LED Color Characteristics

Color quality is an important consideration when evaluating lighting products. This fact sheet reviews the fundamentals regarding light and color, summarizing the most important color issues related to white-light LED systems, including color consistency, stability, tuning, and rendering, as well as chromaticity.

LED Emission Attributes

Individual LED dies, often referred to as chips, emit light in a narrow range of wavelengths, giving the appearance of a monochromatic source. LED lamps and luminaires combine multiple spectral components, which may be produced directly or through phosphor conversion, to create a mixture that appears white to the human eye. In comparison, incandescent lamps have a broad distribution, whereas fluorescent lamps typically rely on a limited, fixed set of phosphors with specific emission characteristics. Figure 1 compares the spectral power distributions (commonly called SPDs) of several example light sources, adjusted for equal lumen output. Importantly, because of the tristimulus sensitivity of our eyes, there are numerous ways that spectral components can be combined to produce the same color of light.

Creating White Light with LEDs

Currently, white light is most often achieved with LEDs using phosphor conversion, but color-mixed systems that use a combination of colored LEDs (typically red, green, and blue [RGB]) are also available. For LEDs, phosphor down-conversion, or the process by which a phosphor absorbs shorter-wavelength energy and emits it at a longer wavelength, is most commonly based on a blue- or violet-emitting die combined with a broad, yellow-emitting phosphor. The ratio of the direct emission and phosphor emission, as well as the phosphor composition itself, can be varied to change the color of the emitted light.

Color-mixed LED sources produce white light by mixing two or more colors, called primaries. Three-primary color-mixed systems are typically considered a minimum for acceptable performance. Additional primaries are sometimes added to improve the color rendering. While color-mixed LED sources have a higher theoretical maximum efficiency, they also currently face the challenge of the "green gap," which is used to describe the low efficiency of green-emitting LEDs. This is sometimes addressed by using a hybrid system, which might combine phosphor-converted LEDs with red LEDs, for example. It is important to remember that color characteristics vary widely, so generalizations about product categories can be
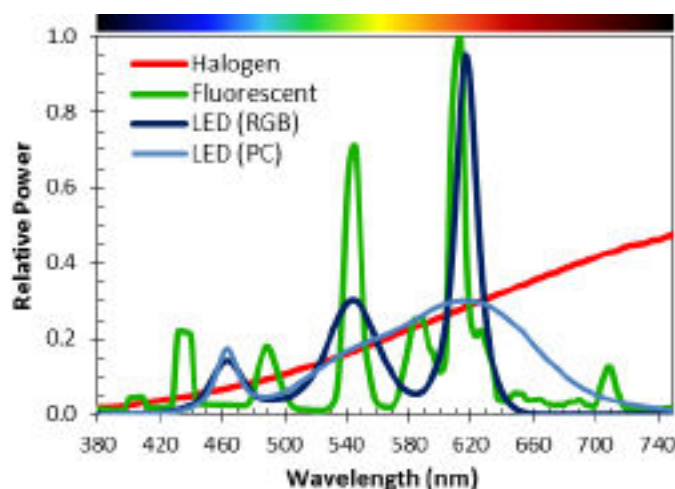

Figure 1. Due to the trichromatic nature of the human visual system, white light of equal appearance can be created with different spectral power distributions. This chart shows the spectral power distributions for four different light sources, including phosphor coated (PC) and a red-green-blue (RGB) color-mixed LED products, all the same color of white light.

misleading. Each individual product should be evaluated on its own merits, regardless of the technology.

For more on the differences between phosphor-converted and color-mixed LED systems, including future performance projections, see the Solid-State Lighting Research and Development Plan, available from: *http://energy.gov/sites/prod/files/2015/06/f22/ssl_rd-plan_may2015_0.pdf*.

Color Tuning

An intriguing benefit of LED technology is the relative ease by which color properties can be adjusted. This applies both to the engineering of a fixed-output product, where the engineer has nearly limitless options, and to color-tunable products, where the end user can adjust the color to his or her liking. Color tuning usually focuses on the color of the light itself, but may also include how the lighting affects the appearance of objects it illuminates (color rendering). Color tuning may be as simple as having two phosphor-converted LEDs of different colors in the same product, which allows combinations of those two LEDs; it could also be as complex as providing a full range of hues that can be mixed from individual primaries. More information on color tuning can be found at: *http://energy.gov/eere/ssl/led-color-tunable-products*.

Color Quality Considerations

Color quality is an encompassing term that includes many individual considerations: the color of the light, the consistency of color from one product to another, the ability to maintain color over time, and the way a light source affects the appearance of objects it illuminates. The building block of all color considerations is the spectral power distribution. A spectral power distribution describes the energy emitted by a source

at the different wavelengths of the visible spectrum (i.e., the colors of the rainbow). The spectral power distribution determines how the light appears, and also how object colors are rendered—since objects do not have color, but simply reflect wavelengths in different proportions. Importantly, it's difficult to look at a spectral power distribution and intuitively understand how the light will look or how colors will be rendered. Thus, we use metrics built around models of human vision that help simplify the complex data in a spectral power distribution to something that is more easily compared and specified.

Color of the Light

Chromaticity

Chromaticity is one way to characterize the color of light. Chromaticity captures the hue and saturation of the light, ignoring the third dimension of human vision (lightness), which is not a property of light sources, only objects. Accordingly, orange and brown could have the same chromaticity coordinates, which is why chromaticity diagrams are used for light sources but three-dimensional color spaces are used for objects.

Each pair of chromaticity coordinates corresponds to a unique color of light, according to a model of human vision. All possible coordinates can be plotted in a chromaticity diagram such as the one shown in Figure 2, and two sources with the same chromaticity coordinates should theoretically appear the same. Is chromaticity a perfect specification for the color of light? Not exactly, because it is a simplification of the spectral power distribution, and relies on models of human vision that were developed for a specific field of view (usually 2°). In different

conditions, or for individual visual systems that vary from the average used in the models, two sources with the exact same chromaticity coordinates can still appear different. This visual difference is usually acceptable, but it is important to recognize.

Aside from specifying the color of light, chromaticity is also used to measure color difference, either between two or more sources (color consistency) or for one source over time (color stability). Different chromaticity diagrams have been developed and standardized by the International Committee on Illumination (CIE), with each attempting to improve on predecessors by more closely matching human color-difference sensitivities for all colors. The goal is for a human perceptual difference to have the same numerical difference throughout the diagram. The CIE 1931 (x, y) chromaticity diagram is still frequently used to specify chromaticity, although the CIE 1976 (u', v') chromaticity diagram is more uniform and therefore superior for determining color difference or color shift ($\Delta u'v'$).

Chromaticity is represented by a pair of numbers, but neither of the numbers alone is related to a perceived attribute. Thus, other numbers that are more intuitive, such as correlated color temperature (CCT) and sometimes the distance to the black-body locus (D_{uv}), are commonly used to communicate the color of the emitted light.

Color Temperature, CCT, and D_{uv}

Color temperature is an important aspect of color appearance related to how "cool" (bluish) or how "warm" (yellowish) nominally white light appears. More specifically, CCT is
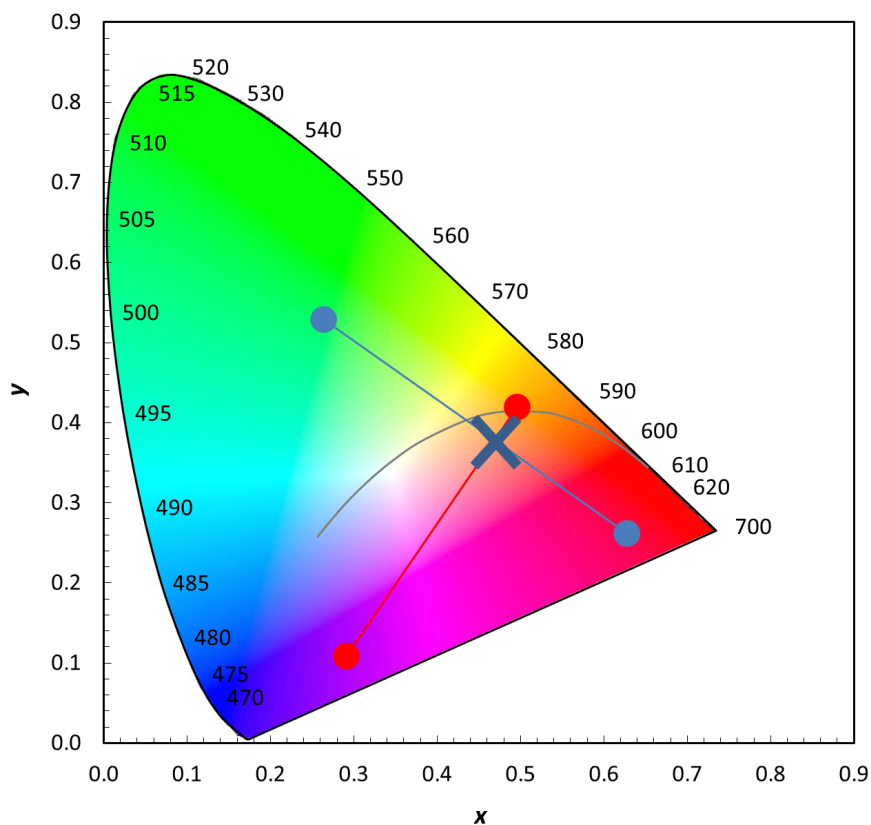

Figure 2. Chromaticity diagrams, such as the CIE 1931 (x, y) chromaticity diagram shown here, are used to provide a numerical specification for the color of light. The colored background is for visual reference only, as chromaticity diagrams do not account for lightness. The numbers around the outside indicate the wavelength of monochromatic light at the coordinate, in nanometers. Colors outside the shaded area are not visible to humans.

A chromaticity diagram can be used for color matching. In this illustration, light sources with the chromaticities indicated by the blue circles and red circles can be mixed to match at the X.

a metric that relates the appearance of a light source to the appearance of a theoretical mass, called a black body, if it were heated to high temperatures. As a black body gets hotter, it turns red, orange, yellow, white, and finally blue. The CCT of a light source, given in kelvin (K), is the temperature at which the heated black body most closely matches the color of the light source in question. It characterizes the appearance of the emitted light, not the color of illuminated objects.

Like all color metrics, CCT distills a complex spectral power distribution to a simple number. This can create discord between numerical measurements and human perception. For example, two sources with the same CCT can look different, one appearing greenish and the other appearing pinkish (see Figure 3), because a constant CCT represents a line in the chromaticity diagram. To address this issue, the American National Standards Institute (ANSI) references D_{uv}, a metric that quantifies the distance between the chromaticity of a given light source and a blackbody radiator of equal CCT. A negative D_{uv} indicates that the source is "below" the blackbody locus, having a purplish tint, whereas a positive D_{uv} indicates that the source is "above" the blackbody locus, having a greenish tint. Together, a specific CCT value and D_{uv} value correspond to a specific pair of chromaticity coordinates.

Frequently, CCT is provided as a nominal designation, such as 2700 K, 3500 K, or 5000 K. These nominal values each cover a range of actual CCT and D_{uv} values, as defined by ANSI in C78.377-2015 and shown in Figure 4A. Two sources with the same nominal CCT can appear very different, and just specifying a nominal CCT will do little to ensure that multiple products emit light that appears similar.

With the 2015 update, ANSI also includes definitions for a different set of bins with more negative D_{uv} values. This reflects recent research suggesting that D_{uv} values lower than those used for the standard nominal designations may be preferred or perceived as more neutral in appearance.

A key aspect of color quality related to human vision is chromatic adaptation. Chromatic adaptation is our eye-brain system's ability to set a consistent white point and discount the illuminant when evaluating the color of objects. That is, regardless of the chromaticity of the light (within limits), we can still identify white, differentiate the color of objects, and make critical evaluations, such as the ripeness of fruits. This functionality is similar to the white balance on a camera, and is a critical evolutionary feature that allows our vision to be effective throughout the day, as the daylight changes.

Color Consistency

To counter variability that is inherent in the manufacturing process, phosphor-converted LEDs are often sorted (binned) post-production, based on chromaticity, lumen output, and sometimes forward voltage. This allows both manufacturers and specifiers of LED lamps and luminaires to receive a more consistent product. Other means for ensuring color consistency, such as binning prior to applying phosphor, may also be used.
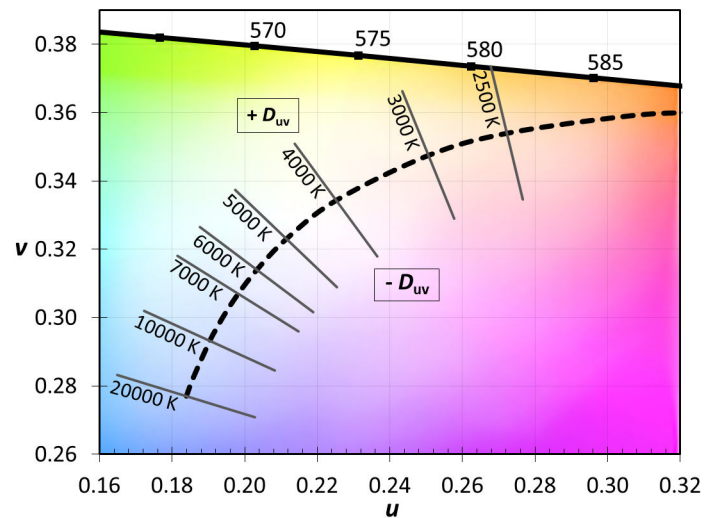

Figure 3. A close-up of the CIE 1960 (u, v) chromaticity diagram, showing lines of constant CCT, which are perpendicular to the blackbody locus in this particular chromaticity diagram. For a given CCT, a source with a positive value for D_{uv} has a chromaticity that falls above the blackbody locus (appearing slightly greenish), whereas a source with a negative value for D_{uv} has a chromaticity that falls below the blackbody locus (appearing slightly pinkish). The lines in this chart represent a D_{uv} range of ± 0.02, which is much greater than ANSI C78.377 tolerances for white light.

The National Electrical Manufacturers Association (NEMA) has published an LED binning standard (SSL-3-2011) based on ANSI C78.377-2008. The variability allowed by the NEMA bin sizes and the ANSI tolerances is roughly equivalent to the chromaticity variation seen in currently available compact fluorescent lamps. Some manufacturers may hold themselves to tighter tolerances. As LED technology continues to mature, binning tolerances remain essential to producing individual products and lighting installations with high color consistency. At this time, there is no industry-wide binning standard for colored LEDs that make up a mixed LED system.

Color Stability (Color Shift)

Separate from color consistency, color stability (also known as chromaticity maintenance, chromaticity shift, or color shift) refers to the ability of a product to maintain its spectral power distribution (and all derived values) over time. As a product ages, the materials in the LED package or optical system can change or deform. Depending on the amount of change, and on whether the change is the same for all products or is different, this can be problematic or simply unsatisfactory, constituting a parametric failure. The color stability of LED products varies based on the engineering of the system, but is usually as good as or better than conventional lighting products over the same time frame. One challenge is that LEDs often have much longer rated lifetimes (based on maintaining light output), so they must maintain their color for much longer too. More information on color shift is available at: *http://apps1.eere.energy.gov/ buildings/publications/pdfs/ssl/2013_gateway_color-maintenance. pdf.*
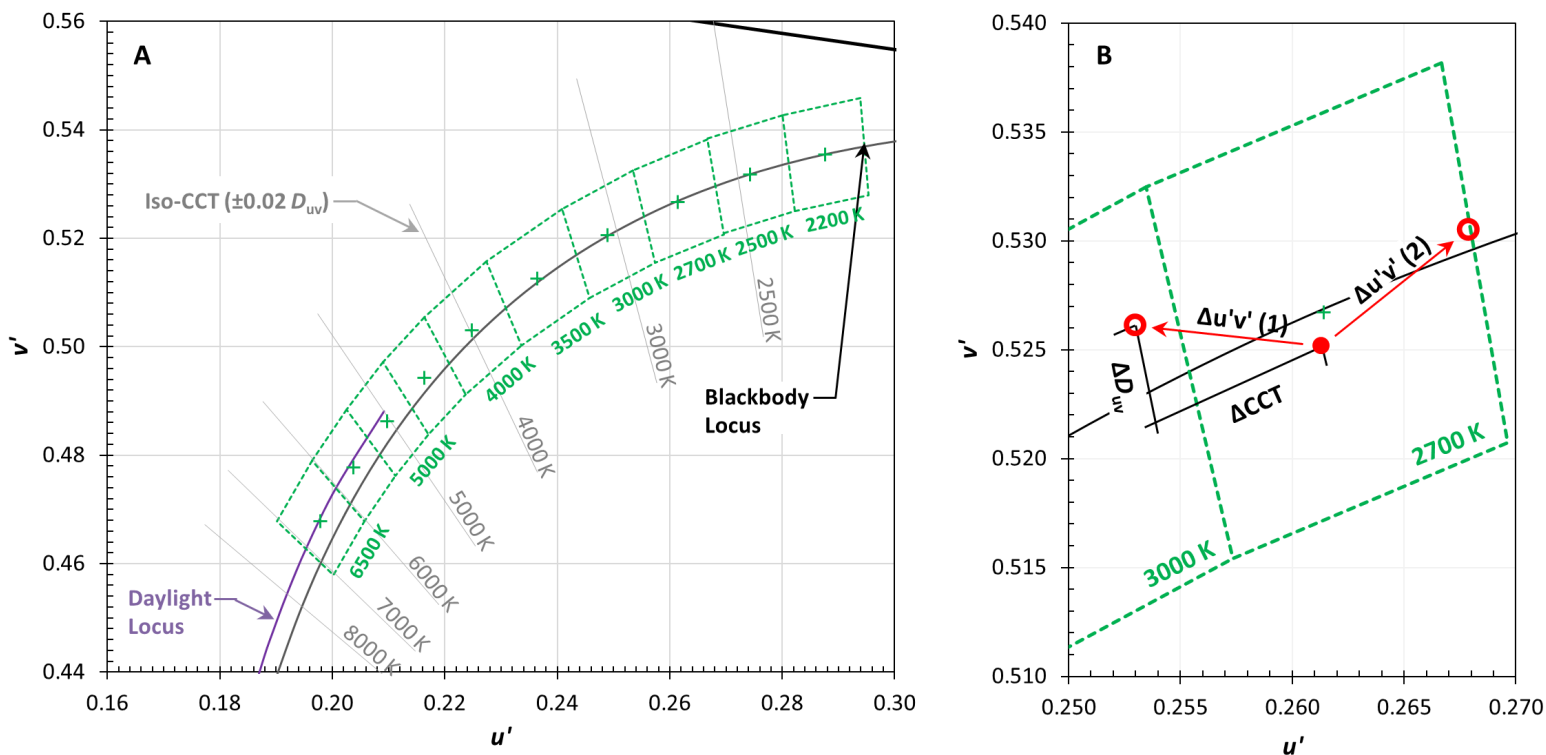
Figure 4. A: ANSI C78.377-2015 defines chromaticity ranges associated with CCT designations for nominally white light sources, using the CIE 1976 (u', v') chromaticity diagram. Each bin represents a range of chromaticity, so sources with the same nominal designation may not match. B: The CIE 1976 (u', v') chromaticity diagram is also used to calculate $\Delta u'v'$, which is a measure of color difference or color change. Importantly, $\Delta u'v'$ is only a measure of the magnitude of the difference, not the direction.

Both color consistency and color stability are best characterized using $\Delta u'v'$, or the difference in chromaticity in the CIE 1976 chromaticity diagram. This value ($\Delta u'v'$) accounts for differences in both CCT and D_{uv}; the difference in either CCT or D_{uv} alone does not characterize the perceptual difference, as shown in Figure 4B. Importantly, $\Delta u'v'$ conveys only the magnitude of the difference, not the direction of the change. Two sources may start at the same chromaticity and have equal $\Delta u'v'$ values but have shifted in opposite directions.

Color Rendition

Color rendition refers to the interaction of a light source's spectral power distribution and the spectral reflectance function of objects, which is a principal factor in determining how the objects will appear (illuminance level, adaptation, surrounding conditions, and other factors play a role as well). This is demonstrated in Figure 5. For any given color of light, there is a wide range of possible color rendering attributes, because there are multiple spectral power distributions that lead to the same appearance of the light. A light source may affect objects' hue, saturation, or lightness, and may do so in various ways for any given hue. Color rendering metrics attempt to quantify this complex change in ways that are easy to understand, yet are meaningful.

At least three aspects of color rendition are relevant to light-source selection and application. These include the accurate

rendition of colors so that they appear as they would under a familiar (reference) source, the rendition of colors such that objects appear more pleasing (sometimes referred to as vivid or flattering), and the ability of a source to allow a subject to distinguish between a large variety of colors when viewed simultaneously. For simplicity, these three facets of color rendering may be called fidelity, preference, and discrimination. The relative significance of these different elements of color rendition depends on the application.

Color Rendition Metrics

With the rise of fluorescent lighting in the 1950s and 1960s, there arose a need to quantify and communicate how color rendering of the new lamps differed from traditional incandescent and other sources. In 1965, the CIE Test-Color Method, commonly referred to as CRI, was first adopted (it was substantially revised in 1974). CRI (or more accurately, the R_a measure) is a measure of average fidelity, or how similarly a light source renders colors compared to the reference source. Among other limitations of outdated color science, CRI doesn't convey any information about the types of differences, just the magnitude of the difference. This means that two sources with the same CRI value can render colors very differently, even to the point where a scene looks appealing under one source and unappealing under the other. See the sidebar for more information on CRI.
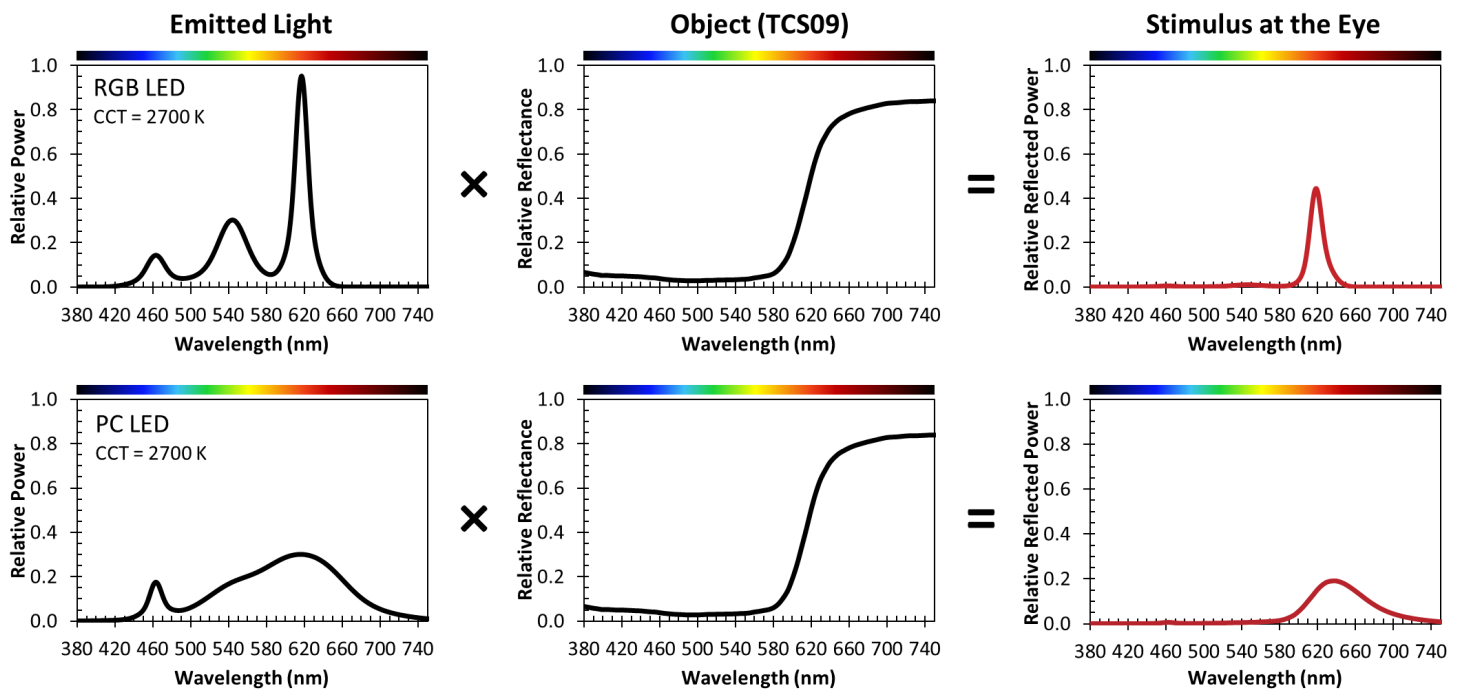
Figure 5. Objects are not inherently colored, but rather reflect different proportions of radiant energy. Thus, if the light incident on a surface changes, the apparent color of the object may change. Depending on the light source, this shift can be small or large. As shown above for example RGB and phosphor-converted (PC) LED sources, the stimulus at the eye can be calculated by multiplying the spectral power distribution of a lamp by the reflectance distribution of the object (in this case TCS09, which is used to calculate R_9).
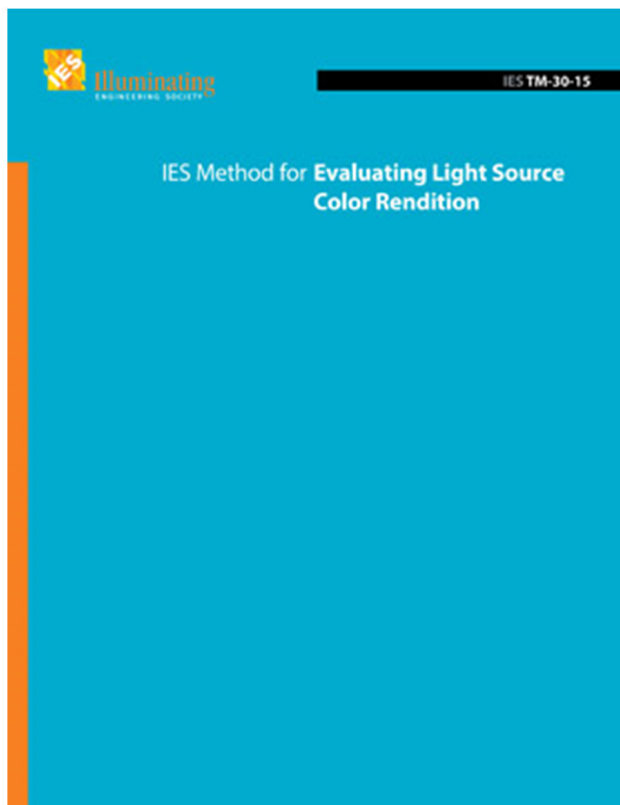
IES TM-30-15 was approved by the IES in 2015. For more information, see: *http://energy.gov/sites/prod/files/2015/12/ f27/tm-30_fact-sheet.pdf*

Since it was adopted, dozens of alternatives to CRI have been proposed. Some of the proposed measures have focused on color preference or related quantifications, whereas others have focused on improving characterizations of color fidelity. However, until recently none of the proposed measures was adopted by the CIE or any other standards-development body. In 2015, the Illuminating Engineering Society (IES) adopted IES-TM-30-15, IES Method for Evaluating Light Source Color Rendition. TM-30 is an evaluation framework, providing a suite of color rendering characterizations that can be used together to understand how a source affects objects' appearance and, ultimately, allow the user to make a more informed decision about the suitability of a product for an application. The core measures are average fidelity (R_f), average gamut (R_g), and detailed information for 16 different hue ranges (presented pictorially in the Color Vector Graphic and numerically with hue-specific values for fidelity and saturation). More information on TM-30 components and their meanings can be found in a separate DOE SSL fact sheet: *http://energy.gov/sites/ prod/files/2015/12/f27/tm-30_fact-sheet.pdf*

At present, TM-30 is not a required standard, and it is not yet included in any minimum color quality specifications, such as those provided by ENERGY STAR. Its components have been proposed to the CIE for recognition as an international standard and direct replacement for CRI, with resolution expected by 2017. Undoubtedly, the transition from CRI to new measures will take time, as regulations and specifications must be

Color Rendering Index

Although standardized alternatives now exist, the CIE Test-Color Method for evaluating color rendering—and its principal component, R_a (commonly known as CRI)—is still embedded within numerous specifications and regulations. Thus, it's important to understand the features and limitations of the tool.

The CIE Test-Color Method utilizes eight standard color samples—having moderate lightness and of approximately equal difference in hue (i.e., equal spacing on a chromaticity diagram)—and six special color samples. For each color sample, the chromaticity under a given (test) source can be compared to the chromaticity under a reference source of equal CCT, allowing for the measurement of color difference, which is then mathematically adjusted and subtracted from 100 (R_i). The principal metric of the CIE system is R_a (commonly called CRI), which averages the R_i scores for the eight standard test colors and typically has a range from 0 to 100, though negative scores are also possible. A score of 100 indicates that the source renders colors in a manner identical to the reference. In general, a source with a CRI in the 70s has been considered acceptable for interior applications, whereas scores in the 80s have been considered good (specification grade) and scores in the 90s excellent. However, recent research has demonstrated that considering only average color fidelity, which CRI characterizes, provides little benefit for selecting the best source for an application. More specifically, reliance on outdated color science often leads to CRI unduly penalizing sources that saturate reds, which is often preferred by consumers. Sources can also be optimized to render the eight color samples used for R_a with high fidelity, but that doesn't necessarily translate to a broader range of colors. The CIE has acknowledged the limitations of CRI: *http://www.cie.co.at/index. php?i_ca_id=981*.

The special color rendering indices, referred to as R_9 through R_{14}, are each based on one of the six special color samples. They are not used for calculation of CRI, but may be used for supplemental analysis when necessary. The "strong red" color sample, TCS09, and its associated fidelity metric, R_9, are especially pertinent, because the rendition of saturated red is particularly important for the appearance of skin tones, among other materials. An R_9 score greater than 0 is generally considered acceptable, because the color space used in the CIE Test-Color Method often causes color shifts in the red region to be exaggerated.

adapted. Because TM-30 offers a more nuanced platform for communicating color, manufacturers and specifiers must also adjust to take advantage of the new capabilities.

Still, many lighting manufacturers are already reporting TM-30 measures, aided by the fact that the measurement used to calculate the value—a spectral power distribution—is already required for calculating a range of lighting metrics. Likewise, specifiers are beginning to use TM-30 to more critically evaluate sources and differentiate between them. Research demonstrating how the evaluation tools can be used to identify preferred sources for given applications is underway. It will take time for a large body of research to develop—as was the case with CRI. A suite of information on TM-30 can be found at *http://energy.gov/eere/ssl/tm-30-frequently-asked-questions*.

The Big Picture

Color quality is just one aspect of lighting. Along with other quality aspects such as flicker, glare, and light level, it must be balanced against such factors as energy use and cost when a product is being evaluated. Other spectral effects are also important to consider, such as the relationship between light and health. In some cases, color quality may be a primary consideration, but in others, it might not be a consideration at all. Some situations call for emphasizing color consistency, whereas in others, value may be maximized by focusing on color rendering. Likewise, in some cases color fidelity may be the key factor, whereas in others, increasing preference by increasing red saturation may be important.

The bottom line is that lighting engineers and specifiers must balance many individual characteristics when evaluating a product. Although it may be desirable, there is no formula that can consolidate the numerous individual attributes and deliver a ranking of the best products—even if the list of attributes is narrowed to only color rendering features. Further, all metrics have limitations that must be understood in order to effectively apply the data.