

Improving Color Image Quality in Medical Photography

Introduction

Medical images have played an important role in the development of modern medicine. They have provided a valuable means for capturing, storing, transmitting and viewing complex visual information. The introduction of digital detectors and display devices in areas such as radiography has made the communication and analysis of the most common of these images easier compared to hard copy film based images. Dedicated monitor specifications and display calibration tools (e.g. the Digital Imaging and Communications in Medicine Grayscale Standard Display Function, DICOM GSDF) were developed to ensure accurate and reproducible presentation of these grayscale images across different monitors to ensure that diagnostic accuracy would not be impacted by differences between displays [1]. Although radiology images sometimes use pseudo-color to enhance specific diagnostic features (e.g., blood flow in Doppler ultrasound), there are no inherently color images in radiology.

The prevalence of inexpensive color digital cameras, however, has dramatically increased the use of color images in other areas of medicine. The growth of telemedicine has further enhanced the utility of these medical records. Clinical specialties such as dermatology, ophthalmology, surgery, pathology, and gastroenterology regularly use visible light images in their practices. Increasingly a host of new imaging tools that acquire and display color images are being incorporated into clinical practice (e.g., fiber-optic based imaging tools to assess ovarian cancer).

The human visual system (HVS) is very effective in recognizing critical features by processing visual attributes such as brightness and color variations in an image. The medical industry has recognized the dependence of the HVS on the relative brightness of image features, and has implemented the DICOM GSDF standard to achieve a perceptually uniform scale for critical grayscale imagery [1]. However, this standard does not address color images. Several groups are currently developing proposals to introduce a medical color imaging process that is compatible with DICOM [2,3].

The color image processing methods and tools can differ depending on the intent of the content and/or the task of the viewer. In some cases (like false color maps in radiology), the need for accurate color reproduction may not be important since the actual colors are not a representation of the real world, but are merely used to highlight additional information that is related to the image. In fields like dermatology and pathology, the image colors provide valuable information that should be properly captured and displayed to the final viewer so they can render the most accurate diagnosis possible. This article addresses the needs of the latter case, and provides guidance to medical color image users for achieving the best possible color reproduction on a digital display.

Modern digital color image workflows can be generically described as illustrated in Figure 1. A more detailed description can be obtained in recent digital camera standards [4,5]. A digital camera captures the scene and applies adjustments to the captured colors that are specific to the camera, to produce more pleasing reproductions. The accuracy of the color information captured by the sensor can be affected by factors such as scene illumination and the camera spectral sensitivities. In simple point-and-shoot cameras, the camera will usually store the processed color image in standard compressed formats like JPEG and TIFF. These compressed formats typically include image enhancements encoded into the data, and the colors are normally processed to be encoded in a standard output-referred color space (such as sRGB) [6]. A color space is a standard numerical representation of colors. The basic process is illustrated in the top section of Figure 1. The bottom section gives possible color correction pathways.

The sRGB color space is typically used as the standard color space in digital camera image encoding since it is expected that the images will likely be viewed on a display that is calibrated for that color space (although that is rarely the case in medicine). Digital single-lens reflex (DSLR) cameras tend to give the user more control over how the color data is processed and formatted. The DSLRs usually offer the user the ability to store the image data in a proprietary RAW format, which is only minimally processed and has larger bit-depth. The RAW format is becoming more common in other digital camera technologies as well. Since image enhancements can make it more difficult to color-correct an image back to accurate scene colors in post-processing, the RAW format can have some advantages. The camera manufacturer's RAW file format is generally specific to the manufacturer, and usually requires the manufacturer's software to process the image. The software may not support processing to accurate colors. However, some open source and third party software can extract the image data from the RAW image files and transform them into images that do attempt to accurately encode scene colors. Using color management, these "scene-referred" images can be viewed on display which enable reasonably accurate viewing of the scene colors.

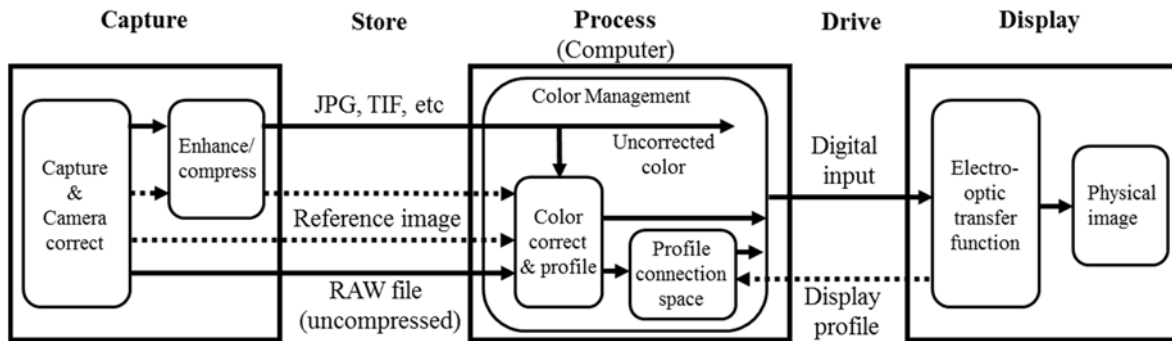


Figure 1: Functional flow diagram of a generic digital color image process, from image capture to a computer monitor. The concept of the reference image is used for workflows where color correction is applied.

Before an image file can be viewed, it must be interpreted and rendered properly on a display. In the simplest case, the image viewing software recognizes the compressed (or uncompressed as the case may be) file format and directly drives the display input. The rendered output image will be determined by the unique electro-optic transfer function of the display. If the display was calibrated to the same color space as the encoded image data (e.g. sRGB), then the image colors will be rendered properly to the viewer. However, since sRGB images have typically been enhanced to produce pleasing reproductions, directly rendering sRGB images without color-correcting can yield significant color errors. Fortunately, the accuracy and reliability of the color images can be significantly improved by employing methods used by professional photographers. It is common practice to place a reference color chart (e.g. Macbeth color chart), with well-characterized colors, in the same scene as the object to be photographed [7,8]. The digital images of the reference color chart taken under the same conditions as the intended object of interest can be compared and used to color-correct the image data. The color correction can be implemented by directly creating a new image file with the corrected color data transformed to a standard color space (e.g. sRGB) for later viewing. Alternatively, the necessary transformation needed to color-correct the original object image can be saved separately as a color preset or profile. This preset or profile must then be applied to the original image prior to being rendered by the display. This open-loop process can work for a fixed image capture and display viewing environment, which is difficult to obtain in a clinical context. However, a more flexible closed-loop process can also be utilized using the open source ICC profile methodology [9]. The ICC framework uses a virtual profile connection space that transforms the color-corrected image to the proper color space used by the output device, either printer or display.

Given the above background information on digital camera color image processing, we will further highlight the factors that contribute to color errors, discuss proper camera and lighting setups for improved color reproduction, review the use of color charts, and dive deeper into the color correction processing. In addition, a summation of our findings is given in terms of a recommended procedure, which describes the industry best practice for improving the rendered color accuracy.

Factors that contribute to color error

As suggested in Figure 1, the image colors viewed on the display can be affected by the actual image capture setup, the image processing, and the physical rendering on the display. Recent research compared the relative contributions of the image color error produced by a point-and-shoot camera, the gamut mapping of the color management system, and the error measured from the display with a factory calibration [10]. The study found that the color errors introduced by the camera capture process were dominant. In addition, since the image capture was at the beginning of the workflow, these color errors would propagate and likely be enhanced by the other error sources. Therefore, taking precautions to improve the image capture process will have the greatest benefit toward minimizing the final color errors viewed on the display. Toward that end, we first identify the factors in the image capture process that produce the most significant color errors.

Lighting uniformity

An important factor that can impact overall image quality as well as color, is the part of the body being imaged. A full body skin scan is feasible, but it may not always show all skin lesions and surfaces with sufficient detail. Enhanced lighting, multiple images and several angles may be helpful. The scalp and other areas with a significant amount of hair may need to have the hair physically displaced or removed, and special lighting may enhance viewing conditions. Mucosal lesions and orifices, including genitalia, often require special attention to lighting and exposure in order to allow examination, especially with respect to reflections off shiny, moist tissues that could impact the colors rendered. The color of the patient's skin itself in combination with the lighting and background conditions may change the color of photographed skin lesions. Therefore, the quality of lighting on the subject area is critical. Directional lighting may be useful to reviewed texture, depth, and morphology. However, this can lead to glare and non-uniform illumination over the image. It has been shown that non-uniform lighting can strongly influence the resulting color error in the image [11].

Spectrum of lighting

The surface color we perceive is dependent on the spectrum of light illuminating the surface. Figure 2 illustrates the difference in the colors of the patches of a Macbeth color chart when illuminated using tungsten and daylight illumination. Therefore, if the illumination spectrum is not the same for each camera capture session, or the spectrum changes over time, then it will be more difficult to compare the colors in images taken under those different illumination spectra. While the photographer can set the camera to adapt to the illumination white point, this is often not sufficient. Previous work has shown that the ability of a camera to adapt a changeable spectrum to a desired white point depends on the camera technology and the spectrum of the light [12]. The performance tends to be best when the light spectrum is not changed. In addition, using illumination with a single continuous broadband spectrum of light is better than using lights with different spectra.



Figure 2: The colors for the patches of a Macbeth color chart with tungsten illumination (top row) and daylight illumination (bottom row).

Camera technology

One would expect that the quality of the images would be dependent on the camera technology. In many cases, color errors can be reduced by using professional camera technology, for example using a DSLR camera instead of a cell phone camera [12]. However, the trend may not always be consistent. The color accuracy depends on the accuracy of the white balance used by the camera, on the camera spectral sensitivities, and on the quality and intent of the processing applied by the camera. Professional cameras tend to be more controllable, but cell phone cameras often use very sophisticated algorithms that in some cases will produce better results for inexperienced users. RAW files offer the greatest degree of control and flexibility of processing, but often the benefits are only realized when processed by experienced professionals.

Subject color

The camera color error generally depends on the spectra of the colors in the subject to be imaged. Although some studies have shown that more saturated colors (like red and blue) tend to have larger color errors, flesh tones can also produce substantial color errors [12]. The HVS is especially sensitive to memory colors, like flesh tones. A study on artificial skin found that CIELAB color differences of $\Delta E^*_{ab}=1.1$ were perceptible and $\Delta E^*_{ab}=3.0$ were at the acceptability threshold for light specimens, and $\Delta E^*_{ab}=1.6$ were perceptible and $\Delta E^*_{ab}=4.4$ were at the acceptability threshold for dark specimens [13,14]. A value of $\Delta E^*_{ab}=1$ is generally considered a just-noticeable difference between colors. This color difference values provide useful guidelines for the amount of color errors that may be tolerable for medical photography.

In addition to the above-mentioned factors, standard photography best practices require that the region of interest that is imaged be properly focused, framed, and have the proper exposure (not over or under-exposed). Figure 3 illustrates some common examples of poor photography technique, some of which can affect the color quality of the image. The top left image has bad white balance, the top right was over-exposed, the bottom left suffers from non-uniform lighting, and bottom right has improper focal depth.



Figure 3: Common examples of poor photography technique.

Achieving consistent color

Prior to offering recommendations on how the image color quality can be improved, we first put forth a set of operational use objectives for the medical photography workflow. The objectives emphasize the need for consistent color reproduction for the evaluation of medical images.

Table 1: Main operational use objectives for medical photography.

1	Visually inspecting a region of interest in a single image in a consistent manner
2	Measuring properties of certain areas in a region of interest of an image in a consistent manner
3	Visually comparing regions of interests in one or several images in a consistent manner
4	Comparing measurements of certain areas in regions of interest in one or several images in a consistent manner.

This is not an exhaustive list, but many other tasks and objectives can be automatically enabled by these 4 items (e.g. good documentation and record keeping, follow-up over long periods, telemedicine, teaching, etc...). The term 'consistent' is admittedly vague and needs to be clarified for a specific application. From the perspective of maintaining the image color integrity, this article strives for color consistency by avoiding or eliminating any factors that can change the perceived color of a patient's image which could affect the interpretation of that image. Although the perceived image color can be affected by the display system and the viewing environment used to view the display, our discussions are limited to the impact of the image capture process and any subsequent color corrections.

In practice, the color inconsistency caused by lighting variability when taking images of the same patient's body part over time can make it difficult to gauge the evolution of the medical condition seen in the images. The same can be said when using different cameras, changing camera settings, different lighting spectra and geometry, etc... The severity and impact of those confounding factors can obviously be very variable and highly dependent on the patient's body part and medical condition. This is illustrated by the examples shown in Fig. 4. For the images of the patient in the top row of Fig. 4, it is not completely clear whether the patient's skin has changed, or that the camera setting and lighting are different. In this case, the patient's skin has indeed changed, but not nearly as much as would be suggested by the image. The background grays are different for each image, which makes it difficult to assess the relative amount of change. The bottom images in Fig. 4 show the case of atypical nevi for skin cancer detection. Not only does the variance in the photography technique between images make it difficult to monitor the change in the cancer, but it is even difficult to know what is the real skin phototype of the patient.



Figure 4: Examples of possible confusion due to photography. The images in each row are of the same patient.

Color is very important, and this places high demands on the consistency of colors in the acquired images. Similarly, when evaluating ulcers, the measured area of an ulcer can be important, and thus the physical resolution in the region of interest must be known so that areas can be consistently compared. It should be noted that consistency does not necessarily mean faithfulness. It is possible to be consistent, and thus comparable, without being an exact copy of how the scene looked to the photographer during acquisition. For example, image colors can be modified to represent the scene as adapted to a certain predefined and fixed light source.

In general, the colors rendered to the viewer are typically either scene-referred or output-referred. The purpose of scene-referred images is to accurately encode the

original scene colors. In contrast, output-referred images encode the preferred image colors for the output device (the display or printer).

In order to enhance the quality and utility of medical photography, it is valuable to ensure colors are captured and rendered in a standard way. This enables all the objectives that involve comparisons (items 2-4 in Table 1), and it usually also makes inspections easier (item 1). One way to achieve this would be to use exactly the same standard set of conditions during the image acquisition (lighting, exposure, calibrated camera, geometry, etc...). This may be difficult in practice.

Another way of achieving this is by using standard hardware, and by calibrating the images to a standard set of conditions after the acquisition. This typically involves adding some kind of color chart with known colorimetric properties in the scene during acquisition. In addition, the color chart can also serve as a scale marker, allowing the accurate determination of dimensions for areas in the imaging plane of the chart.

One of the most important and often challenging factors in ensuring proper photography and consistent color data is homogeneous lighting (see Fig. 3, bottom left). It is clear that if the lighting over a region of interest is uneven, it will be very difficult to obtain sensible color measurements from it, and no amount of calibration is going to help. For accurate color images, uniform diffuse lighting tends to work best. Some examples of diffuse lighting setups are shown in Figure 5. But even with a uniform light source, lighting can still be uneven if the object geometry is not considered (oblique angles, shadows, etc...). It is best to ensure that the region of interest is nominally perpendicular to the optical axis (the line going straight from the camera sensor through the lens to the scene). Clearly this can be quite challenging if the region of interest is not flat (see Fig. 6, middle image in the bottom). Also with wet or moist areas in the region of interest this can lead to specular (mirror-like) reflections, visible as very bright white regions (see also left image in Fig. 6). Normal flash photography tends to make this problem worse, although a special ring-like flash for close-up photography may alleviate this because the incoming light from the flash is coming from the sides. When present, specular reflections can usually be diminished by controlling the polarization of the light, or angling the light source away from the optical axis. Figure 6 also provides a good demonstration of how the white point can be changed through proper color correction. An illustrated guide for medical photography (with color chart) is given in Fig. 7.



Figure 5: Example of diffuse lighting setups using commercial softbox lighting (left), or a homemade lightbox with diffuse walls (right).



Figure 6: Color calibration using a color chart. Top row shows the uncalibrated images, and the bottom row the corresponding images calibrated to a consistent white point. Images courtesy of Dr. S. Van Poucke, Ziekenhuis Oost-Limburg, and the University Hospital Ghent, Belgium.



Figure 7: Illustrated photography guidelines. Top row: good image, image with uneven lighting. Second row: chart not in the same region of interest (toy house). Third row: chart not in view, poor framing (too small). Fourth row: over-exposed and under-exposed image. Bottom row: scene not perpendicular to optical axis, and chart is angled.

As mentioned previously, a good approach to achieving consistent colors is to capture the image using a standard illumination, and calibrate the image to a standard color space after the acquisition. The result is known as a scene-referred image. An important aspect of this is to adapt the white point of the scene to that of the image color space. CIE D65 simulates a daylight spectrum with a white color that is commonly used for color spaces in digital image files and for display technology, and would serve as an obvious standard white point [6,13]. Although it is difficult to accurately realize the D65 spectrum in practice, the color calibration to this idealized spectrum will be better if the actual light source used during the image capture at least approximates this spectrum [12]. It is important to note that the actual color of a subject will depend on the illumination with which it is viewed, so if the spectrum of the capture illumination changes, the colors seen by the camera will change, even with a properly calibrated camera.

Another factor that can help achieve color consistency is the degree of control offered by professional (DSLR or similar) camera technology, particularly RAW files and processing software. Many professional cameras can have the ability to choose a scene white point that more accurately represents the actual scene, which can then be adapted to the standard color space in the subsequent calibration process. These cameras tend to also support the use of RAW image formats. Using RAW camera data instead of the more traditional JPEG workflow can increase the quality of the final color corrected images because some RAW processing applications can be set to avoid the image processing used to perceptually enhance the image. The higher bit-depth and dynamic range of RAW camera data also lessens problems with data saturation (values at the minimum or the maximum of the possible range) due to poor exposure settings or light sources that are very yellow or bluish. Raw processing also allows one to adjust the exposure, white balance and color correction of the image after the image has been captured.

In addition to the more technical aspect of the image capture process, color consistency can in general be improved by using good photography practices. This includes removing distracting jewelry and clothing prior to acquiring the images, as they are distracting visually but could also overlap lesions and potentially distort colors. Although using flash lighting can sometimes eliminate shadows, it can cause white out and distort the colors. Uniform diffuse lighting is generally preferred for color consistency.

Color Correction

For color consistency, our goal for color correction is to display the imaged scene as it looked under the light source used, but chromatically adapted to a standard white point such as CIE D65. The viewing and colorimetric measurements of the resulting scene-referred images will be more consistent over time. This enables qualitative (visual) and quantitative (physical measurements) comparisons to be made, and is more forgiving of different setups. And by adapting to a standard CIE D65 white point, it is easier to render the images on commercial displays. However, in order to adapt the image to the CIE D65 white point, the illumination spectrum needs to be sufficiently similar to D65 so that the camera captures enough colors for the color processing to make an accurate adaption to the target white point.

The color calibration will typically transform the image to a standard color space called sRGB, or provide an ICC input profile to implement correct interpretation [6]. Since sRGB tends to be the native color space used by modern displays, the sRGB image will often display correctly without any further color management (see left side of Figure 8). For calibrated displays or printers that have an ICC output profile, more accurate color rendering may be possible (right side of Figure 8). If the acquired image has a sRGB input ICC profile associated with it, it can then be converted to render the image with potentially greater accuracy on the actual display, with the added benefit of not having modified the original image data.

In order to correct the original image colors to what they should be in the sRGB image file, professional photographers may use a reference color chart. The color chart can be placed side-by-side with the region to be imaged, or it may be imaged before or after the photography session. In the latter case, care must be taken to replicate the same placement and lighting/camera conditions. The image data from the reference color chart can then be used to transform the patient image to scene-referred colors. The center “colorimetric measurements” box in the Figure 8 workflow highlights that the known color data for the reference color chart is used to calibrate the original image. Commercial and open source software is available to perform this color correction, with varying results [12]. The flowchart in Fig. 8 illustrates the main components of the color correction process. The number of stars at the bottom of the figure give a rough indication of the quality of the final displayed image with respect to the objectives outlined in Table 1. Usually the quality of the results increases further if RAW camera data is used instead of the typical JPEG images. This is indicated by the extra red star. A more detailed flowchart of the color correction process, with the color chart in a side-by-side or sequential procedure, is given in Appendix A. A detailed recommended color image capture and color correction procedure is also given in Appendix B.

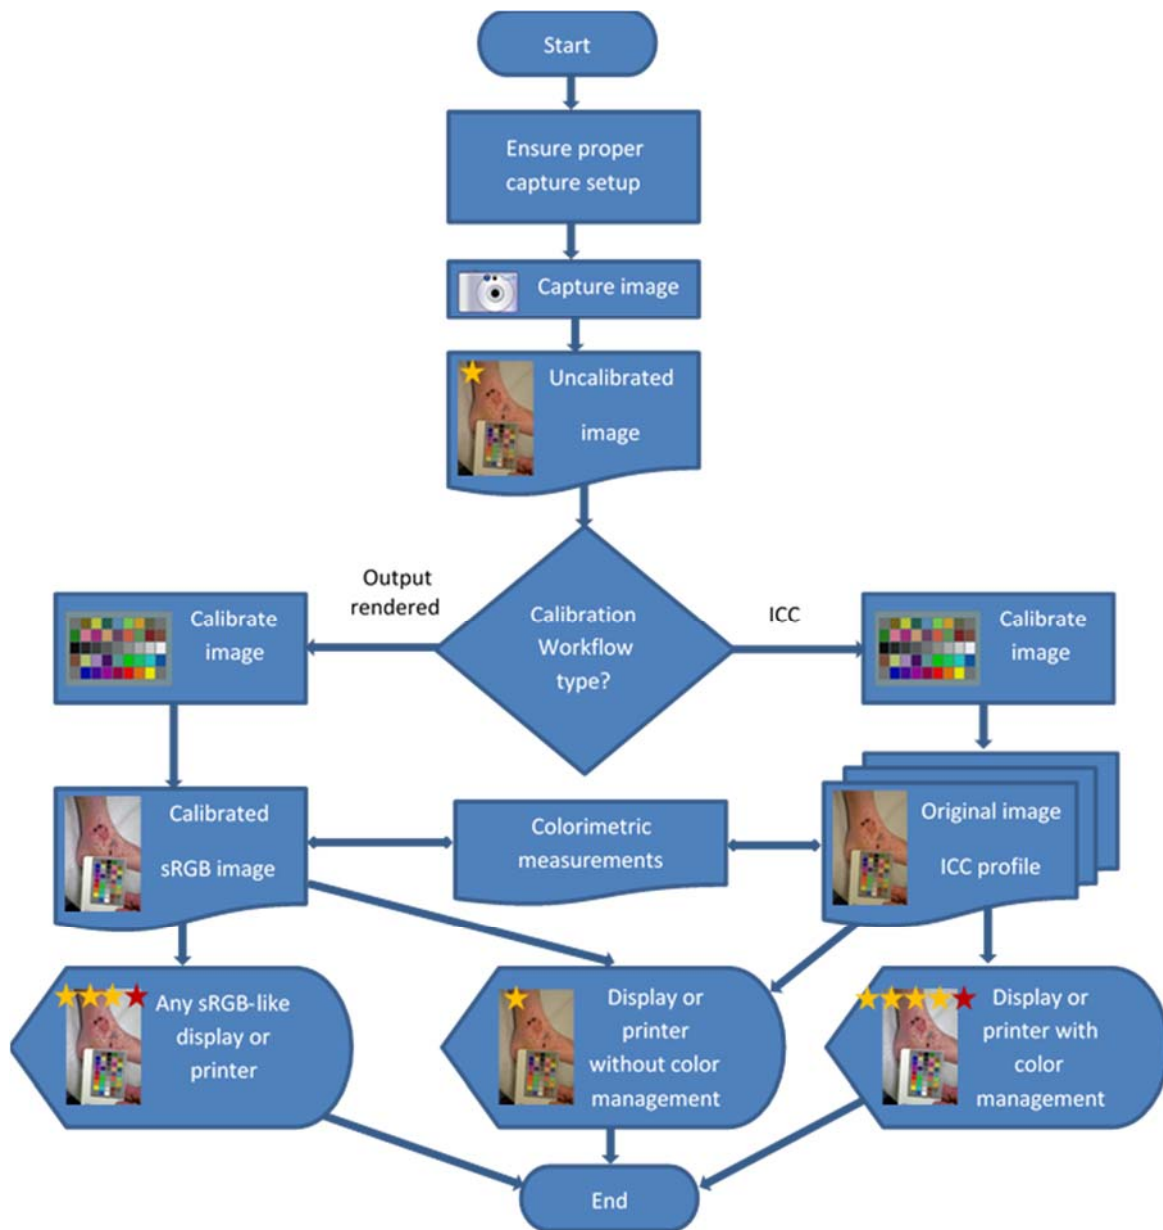


Figure 8: Functional flow diagram of the color correction process, with either direct output-referred rendering or the use of ICC profiles.

General color management considerations

ICC color management provides a basis for consistently communicating the color of the original image throughout subsequent operations of display, archiving and reproduction. The ICC architecture, first published in 1994, is used in most color imaging tasks, especially in professional workflows in photography and graphic arts.

As indicated in the previous section, it is recommended that medical images are processed to scene referred. Then an ICC profile may be assigned, or for convenience they may be converted to a display referred color encoding. sRGB is the most widely used encoding for this purpose, and medical photography subject matter will normally fall within the color gamut of sRGB. In situations where images

are to be printed commercially or displayed on extended-gamut monitors, it may be necessary to convert to a different encoding such as Adobe RGB (1998).

When reviewing images on the display, it is essential that the display be properly calibrated. Although the sRGB color space has been in monitors for a long time, it is not always properly implemented in consumer grade displays. For professional or medical displays, the displays are calibrated according to sRGB specifications, but the calibration only holds for a period of time. Nowadays, off-the-shelf calibrators are widely available, and color management monitors are more accessible. It is therefore recommended for users to calibrate their displays before reviewing color critical images. Some suggested calibration parameters are:

- White Point: D65
- Luminance: 160 cd/m²
- Gamma: 2.2
- Contrast: Panel Native

If the image was captured in sRGB, then it is suggested to calibrate the display to sRGB. But if Adobe RGB was used to capture the image, then it is suggested to calibrate the display to Adobe RGB.

In many simple workflows, especially those used in consumer applications, it is assumed that images are to be interpreted as being encoded and displayed in sRGB, in which case no conversions are required and although it is always recommended to embed the source profile it may not be essential. However, where an image may have a range of purposes, it is recommended that the profile which represents the color encoding of the image (most often sRGB) is embedded by assigning this profile before saving the image. Subsequently it may be required to convert the image so that the same colors are reproduced in a different color encoding. In this case the image is converted using the embedded profile as source, the desired color encoding as destination, and the Media-Relative Colorimetric rendering intent to ensure that color fidelity is preserved relative to the media white.

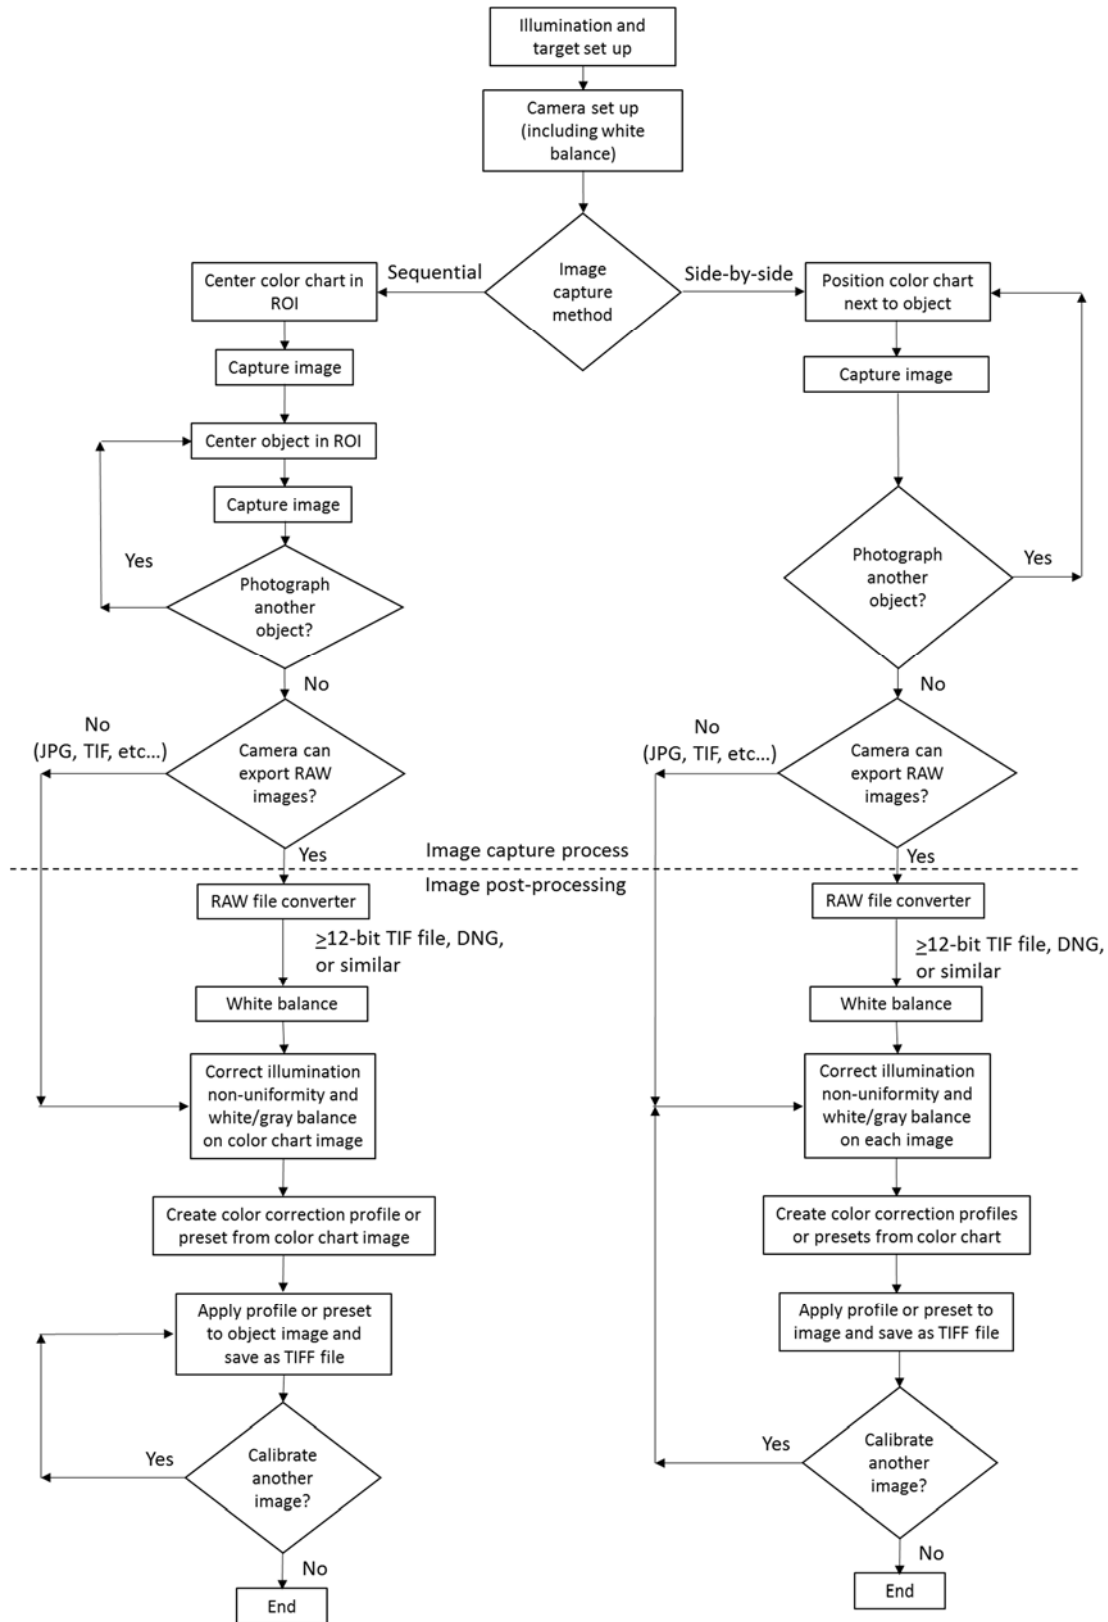
Further color management considerations are given in Appendix C.

Conclusion

Research has shown that a typical image capture workflow for medical photography can yield significant color errors compared to acceptability thresholds for skin tones [12,14]. The more automated technologies, non-uniform lighting and an illumination spectrum that substantially differs from daylight are key factors that contribute to the color errors. These factors can be minimized by the proper implementation of camera technology, and using uniform diffuse illumination that approximates the daylight spectrum. Further improvements in color accuracy can be achieved by applying color correction methods developed for professional product photography and digital archiving, which utilize color reference charts. In developing a medical photography workflow with color correction, we submit that there is great value in emphasizing color consistency between image capture events. By adapting images to a standard CIE D65 white point, the viewer can more easily compare the evolution

of color images rendered on common displays. Detailed workflows and procedures are proposed that realize the concept of standard output-referred rendering for achieving color consistency. If the full guidelines cannot be followed, then a basic set of steps is recommended in Appendix D to help achieve a minimal level of color consistency. The implementation of these procedures and workflows is best executed through closed-loop color-managed systems utilizing color profiles, such as ICC profiles. The application of these color correction methods are expected to improve diagnostic search times and outcomes. However, greater automation of these methods is needed in order to obtain greater acceptance within the medical community [26].

Appendix A: Flowchart of camera image capture and color correction workflow



Appendix B: Recommended color image capture and color correction procedure

This general procedure outlines a recommended digital camera image capture workflow that can be used to improve image color accuracy and consistency. The process is outlined in the flowchart given in Appendix A. The implementation of this workflow would be especially beneficial for use cases where color accuracy is critical, such as dermatology, plastic surgery, pathology, and wound documentation. It should also be noted that since medical photographs are part of a patient's record, they are subject to privacy considerations (e.g., HIPAA). However, most healthcare organizations include a statement in the consent that patients sign when they agree to the medical services being provided, that says they are aware that photographs may be acquired and used as part of their routine medical care and will be included in their record. There is much debate recently, however, as to whether photographs acquired by patients and transmitted to their physicians should become a part of the medical record.

Required equipment:

- Digital color camera with white balancing capability and a minimum resolution of 3 megapixels (MP).
- Reference color test chart. May be a commercial color chart (e.g., from X-Rite, DSC Labs, QPcard, Douglas color card, etc...) or one designed for the application. The color chart should come with the corresponding measured spectral reflectance data, and the chart patch color data should be calculated for the light source to be used.
- Light source and background that can provide uniform hemispherical illumination over the camera field of view. The light source should produce spectrally smooth broadband white light, approximating daylight. Spectrally "spiky" spectra can produce problems.
- Color correction software that can recognize each color in an image of the reference color chart and create a colorimetric calibration profile (HSL Preset file, DNG or ICC profile, or similar), which can be used to color calibrate an image of an object photographed under the same conditions as the reference color chart [15]. Color correction software that does not save calibration files should embed the calibrated RGB values in the image, and export the image file with a tag corresponding to the appropriate standard color space (e.g. sRGB).

Desirable equipment:

- Digital color camera capable of exporting RAW image files, and the ability to perform an in-camera white balance. The camera and lens combination should be flat-field corrected to within 2%.
- A RAW file decoder/converter which is able to import RAW images and export them as ≥ 12 -bit TIF or DNG format. Commercial software (e.g.

Adobe camera RAW, Capture One, Phocus, etc..) is available, as well as open source software (such as Dcraw).

- The software should be able to open RAW files and perform corrections for illumination non-uniformity, white/gray balance, and color.
- Alternatively, color correction software may be used that provides the ability to create ICC profiles to calibrate images to adapted scene colors. Commercial ICC-aware viewing software is available from several companies, in addition to free software (e.g Irfanview and GIMP).

Procedure:

Image capture

1. Set up the illumination and background for photographing the object of interest. The background should be a uniform matte color, ideally a gray with 18% reflectance. The camera field of view, should be adjusted so that it does not extend beyond the gray background. This field of view should be fixed for all photographs.
2. The light source should produce uniform diffuse hemispherical illumination over the field of view, with special attention paid to the lighting uniformity over the image area where colors will be evaluated. This will minimize glare, specular reflections and errors arising from lighting non-uniformity. Examples of diffuse lighting configurations are given in Fig. 5.
3. The object (e.g., body part) of interest and/or reference color chart will define the image region of interest (ROI). For the side-by-side method, the ROI is defined by the object of interest and the color chart placed adjacent to it. In the sequential method, the ROI is defined by the object of interest or the color chart, whichever is larger. The choice of method to some extent depends on the part of the body being imaged and the comfort of the patient. Place a uniform diffuse (ideally 18% reflectance) target in the image plane at the ROI. If the gray target is large enough to fill the entire ROI, then it may be used to compensate for illumination non-uniformity during the image post-processing.
4. Position the camera in front of the gray reference and align the camera so that its optical axis is centered on the gray reference and perpendicular to it. The image ROI should be contained within about half the field of view of the camera. If the sequential method is used, it is best to use a tripod, or similar mechanism, to hold the camera stationary for the remainder of the photographs. If the side-by-side method is used, then a fixture similar to that shown in Fig. B.1 can be used. The side-by-side method is preferred if the illumination is not stable.
5. Use the in-camera white balance function to determine the proper white balance for the given lighting condition, and maintain this white balance

setting for all subsequent photographs. Some cameras have a Preset Manual or Custom white balance mode to obtain and hold that white balance setting. Omit this step if the camera does not have in-camera white balance capability.



Figure B.1: Example fixture used for the side-by-side image capture method.

6. Capture the image of the gray reference in the ROI. If the illuminance is not uniform in the ROI to within 5%, an illumination non-uniformity correction should be applied in the image post-processing. This correction is only valid if the camera setting and lighting conditions are held constant.
7. Place the reference color test chart in the focus plane of the ROI, so that the camera field of view captures all of the colors in the chart. For the sequential method, the optical axis of the camera should be centered on the chart and perpendicular to it. For the side-by-side method, the edge of the color chart is positioned near the center of the camera image (see Fig. B.2). Photographic test charts (such as ColorChecker SG) can be used, although ideally patches should be matte rather than gloss. Custom charts with patches constructed to be similar to the subject of the photography can also be used (e.g., PANTONE SkinTone™ Guide from X-Rite or Douglas color card may be used for skin tones).
8. Set the camera exposure so that the lightest color patch in the test chart is approximately 90% of the camera saturation white.
9. For the sequential method, capture the image of the reference color test chart and export the image in RAW file format, if the camera is capable. Where possible, use a “neutral” mode RAW capture setting, which minimizes any camera visual enhancements. Replace the reference color test chart with the first object to be photographed, center in the image, and capture the image of the target object. Repeat the image capture of subsequent objects in turn (see Appendix A). Export the images in the same RAW file format. The lighting conditions and camera settings should not be changed. If the camera cannot export RAW files, set the camera to use the

- highest quality (least compression) image, use low ISO values, and export images with a tag corresponding to a standard color space (e.g., sRGB).
10. For the side-by-side method, place the color chart adjacent to the object of interest (see Fig. B.2) and capture the image using the “neutral” mode RAW capture setting. Export the image in the RAW file format if possible. Replace the first object of interest with other objects in sequence at the same focus plane. The lighting conditions and camera settings should be unchanged. If the camera cannot export RAW files, set the camera to use the highest quality (least compression) image, use low ISO values, and export images with a tag corresponding to a standard color space (e.g., sRGB).



Figure B.2: Example alignment of the side-by-side image capture method.

Color correction

1. For RAW files, use a RAW image converter/decoder to extract the image information in all files and save them in a standard image format (e.g., ≥ 12 -bit color TIF, or similar files). The file should include the desired white balance.
2. If an illumination non-uniformity correction is deemed necessary, apply the uniformity correction to all reference color chart and object images.
3. Open the image of the reference color chart (for the sequential or side-by-side method). Use the image editing program to ensure that the gray levels are scaled correctly. The gray level scaling will depend on the reference color chart used. However, it is common to use a reference color chart where the whitest color patch is set to an exposure of 90% of the encoding maximum, or RGB= 230, 230, 230 for 8-bit sRGB color images. The darkest patch is then set to an exposure of 4%, or RGB= 10, 10, 10. If the black patch is below this level, then use the current setting or reshoot the photograph with brighter illumination. For the sequential method, the gray level scaling applied to the reference color chart is also applied to all object images taken under the same shoot conditions.
4. The color correction software should automatically find the centers of each color patch of the gray level-scaled reference color chart image, and create

a Hue, Saturation, Luminance (HSL) Preset or color calibration profile (DNG, ICC profile, or similar) based on the known color values of the reference chart. It is recommended that ICC profiles also be created, if it is not already the primary color correction pathway.

5. For the side-by-side method, apply the HSL Preset or color calibration profile to the image and save the new color-corrected image in the desired format (e.g., a high quality TIF file). Repeat the gray level scaling and color correction for each side-by-side image. An example of a color-corrected image is shown in Fig. B.3.



Figure B.2:: Example of color-corrected image using Figure 3 following the side-by-side method.

For the sequential method, import the other photographed objects of interest into the image editing program that is capable of using HSL Presets or color calibration profiles. Apply the HSL Preset or color calibration profile to each image and save the new color-corrected image in the desired format (e.g. a high quality TIF file).

Appendix C: Further color management considerations

ICC color management is universally used to handle transforms between devices and color spaces, as it provides a robust, well-proven and flexible architecture that is widely implemented. One key reason for the success of the ICC profile is the use of a well-defined Profile Connection Space (PCS), which is in effect a virtual color space that all ICC profiles connect to [9]. Since different media have different dynamic ranges, they cannot be matched exactly and so the default procedure in ICC color management is to scale all data to be media-relative. When transformed via the PCS, this has the effect of matching one media white point to another, regardless of its luminance or chromaticity. The case where an exact colorimetric match is required is supported relative to a perfect reflecting diffuser through the ICC-Absolute Colorimetric rendering intent. However, ICC v4 display profiles are required to use the display white point as the adopted white point, so there will be no difference between the media-relative and ICC-absolute transforms in these profiles [16]. For HDR displays or other cases where the display white point is not the adapting white point, ICC v2 display profiles or v2 or v4 output profiles must be used.

Image state and rendering

In the majority of cases, images captured by camera will be enhanced (or ‘color rendered’) so that they appear pleasing when viewed on a display. This may be done by in-camera firmware when saving images in JPEG format, or subsequently in a RAW convertor when processing camera RAW images. Images which have undergone this color rendering process for display viewing (or printing) are said to be ‘output-referred’, while images which are encoded so that the original scene colorimetry is preserved are said to be ‘scene-referred’ [5]. The methods used to render to output-referred image state (such as in-camera processing to produce JPEG images) are generally proprietary and undocumented, which makes it difficult to obtain accurate colorimetry from an output-referred image. If the rendering method is known, it may be possible to invert the transform to recover the original scene colorimetry. However, in some cameras the color rendering processing is scene-specific, in which case the original scene colorimetry can only be recovered using side-by-side chart capture.

Supported transforms

The ICC architecture incorporates processing elements that perform the conversion between data color encoding (such as RGB camera data) and the PCS. These elements include single-channel curves (which can be defined as a gamma value, look-up table or function), matrices and multi-dimensional color look-up tables (CLUTs). In ICC v2 and v4, profiles can be constructed from just a curve per channel and a matrix, or as a series of curves, matrix and CLUT [9].

Scene-referred colorimetry with ICC profiles

When the goal is to capture scene colorimetry, two approaches are possible using ICC color management [17].

1. Custom profile for camera/illumination
2. Standard scene-referred profile

A custom camera profile converts from camera RGB to the ICC PCS. It has to be built for the particular camera (including its exposure settings, lenses etc) and scene illumination. Examples and further information can be found on the ICC web site in the Digital Photography section and the Profiles section [18, 19]. The working space should be set to ProPhotoRGB to reduce clipping, and image data should be 16-bit precision. The media white point can have a maximum of 200 cd/m², twice that of the PCS, which should be adequate for most medical uses although may lead to inaccuracies when very high luminances are captured.

For many applications where scene-referred colorimetry is the goal, it is preferable to convert images to a standard scene-referred profile. If the image data is 16-bit, the linear_RIMM-RGB_v4.icc profile is recommended for this purpose. This profile, and the suggested procedure for converting to it, are documented on the ICC site [20].

The use of the different rendering intents in ICC profiles for different reproduction goals is described in references [17] and [20].

Standard output-referred encoding with ICC profiles

Images in medical photography are created primarily for viewing by medical practitioners on displays. Most displays do not support direct viewing of scene-referred image data, due to differences in scene-referred encoding and display encoding.

An output-referred encoding is one which can be realized on a given reproduction system, and several standard encodings of this type exist. The ICC 3-component encoding registry provides full details of the available standard encodings [23]. The most widely used are sRGB [24] for CRT-like color gamuts, and Adobe RGB (1998) [25] for extended gamuts.

To convert to the chosen standard out-referred encoding, the recommended procedure is to perform a conversion from scene-referred (with the linear_RIMM-RGB_v4.icc profile as source) to the profile for the output encoding, using the Media-relative Colorimetric rendering intent. Profiles for sRGB and Adobe RGB (1998) encodings are widely available, but ICC provides recommended v2 and v4 profiles for sRGB.

Workflow

If a custom profile is generated for a camera capture condition as described above, it can be assigned to an image so that when the image is converted to another color space the profile is used to convert the image data to the PCS before the data is then converted to the color space of the destination profile. After the profile is assigned, an image can be saved with the profile embedded. Most image file formats support embedded ICC profiles, and a full list of such formats is available [21].

Professional image applications support and correctly interpret ICC profiles. However, it should be noted that some applications that provide image previews are not color management-aware, and will therefore render the image RGB data directly

to the screen without interpreting the embedded profile. Some automated workflow applications also strip any embedded profiles on parsing the image.

As discussed above, it is recommended that the color management working space is set to ProPhoto RGB or other large-gamut encoding to avoid gamut clipping. Images are not saved in ProPhoto but converted to either scene-referred or standard output-referred encodings.

The ICC PCS (in ICC v2 and v4) is based on D50 colorimetry, and where the measured data used to characterize a camera, display or printer is not D50 a chromatic adaptation transform must be applied so that all data stored in the profile is D50. When interpreting the colorimetry of an image, it is therefore necessary to invert the chromatic adaptation transform and undo the media-relative scaling in order to determine the original colorimetry of the image.

iccMAX

The recommendations in this section apply primarily to v4 ICC profiles, currently the most widely used version. V2 profiles can also be used for camera profiles, although results may be less predictable. Either v2 or v4 profiles can be used as destination profiles.

ICC has released a new specification, iccMAX, which extends the functionality of the ICC architecture. iccMAX profiles can be used with v4 profiles where the PCS is colorimetric, but iccMAX also supports use of spectral, material and alternate colorimetric PCS, for example where the capture device is multi-spectral, where the input channel represent amounts of different materials rather than solely color, or where it is desired to use a colorimetry other than D50 in the PCS. More information on iccMAX is available [22].

Appendix D: Basic considerations for better color consistency

The best results for color consistency will be achieved by following the detailed recommendations given in Appendix A and B. If it is not possible to completely follow this guidance, then the following minimal guidelines should be followed:

Subject illumination:

1. Use stable and uniform diffuse illumination. Avoid shadows and highlights.
2. Use a broadband white light source that approximates the daylight spectrum.
3. The background should be a neutral gray color.
4. The surface of interest should be roughly perpendicular to the camera's optical axis.
5. Maintain the same lighting conditions for all images.

Camera setup:

1. Set camera to the proper white balance and exposure time. Avoid over or under-exposure. It is better to be a little under-exposed than a little over-exposed.
2. Use a RAW file format with the least image enhancement. If not possible, use uncompressed ≥ 12 bit TIFF files.
3. Properly frame and focus the object surface.
4. Avoid flash photography.
5. Turn off scene and spatially-dependent processing, such as high dynamic range or relighting.
6. Place a reference color chart in the scene for possible color correction.
7. Save images in the sRGB color space with a D65 white point.
8. Maintain the same camera settings (except for exposure time if lighting intensity varies).

Rendering the image:

1. Use image viewing software that properly recognizes the file format.
2. Use an sRGB display.

References

1. NEMA, Digital imaging and communications in medicine (DICOM), part 14: Grayscale Standard Display Function, vol. PS 3.14, National Electrical Manufacturers Association, 2001.
2. A. Badano, C. Revie, A. Casertano, W.-C. Cheng, P. Green, T. Kimpe, E. Krupinski, C. Sisson, S. Skrøvseth, D. Treanor, P. Boynton, D. Clunie, M. Flynn, T. Heki, S. Hewitt, H. Homma, A. Masia, T. Matsui, B. Nagy, M. Nishibori, J. Penczek, T. Schopf, Y. Yagi and H. Yokoi, "Consistency and Standardization of Color in Medical Imaging: a Consensus Report," *Journal of Digital Imaging*, vol. 28, no. 1, pp. 41-52, February 2015.
3. International Color Consortium, "Visualization of medical content on color display systems," April, 2016. [Online]. Available: http://color.org/whitepapers/ICC_White_Paper44_Visualization_of_colour_on_medical_displays.pdf
4. ISO 17321-1:2012, "Graphic technology and photography - Colour characterization of digital still cameras (DSCs) - Part 1: Stimuli, metrology and test procedures.," ISO, 2012.
5. ISO 22028-1:2016, "Photography and graphic technology – Extended colour encodings for digital image storage, manipulation and interchange - Part 1: Architecture and requirements.," ISO, 2016.
6. IEC 61966-2-1:1999, "Multimedia systems and equipment - Colour measurement and management - Part 2-1: Colour management - Default RGB colour space - sRGB.," IEC, 1999.
7. C. S. McCamy, H. Marcus, and J.G. Davidson, "A color rendition chart," *J. Appl. Phot. Eng.*, V2, pp. 95-99, 1976.
8. R. S. Berns, Billmeyer and Saltzman's principles of color technology. 3rd Ed. John Wiley & Sons, New York, 2000.
9. International Color Consortium, "Specification ICC.1:2010; Image technology colour management – Architecture, profile format, and data structure," December, 2010. [Online]. Available: http://www.color.org/specification/ICC1v43_2010-12.pdf
10. J. Penczek and P.A. Boynton, "Display color error in the medical digital image workflow," *SID Symposium Digest*, Society for Information Display, V45, pp. 348-351, 2014.
11. R.W.G. Hunt and M.R. Pointer, *Measuring colour*. 4th Ed. The Wiley-IS&T series in imaging science and technology, 2011.
12. J. Penczek, P.A. Boynton, and J.D. Splett, "Color error in the digital camera image capture process," *J. Digital Imaging*, V27, pp. 182-191, 2014.
13. CIE Technical Report 15: Colorimetry, Note that CIE is Commission Internationale de l'Eclairage (International Commission on Illumination), 2004.

14. R.D. Paravina, G. Majkic, M. del Mar Perez, and S. Kiat-amnuay, "Color difference thresholds of maxillofacial skin replications," J. Prosthodontics, V18, pp. 618-625, 2009.
15. Adobe, "Digital Negative (DNG) Specification," June, 2012. [Online]. Available:
http://www.images.adobe.com/content/dam/Adobe/en/products/photoshop/pdfs/dng_spec_1.4.0.0.pdf
16. International Color Consortium, "Why is the media white point of a display profile always D50?," [Online]. Available:
<http://www.color.org/whyd50.xalter>
17. J. Holmes, "Advanced Color Management for Digital Photography: and possibilities for using ICC profiles," February, 2006. [Online]. Available:
http://www.color.org/documents/AdvColMgmt_for_DP.pdf
18. http://www.color.org/info_profiles2.xalter#digitalphotography
19. <http://www.color.org/profiles.xalter>
20. <http://www.color.org/scene-referred.xalter>
21. http://www.color.org/profile_embedding.xalter
22. <http://www.color.org/iccmax/index.xalter>
23. <http://www.color.org/chardata/rgb/srgb.xalter>
24. <http://www.color.org/chardata/rgb/adobergb.xalter>
25. <http://www.color.org/srgbprofiles.xalter>
26. S. Van Poucke, Y. Vander Haeghen, K. Vissers, T. Meert, and P. Jorens, "Automatic colorimetric calibration of human wounds," BMC Medical Imaging, V10, pp. 1471-2342, 2010.