International Color Consortium®

White Paper #4

Level: Intermediate
Date: December 2004 (revised October 2020)

Color Management - Conceptual Overview, Evolution, Structure & Color Rendering Options

This White Paper provides an overview of some of the key concepts in color management, and describes the evolution from v2 of the ICC specification to the v4 specification in use today. It summarizes the color rendering and re-rendering options provided by the different intents within an ICC profile.

Color Management can be defined as the "communication of the associated data required for unambiguous interpretation of color content data, and application of color data conversions as required to produce the intended reproductions."

Color content may consist of text, line art, graphics, and pictorial images, in raster or vector form, all of which may be color managed. To be successful, color management must consider the characteristics of input and output devices in determining the appropriate color data conversions for these devices.

Evolution

We can identify four distinct phases in the evolution of the understanding of color management. Initially there was what could be described as 'digital color mode', whereby color was expressed terms of the coordinates obtained on devices, in color spaces such as RGB, CMYK and YCC. Subsequently it became common to describe colors by means of their colorimetry, using the well-established CIE system. The move to colorimetric specification of color led to the notion of 'device-independent color' - the idea that a color could be expressed in terms of its colorimetry independently of the device used to create it. Communicating color through CIE colorimetry works well when the viewing conditions are well defined such as it is the case, for example, in Digital Cinema according to the SMPTE standard “Digital Cinema Distribution Master – Image (DCDM)”. If additionally the output device is fixed, a well-defined process for rendering color to the output system and its viewing condition is established. In this case, the device-dependent color coordinates of the output device can be used to communicate colors, such as, for example, with the traditional television model and the reference display described in the user requirements for video monitors in television production, from the European Broadcast Union.

In more recent years, the effect of the viewing conditions on the appearance of color has become more widely appreciated, together with the effect this has on the desired colorimetry of a reproduction. However, the human visual system is still not well understood, and although we have models such as CIECAM02 which are successful for certain viewing conditions, there are as yet no published models that describe appearance robustly.

Today our understanding is based on the different image states in which an image can exist. The desired appearance of an image depends on the output medium and its viewing conditions, and some form of rendering is required to transform an image from one image state (such as scene-referred colorimetry) to another (such as output-referred display or print). This concept of rendering is distinguished from gamut mapping, which can be thought of as primarily an operation to clip a source gamut to a destination gamut of a different (usually smaller) size. The media- or image-specific preference aspect of the mapping can therefore be considered more as an operation to render between different image states.

If we are to successfully render between image states, it is essential that we are able to unambiguously interpret color data and hence it is necessary that the image state at any point in the workflow is known. This type of approach is in fact implicit in traditional photography and graphic arts, where for example a transparency is interpreted in a certain way in order to obtain a pleasing reproduction on a print.

We can then define two types of color management workflow. In the first, we can consider the output device and its viewing condition to be fixed, and thus the intended viewing conditions and mode of viewing, the dynamic range and gamut of the reproduction, and other characteristics of the medium such as the substrate and the type of surface, are all known. In this case we can ship the desired colorimetry to the output device, usually by means of a colorimetric transform to the device encoding.

In the second case, the output device is not completely fixed but is variable in some way, (for example through the option of having different viewing conditions, or through different output media being available). In this case the optimal image appearance may be device dependent, and a successful cross-media or cross-device color transform includes a color rendering between different image states.

Color appearance

Appearance models are frequently useful in imaging applications. Transforms between corresponding colors in different viewing conditions often apply the chromatic adaptation component of a color appearance model. Appearance models also provide more perceptually uniform spaces for gamut mapping, and can be used to model the dependence of colorfulness on absolute luminance. Some device characterization methods also perform error minimization in color appearance coordinates.

However, since the cross media objective is often NOT to reproduce appearance, color rendering approaches that independently use appearance models to deal with viewing condition differences, and gamut mapping to deal with gamut differences, may not be optimal. The primary color rendering task may actually be to alter appearance in order to produce a pleasing reproduction on different media. The changes in colorimetry driven by the appearance model may then be counter to those driven by gamut mapping, making independent optimization ineffective. Moreover, we don't yet have appearance models that robustly describe appearance, particularly for complex images as opposed to uniform stimuli.

Reproduction models

Reproduction models have to simultaneously consider the effects of viewing condition, media limitations, user preferences, and potentially image characteristics in developing optimal color rendering transforms. Such models can be based on an analysis of what is done to image colorimetry by experts in achieving excellent cross media reproductions. They are thus at least partially empirical - but so are appearance models and gamut mapping algorithms. They can add components based on our understanding of the human visual system as this understanding develops. The key to a successful model is simultaneous optimization of all the parameters described above.

Color Imaging Architecture

Unambiguous exchange of color image data requires that the different attributes of color are well-defined. ISO 22028-1 provides definitions of color space encoding, viewing conditions, image state and reference medium.

Color rendering can be applied in either proprietary or standardized ways. Standardization, where applicable, is essential in reducing possible ambiguity, and in achieving inter-system consistency, but it should also be recognized that proprietary methods have the potential for adding value and providing enhanced implementations.

Implementation mechanisms should be aimed at producing standard color encodings (i.e. encodings of the colorimetry of an image on a reference medium, including the associated viewing conditions). An image writer or reader is then required to color render to or from this standard color encoding. Attaching a color profile provides the transforms to be applied to the encoded image data in order to produce image colorimetry in a Profile Connection Space (PCS) describing a specified medium (including its associated viewing conditions). Appropriate transforms to and from the PCS are linked by the Color Management Module (CMM)

Color Rendering Options

In a color reproduction workflow, there are two possible options for handling the color rendering. An intermediate reproduction description provides input-side color re-rendering to some well-defined real or virtual reference medium. Image data is then exchanged and output-side color re-rendering is performed from the reference medium to the actual output medium. Alternatively, a deferred color rendering is achieved by encoding source colorimetry with the medium characteristics and information about the viewing conditions. The color re-rendering capability must be made available at the output stage, so that when final output is selected, color re-rendering is performed directly from source to actual output.

Early Binding and Late Binding are terms used in graphic arts to designate when in the workflow the conversion/separation to CMYK takes place. This workflow usually starts with an intermediate reproduction description created on a computer or produced by a capture device (now almost invariably RGB, although capture directly to CMYK is possible).

Early binding produces an intermediate reproduction description, based on some assumed output device. This (second) intermediate image may need to be color re-rendered to different output devices and media, such as proofs and prints made by different printing processes. It is helpful if early binding images are in some “standard” CMYK color encoding.

Late binding defers the conversion (or ‘separation’) to device values until the actual output device is known. In this case, multiple files may be produced for the different devices.

The advantages of the intermediate reproduction description can be summarized as:

- output is more consistent than with scene-referred exchange (since the desired artistic intent can be communicated in the intermediate image)
- proven in practice by photographers & graphic artists
- commonly used bridging transforms for color re-rendering can be highly tuned and made widely available
- requires less sophisticated processing capability at output

The disadvantages of the intermediate reproduction description are:

- color re-rendering to actual output may be necessary
- may not produce optimal results, particularly if the intermediate image reference medium is very different from the actual output medium
- there is less output-side control of scene-to-picture color rendering
- in the early binding case, assumptions that device values and GCR will or should be maintained when re-purposing may be incorrect

Deferred color rendering has the following advantages:

- output-side control of color rendering and re-rendering are increased
- color rendering or re-rendering is direct to the actual output
- there are no worries that the intermediate image is too different from the actual output
- in the late binding case, decisions involving device value selection (spot color substitution and solids) are deferred until the actual device is known.

The disadvantages of deferred color rendering are:

- less consistent output due to greater color rendering freedom
- a mechanism is needed for preview or proof of the color rendering
- the image creator’s artistic intent may not be maintained
- image data after processing for output is device specific, and can cause difficulties if fed back into open workflows

- the capability to perform color rendering or re-rendering from the source encoding must be available at output
- more hand tuning may be required, if more aggressive automated color rendering and re-rendering algorithms do not produce the desired result.

The current situation

In the case of color rendering (i.e. direct from scene to output-referred image data), the intermediate reproduction description approach dominates today. In most cases this is a standard output-referred exchange, using color encodings such as sRGB, Adobe RGB (1998), ROMM RGB and a standards referenced CMYK such as FOGRA 39. Manually guided deferred color rendering (e.g. camera raw) is becoming increasingly popular, especially in professional markets, although even in this case color rendering is typically to a standard output-referred color image encoding for exchange. Here the concept of the digital negative and positive, in which a master file is archived for subsequent rendering, is relevant.

In the case of color re-rendering (i.e. from an image in one output-referred medium to a reproduction on another output-referred medium) both early and late binding workflows are used, although, the image state is not always communicated. Re-rendering may be performed either by the CMM (using the media-relative colorimetric intent with black point compensation to scale the dynamic range of the first image state to that of the second) or by the profile (using the Perceptual rendering intent to adjust both dynamic range and colorfulness to provide a preferred reproduction for the second medium).

ICC v2 profiles are limited in the performance and reliability of color re-rendering using the perceptual intent, primarily because the dynamic range and color gamut of the first image state is undefined when applying the profile to perform the re-rendering. This problem is addressed in ICC v4 (which has a specified black and white point for the perceptual PCS and a well-defined Perceptual Reference Medium Gamut), although v4 profiles are not yet used in all workflows.

Using the Media-relative Colorimetric intent with black point compensation with v2 profiles deals with at least the first-order dependency of the desired appearance on the intended reproduction medium by means of the dynamic range adjustment, but this approach is not entirely optimal. Advanced CMM-based color re-rendering can overcome this limitation, but the use of such CMMs is not yet common and their required behavior is not standardized. The algorithms required to perform color re-rendering are rapidly evolving, and in some cases, with particularly difficult mappings between color gamuts, the transform must be hand-tuned in order to achieve optimal performance.

For both color rendering and re-rendering, there are two types of implementation in use:

- sRGB is used as an output-referred intermediate reproduction description based on a reference display and viewing conditions
- ICC profiles offer several rendering intents, supporting both color rendering and re-rendering.

sRGB is widely used, especially in consumer devices, and the quality of implementations continues to increase as understanding evolves and the color rendering capability increases.

ICC profiles provide a Perceptual intent based on a reference print intermediate, together with measurement-based Colorimetric intents which enable deferred color rendering by smart (generally proprietary) CMMs. They also enable colorimetric proofing. A degree of standardized color rendering capability is provided by some CMMs through support for Media-relative colorimetric with black point compensation.

The ICC Saturation intent enables proprietary workflows, where the rendering goal is different from that expressed in the Perceptual and Colorimetric intents.

Like sRGB, ICC-based color management is evolving as understanding of the use cases, requirements, rendering methods and color management architecture continues to increase.

ICC v2 issues

Version 2 of the ICC specification had a number of significant short-comings.

- Although the PCS is D50, the chromatic adaptation which had been performed to obtain a media white point in D50 was not required to be defined within the profile, and as a result the chromatic adaptation state of input data was ambiguous.

- The color re-rendering that was required in order to obtain the desired appearance on the PCS reference medium from input data was not defined, and for the perceptual intent there was no standard reference medium. This led to different assumptions about the PCS perceptual dynamic range and color gamut by different profiles.

- Colorimetric intents were not required to be measurement-based, and since in addition measurement methods were not always well-defined, the behavior of the colorimetric intents was unpredictable.

- There was insufficient flexibility in the transforms and color processing models provided within the v2 specification.

As a result of these short-comings, capability limitations and interoperability problems could result.

There were at least three possibilities for input-side color re-rendering in v2:

1. Colorimetric with no black scaling
2. Colorimetric with black scaling
3. Perceptual to some arbitrary reference medium.

Depending on the source image, and the input profile re-rendering, the PCS colorimetry could thus be appropriate for a variety of different media and viewing conditions, and which these were was not identified within the profile.

The different input-side color re-rendering possibilities are illustrated in Figure 1.

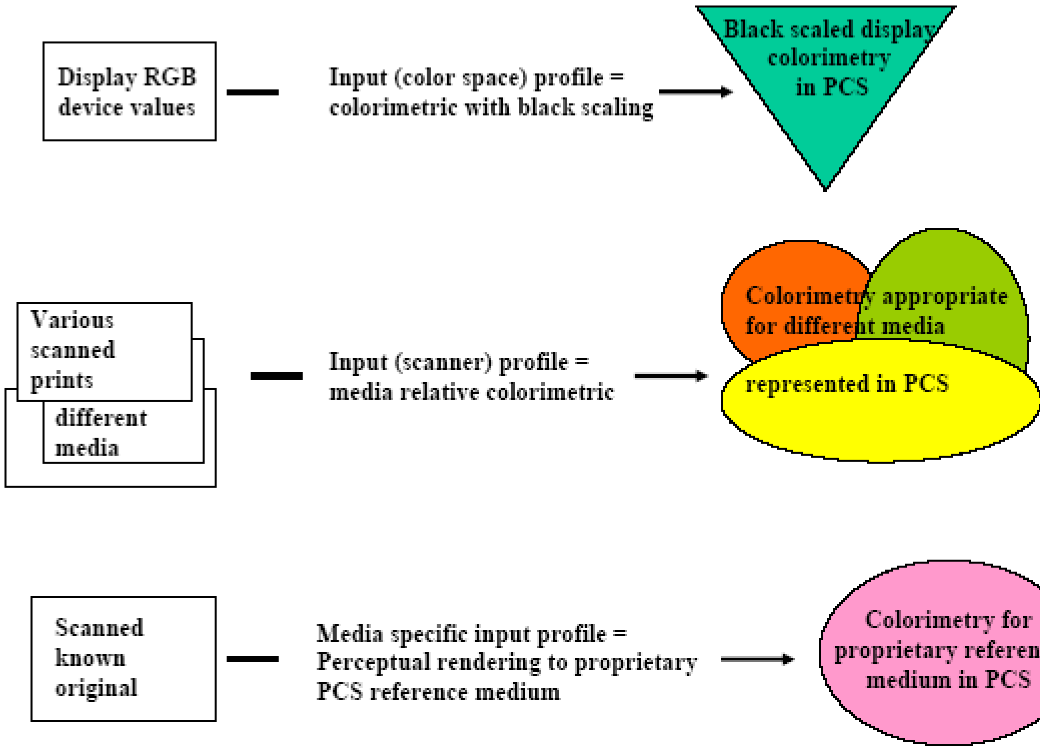

FIGURE 1 V2 INPUT COLOR RE-RENDERING POSSIBILITIES

These multiple input-side re-rendering possibilities lead to a dilemma for v2 output profiles. The perceptual intent of a v2 output profile was supposed to perform a pleasing re-rendering of the PCS image colorimetry to the actual output medium and viewing conditions. However, the output profile creator had no knowledge of the medium and viewing conditions for which the PCS colorimetry was appropriate! It is impossible to create an optimal perceptual rendering without this knowledge, and therefore optimal cross-vendor interoperability is precluded - while the output profile knows the end result, there are in effect many possible starting points in the PCS for a given set of input data., as illustrated in Figure 2.

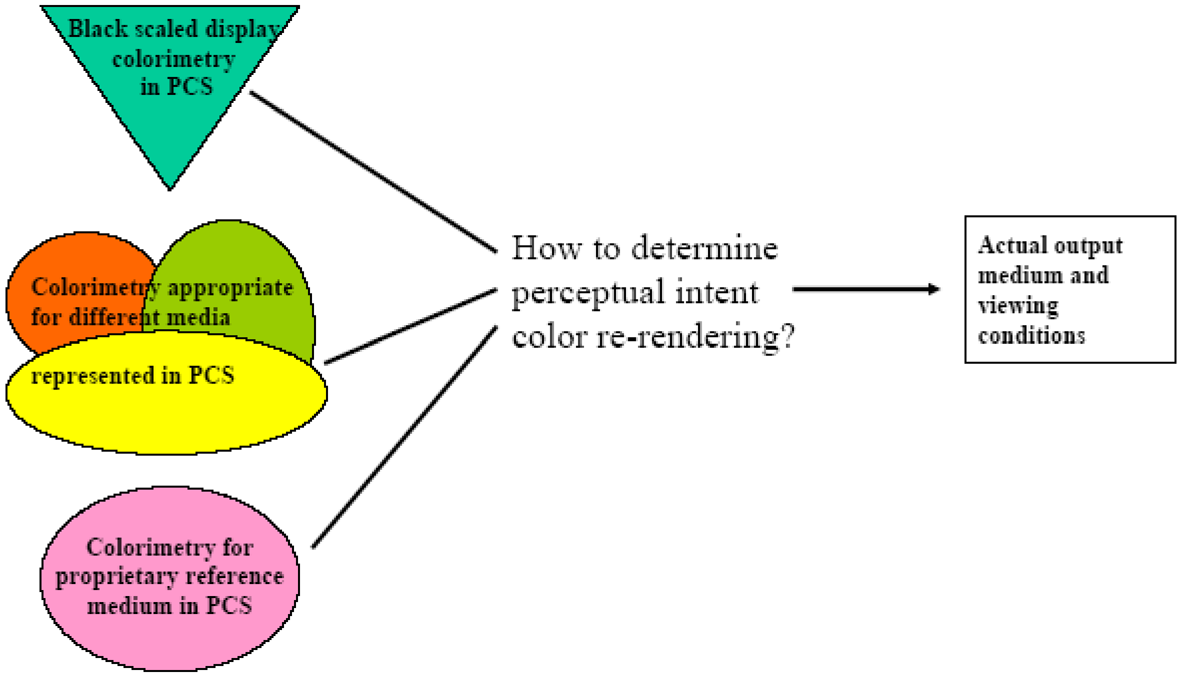

FIGURE 2 AN OUTPUT PROFILE PERCEPTUAL RENDERING INTENT HAS MANY POSSIBLE DATA SOURCES.

The Colorimetric rendering intent in 3v2 also presents implementation issues. In a v2 profile, the source colorimetry may be black scaled or color re-rendered to a proprietary reference medium, in order to enable improved interoperability within a single vendor's products. Because PCS colorimetry may not be accurate relative to the original, CMM cannot rely on the source colorimetry, as represented in the PCS, and as a result v2 profiles will not support advanced CMM color rendering. There are also other issues that arise with v2 profiles because of the ambiguity of the v2 specification and incorrect interpretation of the specification in constructing profiles.

The ICC v4 Solution

In ICC v4, colorimetric rendering intents are measurement based. They can therefore be relied on for proofing, and provide accurate colorimetry for CMM color re-rendering. Specification ambiguities are largely resolved and the text clarified to reduce the occurrence of incorrect implementations. A well-defined reference medium for the Perceptual intent, with an associated gamut known as the Perceptual Reference Medium Gamut (PRMG) ensures cross-vendor interoperability. There is also greatly increased transform capability through extended LUT definitions, such as the lutAtoBtype which incorporates an additional matrix and curve and provides greater mathematical flexibility and an improved definition of 16-bit CIELAB.

ICC v4 Perceptual intent

Significant improvements have been made to the interoperability of the v4 Perceptual path. With the Perceptual Reference Medium Gamut, both input and output profiles can be based on a well-defined intermediate image colorimetry appropriate for the PCS reference medium and viewing conditions. The task of the CMM is thus to connect profiles with the same (or very similar) PCS gamuts, and minimal gamut mapping is required because the image colorimetry in the PCS is matched for the input and the output. Differences between source and output media color gamut and viewing condition are then dealt with consistently within the mapping to or from the reference medium performed by each profile.

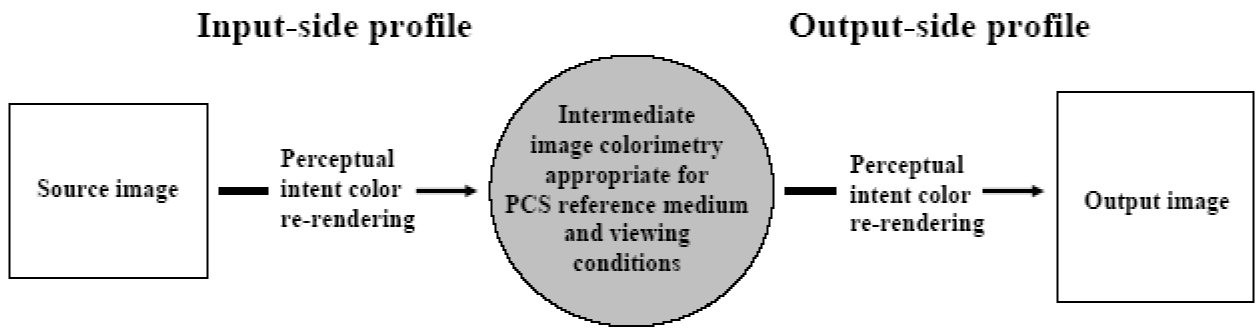

FIGURE 3. PERCEPTUAL INTENT COLOR REPRODUCTION PATH IN ICC V4.

The v4 Perceptual transform includes both the data (typically device value) to PCS colorimetry transform, and color re-rendering to and from the reference medium in the PCS. The re-rendering operation includes consideration of:

- differences in viewing conditions between source and reproduction and their appearance effects
- differences in media characteristics and image state
- color rendering preferences and the attributes of the preferred reproduction on the output medium.

If profiles incorporate all of these considerations, the task of the CMM is simply to connect profiles together to create the transform between source and output data.

The v4 Perceptual transform is useful for general image reproduction across all devices and media. Since color re-rendering operations are typically proprietary, profiles from different sources may produce different “looks”, i.e., color aims. Users may select profiles based on color re-rendering preferences. This was difficult before v4 owing to the v2 issues described above, and a lack of coordination between the different color management components (the operating system, the application and the driver and/or RIP. As differences between actual and reference media decrease, the perceptual and colorimetric intents should converge. Before v4, users were cautious about the perceptual intent because of the inconsistencies with v2. However, it is still important that v4 profiles are correctly constructed and that color management is well coordinated in order to maximize the confidence of users.

ICC v4 Colorimetric intents

The ICC v4 Colorimetric path is illustrated in Figure 4.

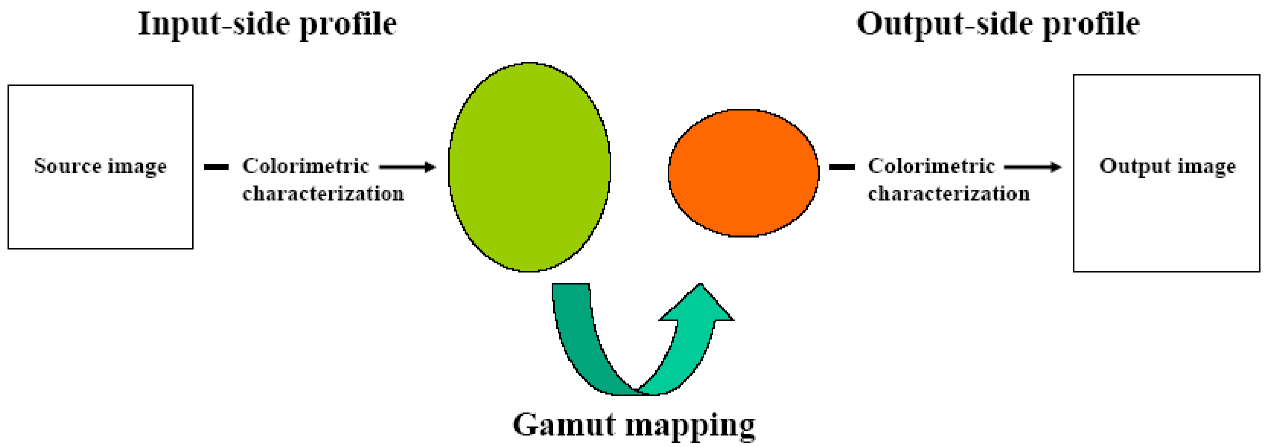

FIGURE 4. ICC V4 COLORIMETRIC PATH

The color gamut mapping performed by a v4 profile has three requirements:

- the input data colorimetry should not be changed within the intersection of the input and output media gamuts
- colors that are outside the source image gamut should not be produced in the output image
- colors in the source image that are outside the output image gamut should be clipped.

A colorimetric transform includes the device data to PCS colorimetry transform, based on measurements made using standard methods (as defined in ISO 13655, and described in ICC White Paper 3). The transform also includes chromatic adaptation to and from the D50 PCS white point, if the data has a different reference white. This allows gamut mapping to be performed directly, if desired. In proofing situations, the extent of gamut mapping required is best minimized by the choice of proofing media. As the chromatic adaptation matrix is included in profile, it is invertible if CMM-based chromatic adaptation is desired. The Colorimetric intent does not include other appearance transforms, in order to avoid unnecessary color appearance model complexity, instability, and other issues mentioned above.

Colorimetric transforms are useful for preview and proofing applications, and in support of CMM-based color rendering. The Media-relative colorimetric with black point compensation (MRC+BPC) provides a standard baseline CMM color rendering that is adequate when the media, substrate and gamut shape differences are not large. This baseline reproduction model includes chromatic adaptation and media white relative colorimetry with black point scaling (on XYZ coordinates). It also includes gamut expansion and compression as required. The current widespread use of MRC+BPC demonstrates the importance of media considerations.

ICC v4 CMM color rendering

In ICC v4, it is possible for color rendering to be performed by the CMM rather than the profile, as illustrated in Figure 5.

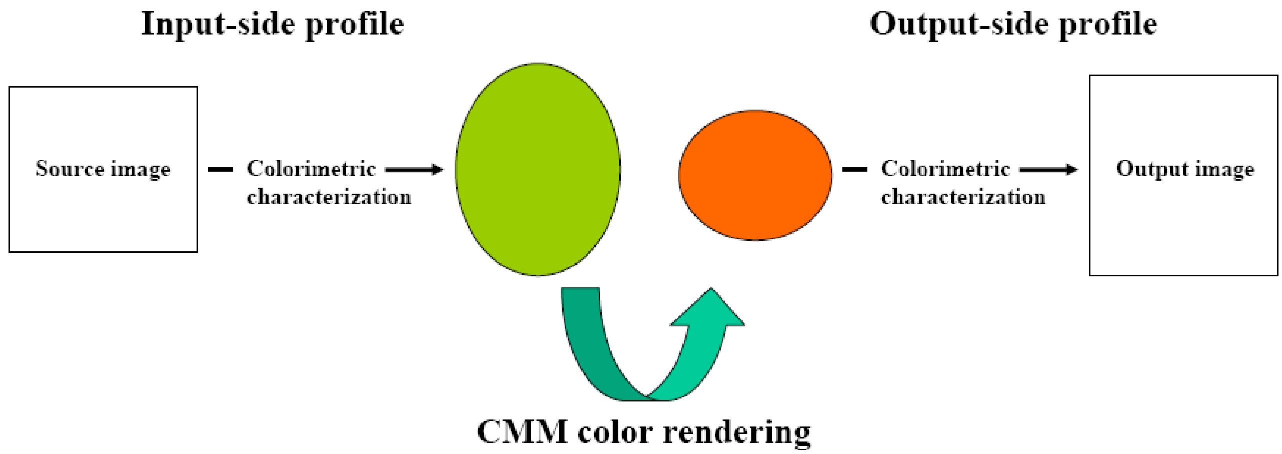

FIGURE 5. ICC V4 CMM COLOR RENDERING

In this scenario, CMM algorithms color re-render source image colorimetry to be appropriate for the actual output medium, taking into consideration source and output medium color gamuts and viewing conditions. They can also support color appearance model based color re-rendering. CMM-based color rendering can take advantage of full output medium gamut, and facilitate user adjustment of color re-rendering at the time of output. For more details of CMM capabilities in ICC v4, please see ICC White Papers 25 and 28.

Moving forward

Current research into color rendering supports both automated perceptual intent transform generation and increased CMM color rendering capability.

High quality ICC v4 tools and profiles have become more widespread and v2 issues are less of a problem. However, we acknowledge that considerable work is still needed to fully coordinate color management across operating systems, applications and devices. ICC and its members and the color management community need to work in a coordinated way to advance all of the technologies described above, building where possible on solid understanding and communication. Clear and unambiguous definitions of color encoding and image state, for example through ISO 22028-1, are key to this process.

The iccMAX specification enables a much wider range of possible functionality than v2 or v4. Users should take care to ensure that the color rendering applied by an iccMAX profile is consistent with the principles outlined in this White Paper.