International Color Consortium®

White Paper #2

Level: Advanced
Date: Jan 2005

Perceptual Rendering Intent Use Case Issues

The perceptual rendering intent is used when a pleasing pictorial color output is desired. [A colorimetric rendering intent is used when an output is to be color matched to its source image.] The perceptual rendering intent is most often used to render photographs of scenes (i.e. views of the three-dimensional world), and when the objective for a reproduction is to obtain the most attractive result on some medium that is different from the original (i.e. re-purposing), rather than to represent the original on the new medium (i.e. as in proofing or re-targeting). Some level of color consistency is required - for example colors should not change hue names. However, with perceptual rendering, if the reproduction medium, for example, allows for greater chroma than the original medium, then chroma may be increased to produce a more pleasing result. Likewise, if the reproduction medium has a smaller color gamut than the original medium, perceptual rendering may alter in-gamut colors to allow for graceful accommodation of the original color gamut through gamut compression. In comparison, colorimetric rendering maintains in-gamut colors across media at the expense of sub-optimal colorfulness on larger gamut reproduction media and clipping artifacts on smaller gamut reproduction media.

Keep in mind that the perceptual rendering intents in ICC profiles provide one approach to perceptual color rendering or re-rendering. There are other ways. Devices such as digital cameras and printers perform embedded (typically proprietary) perceptual renderings to and from standard color encodings like sRGB. In certain workflows, abstract ICC profiles can be used in combination with a colorimetric rendering path through source and destination ICC profiles to perform color re-rendering from source image colorimetry to destination image colorimetry directly in the PCS, before transforming to the destination encoding. Alternatively, a user may apply manual image editing techniques to optimize an image for a particular output condition. Finally, a color management system (CMS) may offer color rendering or re-rendering capabilities beyond that built into any source and destination profiles.

‘Media-relative colorimetric plus black point compensation’ is a simple and widely used perceptual rendering that uses the media-relative colorimetric rendering intent in the source and destination ICC profiles, combined with black point scaling performed by the CMS. Simple media white and black scaling can accommodate differences in dynamic range between an original and a reproduction and (to some extent) differences in color gamut size. In cases where color gamut shapes are roughly similar, and gamut size differences correlate with white and black point differences, media-relative colorimetric plus black point compensation may produce excellent perceptual rendering. However, this approach is not universally available because some CMSs do not support black point compensation. In other cases, more elaborate perceptual transforms are required to produce optimal results, especially when the source and destination media are quite different. The inclusion of an explicit perceptual rendering intent in ICC profiles enables well-defined, repeatable, and high quality perceptual rendering across all ICC based color management systems.

Scene to Reproduction

Scene-to-reproduction perceptual rendering is discussed first because such color rendering must happen first in the capture of natural scenes, and understanding this transformation is helpful in understanding subsequent transformation requirements. However, users should be aware that in typical digital camera workflows scene-to-reproduction perceptual rendering is not accessible to user control. Virtually all digital cameras perform scene-to-reproduction color rendering in the camera. The image file output by the camera does not represent the scene, but rather represents what the camera manufacturer feels will likely be a pleasing reproduction of the scene. This reproduction typically includes alterations of the scene colorimetry, including highlight compression, and midtone contrast and colorfulness enhancements as discussed below.

Likewise, camera raw processing applications typically embed scene-to-reproduction color rendering. While it is possible to create true scene-referred images from camera raw image data, most camera raw processing applications do not do this. Camera profiling applications include scene-to-PCS color rendering but may not offer user controls (note that with some camera profiling applications the scene color analysis accuracy is limited more by the accuracy of the target based characterization method than by intentional preferential alterations).

In the future, it is expected that users will have more access to scene-referred image data, thereby gaining more explicit control over scene-to-reproduction color rendering. At present, this material is primarily intended as background, and for custom workflows where special camera modes or processing applications are used to enable true scene-referred image creation, followed by scene-to-reproduction color rendering.

The **essential process** in any scene to reproduction (scene-referred-to-reference-output-referred) perceptual transformation is a coordinated combination of color appearance adaptation, preference adjustments and gamut mapping. This perceptual rendering intent *color-rendering* transformation is used to map scenes to the fixed range of a reproduction in a pleasing way. When a source image is scene-referred, the Device-to-PCS perceptual transform performs a perceptual rendering from the scene to the perceptual intent reference medium. Note that in an ICC Version 4 compliant [scene-referred] input profile (e.g., a digital camera input profile), the reference-output-referred-to-scene-referred PCS-to-Device perceptual rendering intent transform should invert (i.e., undo) that profile's own Device-to-PCS perceptual rendering intent transform.

Commonly, the color appearance adaptation portion of a perceptual color-rendering transformation includes adaptation from the scene adopted white (both the chromaticity and luminance) to the adopted white of the reproduction. Reproduction constraints and color appearance preferences determine the mapping of the adopted-white-adapted scene colorimetry to produce a pleasing reproduction. For example, if the scene luminances are much higher than those of the reproduction in the anticipated viewing conditions, a chroma boost may be necessary to maintain the appropriate colorfulness. The anticipated surround of the reproduction can affect the desired contrast, with darker surrounds requiring higher contrast. Preferences play a significant role in determining this mapping, as viewers tend to prefer increased colorfulness and contrast in reproductions, to the extent that the increases do not look unnatural. Ideally, mappings are determined on a scene and output medium specific basis, implying image-specific perceptual intents. In production workflows fixed mappings that work reasonably well for most scenes are often used. These mappings typically boost the scene gamma and mid-tone contrast. For example, film reproduction systems have a mid-tone gamma greater than unity (~ 1.2 to 1.6, depending on the anticipated output medium) combined with highlight and shadow roll-offs. This s-shaped mapping allows film systems to accept both low and high dynamic range scenes, while maintaining preferred mid-tone contrast and colorfulness. Likewise, video systems have a system gamma of ~1.2 to 1.4 and some highlight compression (at least in high-end systems).

The preference adjustment portion of a perceptual color-rendering transformation often includes preferential expansion or compression of the source gamut and dynamic range to match that of a particular output (visualization) medium. Source scene gamut expansion and compression may be determined based on the potential scene extent from a particular digitization source device. Alternatively, in scene-specific color-rendering cases, each specific source scene gamut extent may be evaluated and preferentially expanded or compressed to match the output medium. In some cases, preferential mappings also explicitly consider the reproduction of memory colors. Following such appearance-preference mapping, it may be necessary to apply gamut mapping to bring the remapped colors to within the actual gamut of the destination medium. Ideally the appearance-preference mapping would accomplish this, but practically, a following gamut mapping operation may be required. Note that the perceptual rendering intent color rendering provided in Version 4 input profiles targets the ICC perceptual intent reference medium.

Optimal preference mappings differ for scenes of low, medium, and high dynamic range, key, and gamut extent. Some scenes have colors out to the spectral locus (and beyond, after chromatic adaptation) and have very high luminance (dynamic) ranges, however, many scenes do not. In fact, most scenes have dynamic ranges (and gamuts) smaller than the 288:1 of the ICC perceptual intent reference medium. ICC profiles are often used in capture-condition or visualization-condition (i.e., image state) specific – rather than image specific – workflows. With these workflows, customizing the choice of rendering intent is one way to adapt the use of an ICC profile to a particular scene or color object.

It should be noted that the capture digitization of an original (two-dimensional) artwork or photograph is different from the capture of a scene, which is a view of the natural (three-dimensional) world. The discussion above relates to the capture of scenes. The capture of originals, even using a digital camera, falls under re-targeting or re-purposing as discussed below. Perceptual rendering intents for scene capture will generally not be appropriate for the capture of two-dimensional originals.

Re-targeting and Re-purposing

After data is color rendered to a particular reference-output-referred or actual-output-referred *first visualization* condition, i.e., output-referred image state, it may be necessary to transform the data for a *second visualization*. For example, in a typical digital camera workflow, the ‘pleasing reproduction of the scene’ produced by the camera is targeted for viewing on a softcopy display. That display-referred data may be color re-rendered when a print output is desired. Two scenarios are defined regarding the color re-rendering transformation. When the second visualization is intended to represent or match the original first visualization, this is called *re-targeting*. Re-targeting is typical for ‘proofing.’ When the second visualization is independent of (i.e., not constrained by) the first visualization and can be optimized for the second visualization condition, this is called *re-purposing.* Keep in mind that both re-targeting and re-purposing are intended to operate on source images that are already in a picture-referred image state (either original- or output-referred, but not scene-referred).

In re-targeting, the Device-to-PCS media-relative colorimetric transform of the first visualization output or display profile is sequenced with the PCS-to-Device media-relative colorimetric transform of a second visualization output or display profile. [Absolute colorimetric intents can be used when the color of the target substrate from the first visualization is to be carried through to the second visualization.] No new or revised image state preferential rendering is called for in re-targeting. The accuracy of the representation through the second visualization condition will be proportional to the capability of the second visualization condition to match the first visualization condition (e.g., gamut volume shape, luminance range, and color differentiation).

In re-purposing, the first concern is to remove the constraints in the color data that were induced by prior perceptual rendering for a particular visualization condition (constraints preferentially based on a color aim determined as a function of prior source and destination image states). It is problematic that the constraints induced by a first preferential color rendering cannot be determined by examining color data after it has been so rendered. Color aim preferential rendering behavior is also not easily determined by examining the perceptual rendering intent transform of an output profile. Further, preferential capabilities in a CMS may have contributed to the first visualization, and can be difficult to extract in preparation for a later visualization.

In support of re-purposing, the ICC Version 4 specification places a new emphasis on perceptual rendering intent transformations:

- In ICC Version 4 compliant [actual-output-referred] output profiles, the actual-output-referred-to-reference-output-referred Device-to-PCS perceptual rendering intent transform should invert (i.e., undo) that profile's own PCS-to-Device perceptual rendering intent transform, to allow for re-purposing from the ICC perceptual intent reference medium.

- In ICC Version 4 compliant [original-referred] color space encoding and scanner input profiles (e.g., an sRGB profile, a document scanner profile), the Device-to-PCS perceptual rendering intent transform should color re-render the original to an appropriate ICC perceptual intent reference medium representation (i.e., transform from the device, or encoding, medium image state to the ICC perceptual intent reference medium image state).

- In ICC Version 4 compliant [original-referred] color space encoding and scanner input profiles (e.g., an sRGB profile, a document scanner profile), the PCS-to-Device perceptual rendering intent transform should color re-render back to the original (i.e., transform from the ICC perceptual intent reference medium image state to the device, or encoding, medium image state) to allow for a new re-purposing directly from the original-referred image state. Note that in order to provide for lossless round-trip this PCS-to-Device perceptual rendering intent transform should be an inverse of the Device-to-PCS perceptual rendering intent transform.

With Version 4 ICC profiles, re-purposing can be accomplished by sequencing the Device-to-PCS perceptual rendering intent transform of a “source” first visualization output profile with the PCS-to-Device perceptual rendering intent transform of a second visualization output profile. The Device-to-PCS perceptual transform from the source output profile "undoes" the previous perceptual color re-rendering from the perceptual intent reference medium to the source profile's actual output medium.

Note that use of the perceptual “undo” is appropriate only if the first visualization resulted from a perceptual rendering transformation. The rule of thumb is that the inverse of the rendering intent that was used to produce a particular visualization should be used to “undo” that visualization. Also note that even with the improved support in compliant Version 4 ICC profiles, subsequent visualizations can be constrained by loss of color detail in earlier transformations.

For re-purposing in general, when the destination output-referred image state gamut and viewing environment condition are “like” that of a source output-referred image state, then a colorimetric intent, with no preferential adjustment, may achieve acceptable results. (In fact, if the source and destination media are similar to the ICC perceptual intent reference medium, there should be little difference between the colorimetric and perceptual intent transforms.) On the other hand, when there are significantly different gamut constraints, and/or viewing environments, then a perceptual rendering intent, with inherent preference adjustments, can improve results. PDF/X-3 files, containing a fully populated (complete sets of PCS-to-Device and Device-to-PCS transforms) ICC output profile that describes the PDF Output Intent, support this type of re-purposing.

The goal with version 4 ICC profiles is to enable blind use of perceptual intents for re-purposing. It is expected that as version 4 profiling tools become more capable in generating quality perceptual color re-rendering transforms, this goal will be realized. However, in critical applications with media that are quite different from the perceptual intent reference medium, sophisticated users may find that careful, controlled application of colorimetric intents, abstract profiles, and CMS color rendering can produce better results.

Preserving an Artistic Intent through Multiple Visualizations

Preserving an artistic intent through multiple visualizations can require a combination of re-targeting and re-purposing approaches. The approach that is most likely to produce the best results in a particular situation depends on the similarities of the various actual media to each other, and to the perceptual intent reference medium. When multiple independently optimized visualizations are planned in advance, alternative approaches can be considered. If a specific artistic intent is desired, particular care should be taken with the first visualization.

A large gamut output-referred source image can be obtained by first applying the appropriate perceptual intent transform to color render scene-referred image data to the ICC perceptual intent reference medium, and then transforming the colorimetry of that reference-output-referred first visualization image to an appropriate storage color encoding such as ROMM/ProPhoto RGB. (Note that for a color encoding to be appropriate for this use the encoding image state will match the ICC perceptual intent reference medium image state, and the profile for that color encoding will have identical perceptual and colorimetric rendering intents). Alternatively, after using an appropriate perceptual intent transform to color render scene-referred image data to the perceptual intent reference medium, a first ‘actual’ visualization can be obtained by using an appropriate perceptual intent color re-rendering transform to re-render from the perceptual intent reference medium to the medium of a large gamut output device. Using such a “superset” first visualization as the source for subsequent visualizations can improve the optimization for each subsequent visualization, while maintaining color fidelity with the intended artistic intent.

When a color rendering to a first visualization represents a "master" image, including the artistic intent of the image creator, subsequent color transformations should not "undo" the initial perceptual intent color rendering. A subsequent actual output-referred visualization can be produced via a re-targeting approach (i.e., using colorimetric transforms) when the actual output medium is “like” the master image medium. When a subsequent actual output medium is dissimilar to the master image medium, the approach most likely to produce the best results depends on the relations of the media to each other, and to the perceptual intent reference medium.. If the master image is targeted to the perceptual intent reference medium and an actual output medium is dissimilar to the perceptual intent reference medium, then the perceptual intent transform of the actual output destination profile should be used to color re-render from the perceptual intent reference medium to the actual output medium.

If the master image is targeted to a first actual visualization medium that is substantially different from the perceptual intent reference medium, and a subsequent actual output medium is similar to the first visualization medium, then a Device-to-PCS colorimetric intent transform from the source (first) output profile should be used to re-encode the master image in ICC PCS. Then the PCS-to-Device colorimetric intent transform of the subsequent actual output destination profile should be used to create device values for the subsequent actual output medium.

The case where the first actual visualization medium, the perceptual intent reference medium, and the subsequent actual output medium are all substantially different from each other is the most challenging for color management. Ideally, the first actual visualization medium profile should perform color re-rendering to the perceptual intent reference medium, followed by color re-rendering from the perceptual intent reference medium to the subsequent actual output medium. However, it is possible, perhaps likely, that the input and output profile perceptual color re-renderings may not be complementary with each other to preserve the master image artistic intent. In that case using specifically tuned DeviceLink profiles to transform directly between the first visualization and each subsequent visualization may produce better results.

Note that when no related artistic intent is required among the multiple visualizations, then more flexibility in the final output can be obtained by retaining capture-referred (scene- or original-referred wide gamut RGB) data to use as the source for each independent visualization color rendering or re-rendering. This enables maximum flexibility for each visualization. It should be noted that this approach can produce significantly different versions of the same image, as scene-to-picture color rendering can be quite aggressive, and involve choices such as overall lightness, contrast, tone and saturation that go beyond the optimization of the scene to some output medium.

Additional Rendering Intent sequence examples

Visualization of the ICC perceptual intent reference medium image: When it is desirable to visualize the perceptual intent reference medium rendition of a color image directly, a visualization device with capability matching or exceeding the perceptual intent reference medium is required. Given that, one can use media-relative colorimetric rendering from the PCS, re-targeting the perceptual intent reference medium image to the actual output device (after correct perceptual rendering to the perceptual intent reference medium). Such visualizations should then be viewed in the reference viewing conditions (ISO 3664 condition P2) to produce the appropriate appearance.

Image specific preferential color rendering: As discussed above, image specific profiles and/or rendering intents can be used to obtain optimized preferential color-renderings from the capture-referred state to the reference-output-referred ICC perceptual intent reference medium. Use of image specific color-renderings should consider the need for color appearance compatibility across the various color objects intended for a particular document.

Color rendering or re-rendering from an ambiguous image state RGB color encoding: The first decision in this situation is, “Has the color data been previously color rendered to an output-referred state?” Certain RGB encodings inherently carry with them a particular image state: sRGB is output-referred for monitor viewing; ROMM/ProPhoto RGB is output-referred for the ICC perceptual reference medium print condition. It can be helpful to understand the use case or workflow that produced the RGB data when inferring the color-rendering image state condition. Typically, RGB data that is exchanged will have been color rendered to a first visualization and can be considered output-referred. However, beyond that it may be difficult to determine whether the RGB data is optimized for print or monitor viewing. When color re-rendering from an RGB working space, both the image state of the data, and the medium to which it may have been previously “color rendered” can affect the outcome of a subsequent color re-rendering. Keep in mind that manual adjustments may have been applied to optimize the data for a particular visualization. Caution is required because repeating a scene-referred-to-output-referred perceptual rendering intent transformation (as described above) will degrade image quality, as will applying an inappropriate color re-rendering transformation.

A source rendering intent can be selected to be appropriate for the image data in a particular working space. For example, prior to printing typical sRGB image data, it should be re-purposed from its display-referred state to the reference-output-referred image state corresponding to the ICC perceptual intent reference medium. On the other hand, if a user has been editing Adobe RGB image data to produce a desired appearance on a print media, a relative colorimetric source rendering intent may be appropriate when transforming for print.

When selecting the 'next visualization' destination rendering intent for a previously color-rendered (output-referred) RGB encoded image, as above, color re-rendering from the perceptual intent reference medium to an actual-output visualization encoding can be media-relative, or absolute colorimetric when the actual-output visualization gamut extent and tone range are similar to the reference medium gamut extent and tone range. When the actual output visualization gamut extent and tone range are significantly different from the reference medium gamut extent and tone range then perceptual rendering may provide an improved result.

Color re-rendering of computer generated imagery: Use of the perceptual rendering intent in reproducing computer generated color infers the computer display as the 'original' capture device. The computer display 'synthetic original' (original-referred image state) can be preferentially color re-rendered to the ICC perceptual intent reference medium using the perceptual rendering intent of a Version 4 compliant input profile for the computer display. Consideration of the rendering intent to use from the reference medium to the 'next visualization' actual output encoding is similar to that discussed above.