefully controlled environment.

Color Management Approaches

Some solutions for employing reproduction models in color management are therefore as follows:

Conventional Color Management

The job of conventional color management would essentially remain the same as it is now - to take one reproduction description and transform it as appropriate for the intended application. The algorithms used by conventional color management would remain open and non-proprietary, so that everyone would know exactly what they were doing. This is essential for interoperability. It would also be possible to maintain two paths in conventional color management: a colorimetric path, where the profiles describe device value to colorimetry transformations and the gamut mapping is done by the CMM; and a perceptual path, where the input profile transforms input device values to the colorimetry of a specific reproduction description in the

PCS, and the output profile takes this description and produces appropriate, but not necessarily colorimetrically accurate, device values for the output device. The latter path allows the CMM to be "dumb", and is analogous to the sRGB workflow, except the image data is never transformed to the PCS reproduction description.

The assumption when using conventional color management would be that the reproduction description provided to the non-proprietary reproduction model was well behaved and falls into one of the three classes of reproduction descriptions described in the introduction. It is assumed that, with these classes of reproductions, a non-proprietary solution will perform reasonably well. The overall color management system has the option of whether to use proprietary or non-proprietary reproduction models or some combination. For example, a proprietary algorithm might be used to produce a PCS reproduction description of digital camera image data by producing an image specific profile. An output profile could then be used to transform the PCS description for the output device. If the output profile had a colorimetric rendering intent, the PCS description would need to describe the desired reproduction colorimetry on the actual output medium. If the output profile had a perceptual rendering intent, the PCS description would need to be a description of the reproduction that the output profile was assuming.

Proprietary Color Management

The job of proprietary color management is to produce reproduction descriptions that can be appropriately dealt with by conventional color management. There are three ways (and possibly more) in which proprietary approaches and conventional color management can be integrated:

Profiles on the Fly

With the "profiles on the fly" approach, a proprietary software application creates profiles as needed for use by a conventional color management system. This could be accomplished using either an external application that gets the profiles ready before use, or with a plug-in application that integrates itself with conventional color management. In either case, the application determines image specific profiles based on image information, input device information, and/or output device information. In some cases these may be device link profiles.

The requirements for implementation of this approach include: methods for communicating the device information necessary to create the profiles, an exact definition of what the PCS is, a way to integrate the appropriate plug-ins into the CMM (if the plug-in approach is used), and recognition of this approach as the recommended method of integrating proprietary reproduction models with conventional color management. A severe drawback of this approach is its inability to support proprietary algorithms where the color processing varies within an image.

Competing CMMs

The "competing CMMs" approach consists of a bunch of proprietary CMMs all using a compatible profile format (sort of what we have now). This approach can be cleaned up by requiring measurement based profiles and isolating the competition to the CMM. For this approach to work it will be necessary to standardize the image metadata carried with the profiles so that any proprietary CMM will have all the information it needs to function. The disadvantages of this approach are that a lot of work needs to be done to get really well defined colorimetric input and output profiles, and the fact that different CMMs will produce different results.

Two-Stage Color Management

This approach segments the processing problem. The first stage uses proprietary algorithms to produce an initial reproduction description. The second stage follows the conventional color management strategy to produce the desired appearance intent in another reproduction. The proprietary algorithms in the first stage can choose how much to defer to the well-defined and public second stage. Figure 1 contains a flow chart diagram of a two-stage color management as applied to digital photography.

Requirements for the two-stage approach include clarification of what the PCS is, communication of the image metadata necessary for the proprietary first stage to function, and mandatory inclusion of colorimetric output profiles so that the first stage can perform all the reproduction model processing if desired.

Conclusions

At present, our knowledge of how humans perceive color in complex scenes and reproductions is insufficient to construct a generic color management system that performs optimized color processing for all applications. It is sometimes necessary to use proprietary reproduction models. The issue of artistic intent makes it questionable whether it will ever be possible to completely avoid proprietary algorithms. However, several methods allow proprietary algorithms to be implemented in an open systems context. All that is necessary is to recognize that optimized color processing is image, media, and viewing condition specific, and that reproduction descriptions must be interpreted in the context of the media and viewing conditions for which they are intended.

Current manifestations of color management can fit nicely into a more global context. In fact, placing them in context helps clarify the nature of different implementations. However in some cases, proprietary algorithms and additional information are necessary to produce good results. There is a need to obtain consensus on the exact nature of an open systems structure that allows both proprietary and open approaches to be employed.

References

1. P.M. Hubel, J. Holm, & G.D. Finlayson, "Illuminant Estimation and Color Correction," *Colour Imaging in Multimedia*, Lindsay MacDonald, editor, Wiley, London, 1999.
2. J. Holm, "Exposure-Speed Relations and Tone Reproduction," *Proceedings, 47th Annual Convention of the Society for Imaging Science & Technology*, p. 641-648 (1994).
3. J. Holm, "Photographic Tone and Colour Reproduction Goals," *Proceedings, CIE Expert Symposium '96: Colour Standards for Image Technology*, p. 51-56 (1996).
4. J. Holm, "A Strategy for Pictorial Digital Image Processing (PDIP)," *Proceedings, IS&T/SID Fourth Color Imaging Conference: Color Science, Systems, and Applications*, p. 194-201 (1996).
5. J. Holm, "Issues Relating to the Transformation of Sensor Data into Standard Color Spaces," *Recent Progress in Color Management and Communications*, Robert Buckley, editor, IS&T, Springfield, VA, 1998, p. 193-197.
6. J. Holm, "Adjusting for the Scene Adopted White," *Proceedings, IS&T PICS Conference*, p. 158-162 (1999).
7. E.H. Land, "The Retinex Theory of Color Vision," *Scientific American*, p. 108-129 (1977).
8. X. Zhang, multiple papers at: http://white.stanford.edu/html/xmei/scielab/scielab.html
9. S.N. Pattanaik, M.D. Fairchild, J.A. Ferwerda, & D.P. Greenberg, "Multiscale Model of Adaptation, Spatial Vision, and Color Appearance," *Proceedings, IS&T/SID Sixth Color Imaging Conference: Color Science, Systems, and Applications*, p. 2-7 (1998).
10. IEC 61966-2-2 WD 4, Multimedia Systems and Equipment - Colour Measurement and Management - Part 2-2: Colour Management - Extended Precision RGB Colour Space - sRGB-X
11. ISO 17321 WD 3.1, Graphic Technology and Photography - Colour characterisation of digital still cameras (DSCs) using colour targets and spectral illumination, March 1999.

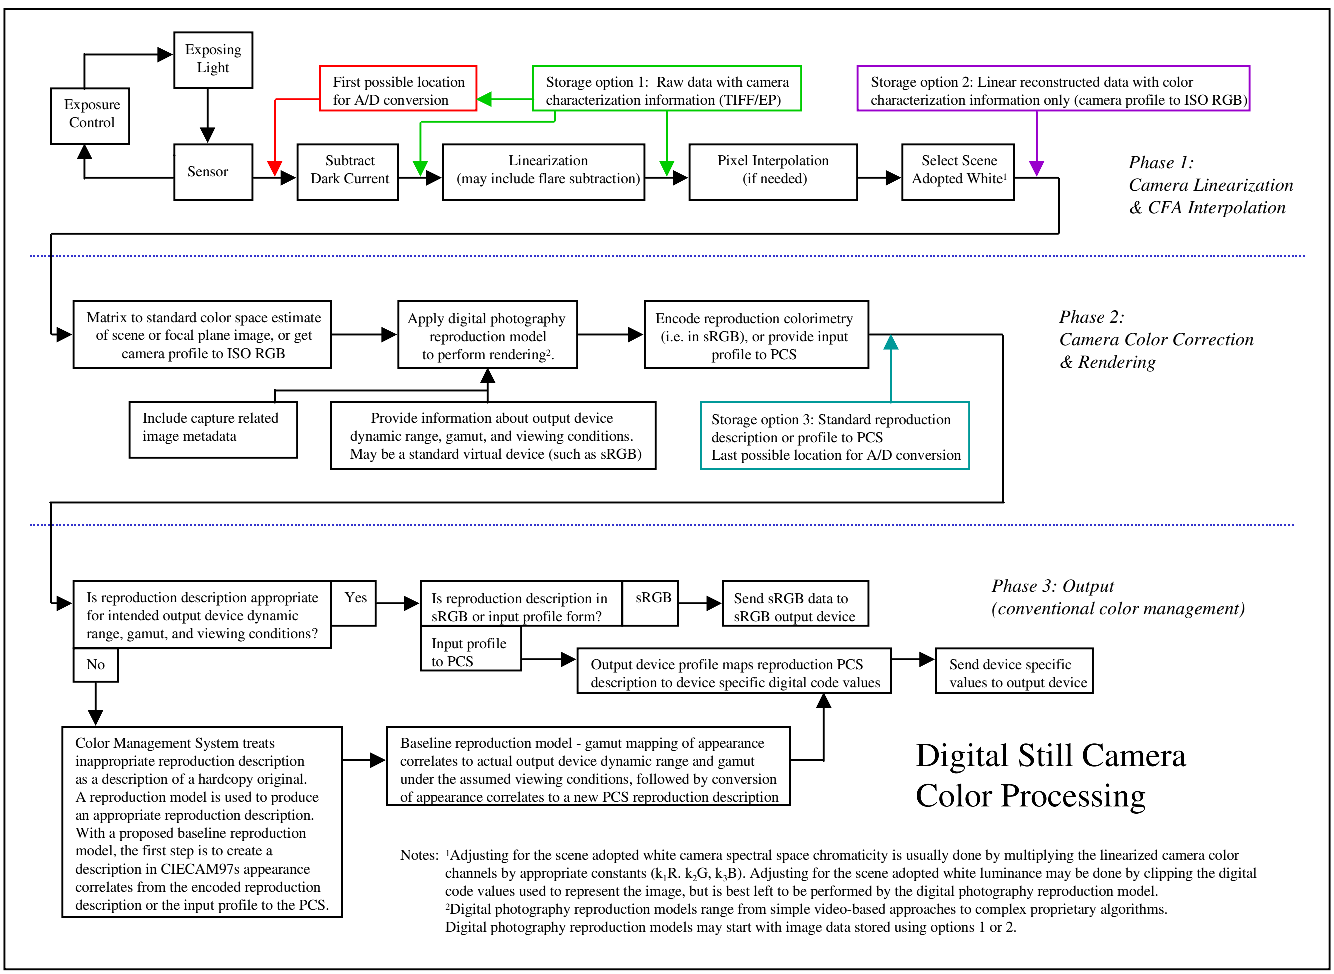

Figure 1: Flow chart diagram of Two-Stage Color Management

Plate 1a: A photographic reproduction model is used to capture a high-key scene. The reproduction appears too dark because the model assumes that the mean scene luminance should be placed at an 18% gray reflectance.

Plate 2a: A photographic reproduction model is used to capture a high dynamic range scene. Shadows are blocked up in the reproduction because the model assumes a scene dynamic range of 160:1, and this scene is 1400:1.

Plate 1b: The result obtained using a more sophisticated proprietary reproduction model. The model used still applies a global transform, but the nature of the transform is determined using several image based statistical parameters.

Plate 2b: A global transform determined using the proprietary reproduction model is applied. This transform compresses the dynamic range of the scene to that of the output medium.

Plate 3a: A simple video reproduction model is used to capture a 180:1 scene. Although this scene does not have an extremely high dynamic range, it is severely clipped because the model system gamma of about 1.2 limits the range captured.

Plate 4a: A colorimetric reproduction of a 24:1 scene. Even though the colorimetry of the print is representative of the original scene (system gamma of unity), the reproduction is not optimally pleasing.

Plate 3b: A global transform determined using the proprietary reproduction model is applied. Detail is maintained throughout the entire scene dynamic range and the color reproduction is still pleasing.

Plate 4b: A global transform determined using the proprietary reproduction model is applied. The more pleasing color reproduction may indicate that the human visual system includes scene dynamic range adaptation.