Postings: from the desk of Jim Brodrick

I've written before about the color quality of white light, and how it's been one of the key challenges facing SSL. Two metrics have traditionally been used to convey white light's color quality: Correlated Color Temperature (CCT), which describes the hue of the light itself, and the Color Rendering Index (CRI), which indicates how well a light source renders the colors of illuminated objects. CCT alone is inadequate for describing SSL, because two LED light sources with identical CCTs can render object colors very differently due to the differences in spectra. And CRI value is poor at predicting the quality of the appearance of saturated red objects, and doesn't correspond well to human perception of color quality. Using these two metrics together only gets us in the ballpark of understanding white-light source color; it's still necessary to see the light source with our own eyes to evaluate it for a particular application.

That's because other things come into play. These include color shifts that increase color saturation, which can enhance our ability to distinguish between colors (color contrast) and can often make objects more visually appealing. In addition, because a visual scene appears less colorful at lower illuminance levels, sources that enhance color slightly in indoor lighting may make those colors look truer to how they would appear in the higher illuminances of daylight.

To address these and other weaknesses of the CRI, the National Institute of Standards and Technology (NIST), with input from the lighting industry and the International Commission on Illumination (CIE), has developed a new metric that more accurately communicates the color quality of lighting products. Called the Color Quality Scale (CQS), it correlates well with visual evaluation and indicates more appropriately how people actually see and prefer the color rendering of the light source.

To understand why the CQS is the superior tool, it's necessary to look a bit more closely at the CRI, which is an inconsistent predictor of color rendering ability because it's based on eight pastel colors that don't represent the full range of colors that we experience. The CRI is especially poor at predicting the rendering of complex colors – such as skin tones, wood finishes, fresh produce, and woven fabrics – because it doesn't give critical long-wavelength reds (like the deep red of apples) enough weight in its calculation. Thus, CRI scores can be good even if the color rendering of skin, oriental carpets, or purple glass – all of which need long-wavelength red to look appealing – is poor.

By contrast, the CQS uses a set of 15 color samples that have much higher chroma (deep color) than those of the CRI and span the entire hue circle in approximately even spacing. And while color differences for all reflective samples in the CRI are averaged, which makes possible a high score despite poor rendering of one or two colors, the CQS uses a different calculation to ensure that any large hue shifts significantly impact the overall score.

There are other problems with the CRI that the CQS addresses, such as the fact that the CRI penalizes light sources that enhance color contrast, which is preferred by consumers and can be accomplished with LED lighting using RGB peaks. But I won't bore you with any more details here. Suffice it to say that while the CRI attempts to describe color fidelity, it ignores other aspects of color quality, such as color discrimination and observer preferences. The CQS is designed to incorporate these other aspects of color appearance.

To compensate for the CRI's shortcomings, some people advocate supplementing it with other metrics, such as R9. But this is impractical. The bottom line is that we need a single metric that describes the color quality of white light, without leaving any gaping holes. I've had an opportunity to visit NIST's color laboratory and see the CQS in action, and I can tell you it fills the bill. Regardless of the type of light source, the CQS represents the color rendering qualities of white light more accurately than the CRI and is a far better predictor for colors that have a high red content, such as skin color and wood finishes – which is one of the CRI's major weaknesses.

In my mind, and to many others in the field, switching to the CQS should be a slam-dunk. But the CIE technical committee on color rendition of white light sources (TC 1-69) is instead considering an alternative metric that's very similar to the traditional CRI and that won't solve the problems of characterizing color quality for SSL sources nearly as well as the CQS would. The CRI has been widely used in the lighting industry for more than 40 years, and lamp manufacturers are concerned about changes in its scores for existing lamp products – despite the fact that the CQS has been shown to maintain consistency with CRI scores within a few points for most of these lamps, but may improve several points for sources that render colors especially well.

The reality is that the lighting industry needs a metric that is both accurate and easy to use. My hope is that the CIE will vote to adopt the CQS soon, so that LED manufacturers will have an added incentive to tune their products to deliver excellent color rendering that's flattering to people and objects, instead of aiming at highest efficacy at the expense of good color.

The issue is thorny and complex, and has been neatly summarized and put into perspective by Yoshi Ohno and Wendy

Davis of NIST, in a paper that's posted online. What's at stake is something that will have considerable impact on solid-state lighting in the years to come. I welcome your thoughts on the topic (or on any other), which you can send me at postings@lightingfacts.com.

James R. Brodrick