Light Profile Uniformity in Linear Lighting Applications

Michael J. Murdoch[1], Susanne Seitinger[2], Eric Roth[2], Peter Goldstein[2], Ulrich Engelke[3]
[1]Philips Research, Eindhoven, The Netherlands; [2]Philips Color Kinetics, Burlington, MA, USA;
[3]Commonwealth Scientific and Industrial Research Organisation (CSIRO), Hobart, Australia

Abstract

LED lighting in a linear form factor can provide valuable accent lighting and vertical illumination in architectural cove and graze applications. Ideally, linear lights provide a smooth gradient of light onto a surface without disturbances along their length. However, socket shadows between discrete fixtures and multiple LEDs per fixture can lead to visible nonuniformities, often worse near the light source but mixed to uniformity further away. Perceived mixing distances in the illumination patterns of a variety of linear light sources were assessed visually, and objective metrics based on illumination measurements were developed. Accurate predictions of mixing distances are shown based on trends in CIE DE2000 color errors computed between the measured light pattern and analogous Gaussian-weighted local neighborhood regions. Effects of variations in parameters of DE threshold and neighborhood size and shape are discussed.

Background

Visually interesting lighting is becoming more widespread, driven by solid-state technologies in a variety of form factors. This means that beyond simple functional light, lighting systems are used to highlight architectural details, provide luminance and contrast on vertical surfaces, and create atmosphere or ambience. Typical applications are flood lights that wash a building façade with colored light, linear lights that provide a highlight in a cove or graze a wall to bring out texture, spot lights that highlight art objects or building details, and sequences of downlights that create scalloped patterns along a wall. Further details on lighting technology and applications may be found in [1]. As these applications have become more common, awareness of quality of light and the value of visual interest has increased. Vertical illumination, that is light falling on vertical surfaces such as interior walls, is valuable for many reasons, for example as a contributor to percepts of clarity and spaciousness in architectural spaces [2] and to introduce color and contrast necessary for atmosphere creation [3], [4], [5]. Linear light sources can be a valuable source of vertical illumination in architecture.

What may be simply thought of as beams of light on surfaces requires care in implementation, lest visual artifacts arise. This paper discusses ideal light profiles, the deviations from ideal which can be considered artifacts, and the perception and measurement of nonuniformities in the light profiles of linear light sources used for architectural cove or grazing applications.

Ideal Light Profile

While it may be impossible to define an ideal light profile given the myriad applications which may be found, it is worth considering what is generally good for the case of linear lights relatively close to a wall. In this situation, linear lights generally provide a light pattern that is consistent along the source's linear orientation and a gradient perpendicular to it. In some applications the gradient is short, to provide a small accent, and in others a longer gradient is used to allow the light to reach further onto the wall. In most cases, an assertion is that the light profile should be invariant in the longitudinal direction—that is, parallel to the source's linear orientation. In the lateral direction, a monotonic, smooth gradient of light fading away from the source is a logical starting point. Perhaps this falloff should be linear in luminance or Lightness, or with a Gaussian or other sigmoid profile, but such specific aims are difficult to justify. More justifiable is the general goal of a monotonic, smooth (for example with continuous first and second derivatives) intensity gradient.

Deviations from an ideal light profile could be considered artifacts. In practical applications, linear lights may exhibit artifacts in the longitudinal direction because they are necessarily composed of multiple discrete light sources and/or multiple discrete fixtures. Visible artifacts may include luminance variations between LEDs, socket shadows between fixtures, and color variations where the light contributions from different, particularly differently-colored, LEDs are not fully mixed. In the lateral direction, artifacts such as a lack of smoothness or monotonicity may be seen, like bumps in the expected gradient. Interestingly, color and luminance artifacts in the longitudinal direction generally blend together to uniformity further from the source in the lateral direction. Perception of the mixing or blending of artifacts is particularly interesting and was studied in the present example.

Uniformity

Color uniformity has been studied in many imaging and lighting applications. In lighting, it is common to quantify the uniformity of a light source or light pattern based on chromaticity differences. Exemplary is the Energy Star chromatic uniformity norm for solid state lighting which requires that chromaticity measured in different angles (or equivalently different spatial locations in the light pattern) be no further than 0.004 from the luminance-weighted average chromaticity point on the CIE 1976 (u', v') diagram [6]. The use of thresholds in delta-u'v' is closely related to the work of MacAdam, who in 1942 published standard deviation of color mixing (SDCM) ellipses drawn in CIE 1931 xy chromaticity space, showing the relative sizes of perceptually-equal color differences on the perceptually-nonuniform xy diagram [7]. Despite SDCM being based on one observer, despite the better but still imperfect uniformity of u'v' chromaticity, and despite the lack of any spatial size or proximity consideration in similar metrics, they persist in specifications, LED binning, etc.

Many other approaches to quantifying uniformity have been proposed in both lighting and imaging. Contrast sensitivity functions (typically only luminance-based, but with some limited data on chromaticity-based) can be used to estimate visual sensitivity to frequency bands present in a stimulus, whether that is a print, display, or light profile on a wall. Mura, mottle, and other nonuniformities in LCD and OLED displays have been quantified empirically.

Measurement & Simulation of Light Profiles

Crucial to modeling the perception of artifacts and nonuniformities in light profiles is the physical measurement of a light profile on a surface. Light on a surface can of course be quantified in two orientations: in terms of illumination, meaning the amount of light reaching the surface, or in terms of luminance, meaning the amount of light reflected from the surface in a particular direction. For a perfectly diffuse (Lambertian) reflector, these two differ by a factor of ρ/π, where ρ is the reflectance factor and π accounts for the directionality of luminance.

It is worth noting that goniometry is a common source of angular photometry data for lamps and luminaires. Because goniometry treats the source as a point from which all rays diverge, and because it typically measures in the “far” field, goniometry data may be equivalent for a light profile reaching a relatively distant plane but are unlikely to be similar in the near field. Deviations from the distant point assumption, such as extended luminaire size, multiple light sources, and multiple optics (which are all typical for linear sources) mean that much interesting near-field activity like crossing rays and source mixing is missed. This is a problem for modeling the behavior of linear light sources close to a wall, where they are most often used.

In the present work, characterizations of light patterns were made by taking an array of measurements over the surface of interest. These were done colorimetrically, in both luminance and illuminance orientations. Each is tricky for different reasons: with luminance measurements, the surface reflecting the light has a huge influence on what the instrument records; with illuminance measurements, there is no surface to worry about but the instrument must be reliably moved throughout the plane of interest. Each is interesting for different reasons, as well: a luminance measurement more directly represents what the observer actually sees in application, light on a wall, assuming the wall is representative; an illuminance measurement specifically characterizes the light source itself without regard to the wall, so it is conceptually more direct.

Luminance Measurement

Measurement in the luminance orientation is in essence just a picture of the light pattern on a wall. However, the wall itself has a huge influence on the light it reflects: a selective spectral reflectance will change the color of the light; a non-Lambertian angular reflectance profile will affect the measurement depending on the luminaire and instrument location; texture on the wall will introduce localized variations, shadowing, or other problems; a lack of flatness will imply that the slice of the beam being measured is not planar; and a nonuniformity in the wall (different color or texture) will confound any nonuniformity in the light beam itself. An ideal wall would be perfectly flat, perfectly smooth, and perfectly Lambertian, with a uniform and non-selective spectral reflectance. A flat, matte-painted wall or a tensioned projection screen with gain less than one may approximate this ideal.

In our work, luminance-oriented colorimetry over an array of points on a surface was measured with an imaging colorimeter, specifically a Radiant Imaging ProMetric PM-1423 with a 35mm f/2 lens. This device directly delivers images with 14-bit data in 3 channels corresponding to CIE 1931 XYZ tristimulus values. At a distance of 2.5m (corresponding to a typical viewing distance for indoor fixtures), the measurement area was 1.0 x 0.7m. With this arrangement the camera’s 1024 x 768 pixels correspond roughly to 1 pixel per mm on the wall.

Illuminance Measurement

Measuring in the illuminance orientation requires an instrument at the point being illuminated, oriented toward the light source rather than toward the light pattern. To create an image-like array of measurements, the instrument is moved in a virtual plane, sampling the illuminance at the implied surface. In our work, this was done using a StellarNet Black-Comet spectrophotometer on a computer-controlled X-Y transport stage. The x-y spacing of measurements can be chosen arbitrarily using the X-Y stage, with obvious implication for measurement time; in our work 1cm was sufficient.

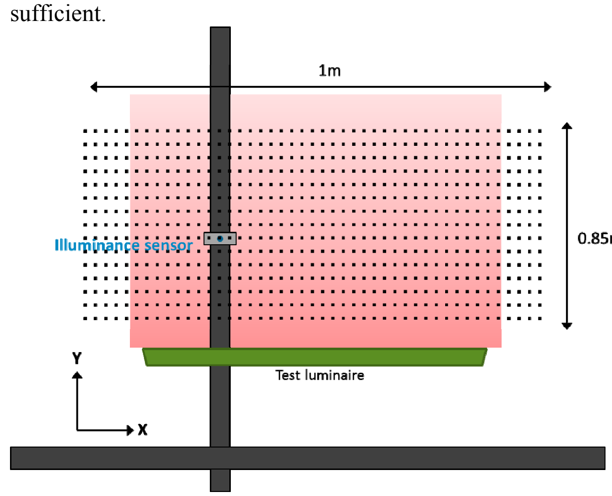

Figure 1: Measurement geometry for illuminance measurement: a spectral illuminance sensor is moved in a virtual plane by an X-Y transport stage to capture an array of measurements.

Luminance & Illuminance Simulation

Luminance or illuminance distributions may be simulated with a variety of physically-based path-tracing tools, including optical engineering packages such as LightTools[1] or image rendering packages such as Indigo[2] or Maxwell[3]. With LightTools, a detector plane takes the place of a motion-controlled illuminance sensor and can produce a sampling at arbitrary resolution of any plane in a light distribution. In image rendering programs, a plane that is part of the virtual scene geometry may be photographed in a virtual analogy of the luminance measurement described above.

---

[1] Synopsys Optical Solutions: http://optics.synopsys.com/lighttools/
[2] Glare Technologies: http://www.indigorenderer.com/
[3] Next Limit Technologies: http://www.maxwellrender.com/

Simulations in either category may model the full light path from LED die, through all optics, to the plane of interest, or they may be based on angular measurements such as goniometry. As was mentioned previously, because goniometry typically treats the light distribution as if it comes from a point, it is likely not sufficient for near-to-far-field light profiles such as the linear lights discussed here.

Experiment: Linear Light Sources

Linear light sources have become common in architectural lighting for cove and graze applications, where they are mounted close to a wall or other surface and provide a gradient light profile onto it. Generally with this type of fixture, the light profile is less uniform near the light source and more uniform further away. In this experiment, mixing distance was assessed visually for a variety of linear light fixtures, and these results were used as a basis to develop and evaluate objective metrics computed from analogous physical illuminance measurements.

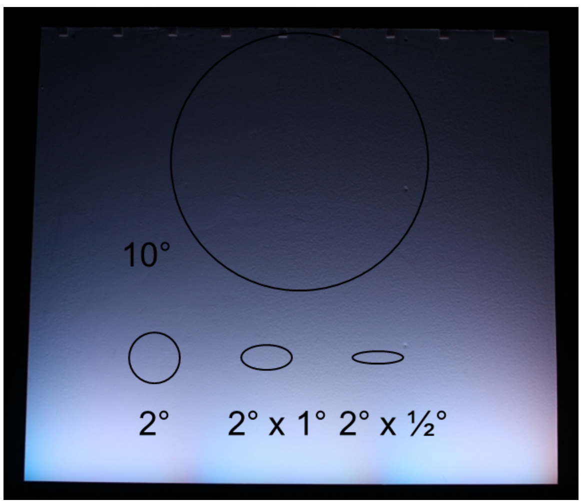

Figure 2: Photo of stimulus with illustration of surround regions: for size reference, overlaid ellipses are labeled in visual angle, assuming a 2.5m viewing distance, on the 90x90cm wall used in the visual experiment.

Mixing Distance Experiment

The visual experiment was conducted with 12 sets of linear fixtures illuminating a flat wall, such as can be seen in the annotated photo in Figure 2. The observers viewed each light profile from a fixed distance of 2.5m, and their task was to identify the "mixing distance," or the height above the fixtures where they saw that the color and/or luminance variations evened out. 12 different linear fixture types were used, 6 of which employed separate red, green, and blue LEDs, and 6 of which employed fixed white LEDs. The selected fixtures varied in beam angle, color temperature, and type of optics, providing a wide range of visual quality in the experiment, valuable for exercising an objective metric. The RGB and white fixtures were tested in separate experimental sessions: 24 observers participated in the former and 25 in the latter, with an overlap of 14 individuals. In this pool of 35 observers, the average (and median) age was 39 years, and the male to female ratio was 26:9.

A visual summary of this experiment can be found in Figure 3, which shows the twelve stimuli (measured colorimetry converted to sRGB images) with mixing distance experiment results overlaid. The yellow curves show the smoothed histogram of observers' mixing distance assessments, and the white line shows the mean mixing distance over observers with a 95% confidence interval. It can be seen that different stimuli exhibit widely different mean mixing distances, and the inter-observer variability shown by the histogram spread and confidence intervals is wider for RGB fixtures (G-L) than for white fixtures (A-F).

Objective Metric

An objective metric was desired, computable from measurements, which can predict the visual mixing distance results. The same 12 linear light fixtures were measured colorimetrically using the X-Y transport-driven illuminance instrument described earlier. A measurement grid covering 94x80cm with spacing of 1cm characterized a region of the light profile corresponding to the wall used in the visual experiment. From the resulting raw measurements, a variety of candidate metrics were computed and compared.

Predictions of the visibility of color differences can be based on either or both spatial sensitivity or color difference sensitivity. Both are relevant to the present situation, but the spatial size of the color errors is relatively large (spatial frequency less than half a cycle per degree) so color difference metrics were chosen. Generalizing this, local color error metrics can be thought of as the difference between a spot measurement and its neighborhood, where the neighborhood could be the entire pattern of light, the adjacent measurement, or a weighted average of a surrounding region. Further, the color error metric can be a raw colorimetric difference such as delta-u'v', or it can be a perceptual color difference such as CIE DE2000.

Because delta-u'v' ignores luminance variations, we focused on CIE DE2000, computed for every value in the measurement grid: the measured color compared with the Gaussian-weighted average of a neighborhood region. Note that for all CIELAB computations the reference white used was the luminance-weighted average chromaticity with average luminance. In practice this resulted in L* values ranging from about 20 to about 120. The estimation of mixing distance from DE2000 values is shown in Figure 4. The left plot shows contours of local DE2000 values versus x-y position in the light pattern. The right plot shows summary statistics of these DE2000 values versus vertical position: the smoothed mean (red) in a pink region indicating the 10^{th} and 90th percentiles. The percentile range is simply informational, but the mixing distance estimation is where the mean curve crosses a DE threshold value of 1.5, computed via a simple zero-finding algorithm. In this example the estimated mixing distance (vertical dashed line) is almost exactly equal to the mean assessed mixing distance (vertical black line, with 95% confidence interval shown).

In these analyses, the neighborhood region size, statistical metric (i.e. mean, median, or another percentile), and metric threshold all have strong effects on the mixing distance estimates.

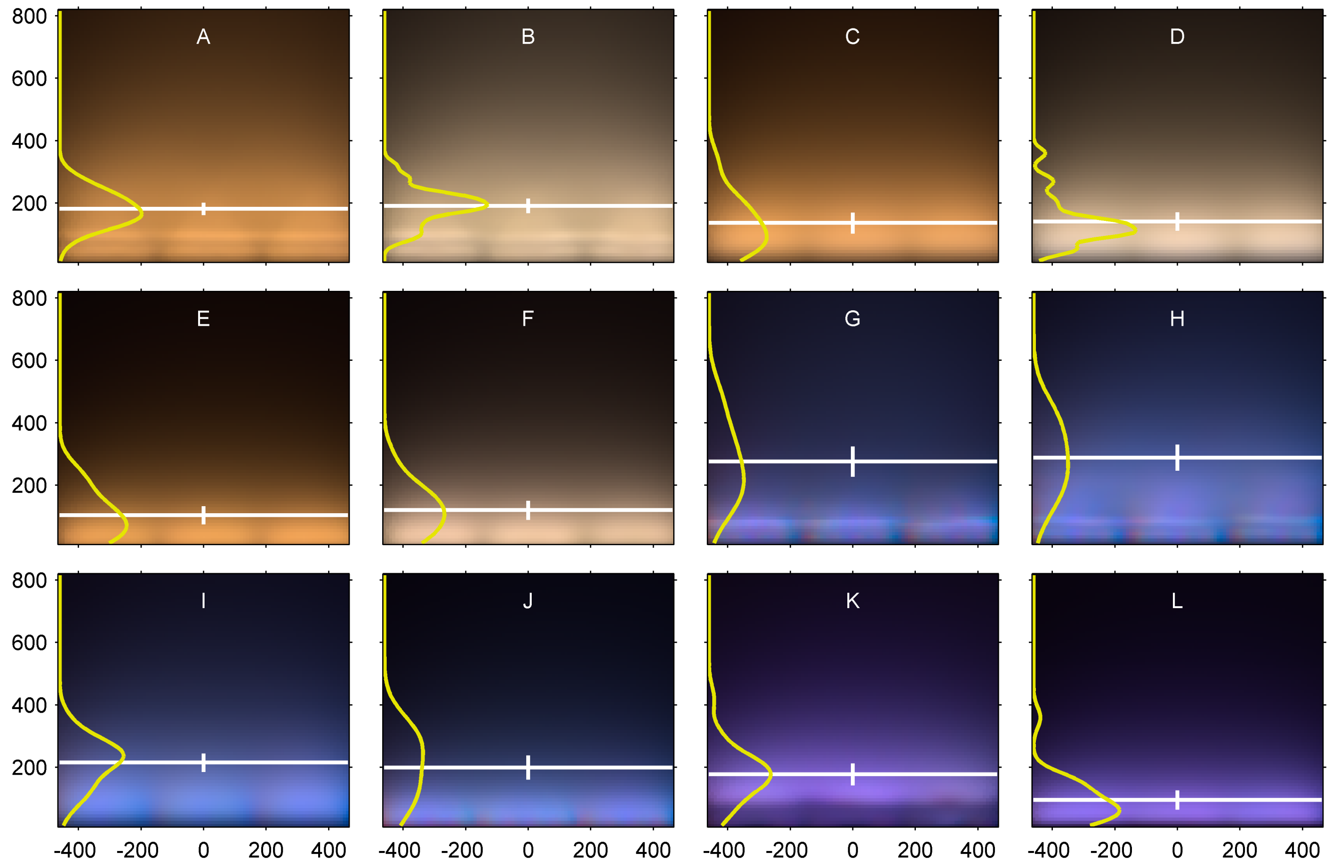

Figure 3: Experimental stimuli: images computed from illuminance measurements show the twelve light profiles used in the visual experiment, on axes in mm, overlaid with the mean assessed mixing distance with 95% CI (white) and smoothed histograms of observers' estimates (yellow). Fixtures A-F use only white LEDs at one of two correlated color temperatures, while G-L use RGB LEDs.

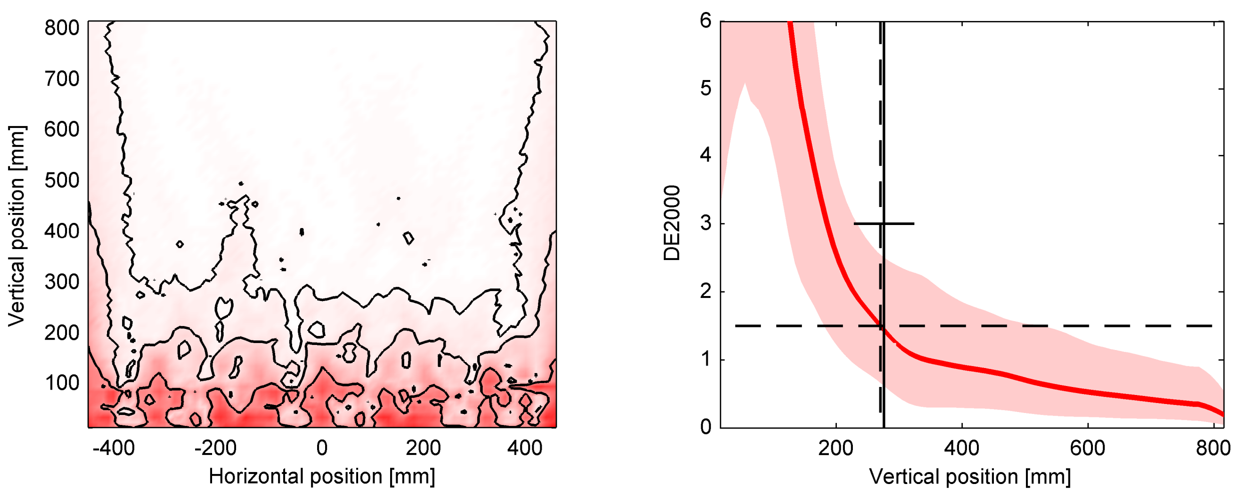

Figure 4: Example mixing distance computation for sample G: at left, DE2000 between each measured point and its neighborhood is shown versus x-y position, shaded with redder colors indicating larger local differences and with contours at DE values of 1, 3, and 10. Large deviations near the bottom of the light pattern diminish going upward. At right, the summary trend of DE2000 values versus vertical position is shown: the mean value (red) is drawn over a pink region indicating the range from 10th to 90th percentiles. The mixing distance is estimated where the mean curve crosses a threshold DE value of 1.5. In this example the estimated mixing distance (vertical dashed line) is slightly lower than the actual mean assessed mixing distance (vertical black line, with 95% CI).

The choice of neighborhood region is critical: of course, the larger the region, the more sensitive the measure. Further, because of both the strong vertical luminance gradient present and the observers' task to identify the mixing distance as a vertical position, an asymmetric region was considered, thereby giving more weight to horizontal neighbors and less weight to vertical. This took the form of a 2D Gaussian weighting function with different widths (standard deviations) in the x and y directions. The metric and threshold also affect the estimate: most statistical summary metrics behave roughly monotonically as shown in Figure 4, but some are noisier than others. It may seem obvious to choose a Delta E value of 1 as a threshold; however, split-field color matching like that which led to the definition of a Delta E is a very different task than assessing color differences with fuzzy edges in a non-uniform field. In fact for the RGB fixtures which have some fairly chromatic activity near the fixture itself, it appears that a threshold higher than one is better.

Results

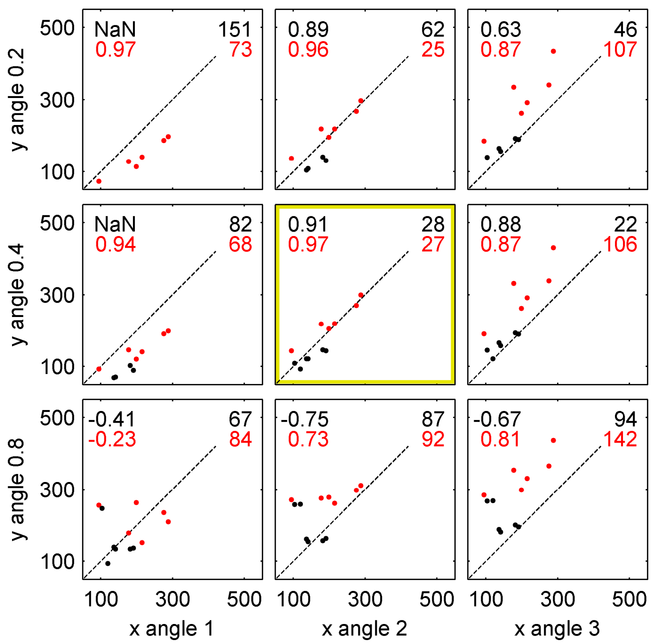

Figure 5: Analysis highlights for RGB (red) and white (black): Each subplot shows the estimated vs. actual mixing distances in mm for 6 RGB (red dots) and 6 white fixtures (black dots) along with Pearson correlation values and RMSE values (left and right, respectively, colored accordingly) for a threshold mean local DE2000 of 1.5. Each subplot corresponds to a neighborhood region of angular Gaussian size: x angles of 1, 2, and 3 degrees visual angle, and y angles of 0.2, 0.4, and 0.8 degrees. The plot at (x, y) angles (2, 0.4) shows the best combination of high correlation and low RMSE for RGB and white.

Analyses were done with many variations of region size and shape and metric threshold, and fits were assessed using Pearson correlation and RMS. For the RGB fixtures, the best combination of parameters was a DE2000 threshold of 2 and an angular Gaussian size in (x, y) of (2.5, 0.4) (see Figure 2 for a visual size reference), resulting in an RMSE level of 20mm, which is less than the mean 95% confidence interval on the observed mixing distance for RGB fixtures, 37mm. For white, the best fit was attained with a DE2000 threshold of 1 and a Gaussian size in (x, y) of (2, 0.2), giving an RMSE of 13mm, less than half the mean 95% confidence interval on the observed mixing distance for white fixtures, 27mm. An effort was made to find a single threshold and region definition which could be used for both RGB and white. With a DE2000 threshold of 1.5 and a Gaussian (x, y) angular size of (2, 0.4), RMSE levels of 28mm for white and 27mm for RGB were obtained. This single best configuration is thus slightly less accurate for RGB and moderately less accurate for white, but a good single compromise for both.

To illustrate the observed trends, Figure 5 shows a sampling of the results obtained with a DE2000 threshold of 1.5. Each subplot shows the estimated versus actual mixing distances (if perfect, these would follow the dashed 1:1 line) for a Gaussian-weighted region with specific x and y standard deviations, with indications of the corresponding correlation (higher is better) and RMSE (lower is better) values. The optimum Gaussian size of (2, 0.4) is highlighted. It is apparent that in general the RGB fixtures (red) are overestimated while the white (black) are underestimated. Further, as x angle increases, mixing distance becomes more overestimated, even if in many cases the correlation is high. As y angle deviates from the optimum, correlations decrease and in some cases an estimate cannot be found because the expected threshold is not reached – indicated where the correlation value is given as NaN.

Discussion

The preceding analysis shows that it is possible to accurately predict visually assessed mixing distance. For RGB and white fixtures, the mixing distance is estimated from colorimetric data with an RMS error less than the 95% confidence intervals on the visually assessed mixing distance. Some aspects of this approach are worth probing further.

RGB versus Luminance Mixing

The first point to consider is that, while a single DE2000 threshold can be used with good results for all fixtures, the DE2000 threshold used to reach the most accurate estimate of mixing distance was different for RGB and white fixtures. Making such a distinction with a relatively small set of light fixtures may be premature, but it is worth considering if it is more valuable to have a single metric for all, or tailored metrics with slightly more accuracy per category. In application there is strong value in the simplicity of the former, but it is worth further thinking about why, fundamentally, we might expect a difference in perception between the nonuniformity and mixing of white and RGB. It may be that because of the significant amount of chromatic variation visible in the lowest part of the RGB light profiles, adaptation to such a chromatic variation may make the more-mixed region above look less chromatic than it might have otherwise. Another possible contributor is the luminance falloff, typically to 10-20% of the max luminance near the assessed mixing distance, which may make the color differences less noticeable.

Effect of Adaptation

Adaptation is always difficult to estimate in lighting applications. Unlike in prints, where paper white may usually be safely assumed to be a reference white, reference white in a gradient light profile is more ambiguous. Using the average luminance with the luminance-weighted average chromaticity is a justifiable starting point, but the larger the light profile, the more likely the eye adapts to local color and intensity as it scans. A color appearance model which takes surround and background into account might make this more accurate. That being said, a slightly different reference white wouldn't wildly change the magnitude of Delta E computations.

Another consideration is the use of chromaticity-based color difference metrics, such as delta-u'v' or SDCM, which don't explicitly require a reference white. Though, SDCM involves an implied illuminant C white reference. Many uniformity and tolerance studies have used these metrics, however their main limitation (alternatively, main benefit for some applications) is that they ignore the luminance dimension. In the present example, we found that delta-u'v' predicts the RGB fixture mixing distance quite well, but of course because there is scant chromaticity variation in the white fixtures, the metric stumbles. A metric which fails for a whole class of light fixtures is of limited use.

Visibility versus Acceptability

A distinction which is always important, especially to customers and marketers, is that between the visibility and acceptability of an artifact or shortcoming. An acceptability threshold is usually some factor higher than a visibility threshold, but that factor can depend strongly on the audience deciding what is acceptable as well as the application, form factor, and cost/quality expectations. Visibility experiments are much more robust to observer populations and thus a better place in which to build models. Later, the acceptability factor for a given target audience may be determined with a much simplified experiment.

Related to this, it is important to remember the impact of the wall illuminated by the light profile in question. In many architectural applications, stucco, brick, stone, or other "busy" textures are present, which significantly reduces the visibility of all types of local artifacts. Analogously, it is known that uniformity errors in display backlights are much less detectable in natural images than they are in uniform fields. In both cases, as with acceptability, there is generally a factor by which thresholds may be increased for such applications.

Future Directions

The metrics discussed in this paper provide an excellent match to the visual experiments based on the measured data for these light profiles; however, generalization can be risky and improvements are always possible. Localized CIE DE2000 is clearly a robust and perceptual metric, but other possibilities could be considered. Metrics based on color gradients, [chromatic] contrast sensitivity to spatial frequency content, or a comparison to a target ideal light profile may be worth exploring. Further, other form factors and light profiles are valuable to understand, such as point lights with radial light patterns. A good goal would be a single model of the visibility of nonuniformity artifacts which is applicable to all types and directions of gradients.

Conclusion

In this work, CIE DE2000-based local color difference metrics have been developed which accurately model visually-assessed mixing distances in linear light profiles. However, the strong effect of model parameter settings on the fits signal caution for the immediate generalization of these results. Absolute best model fits for white and RGB LED linear light fixtures were achieved using different DE2000 thresholds of 1 and 2 respectively, while a slightly poorer fit was reached using a single DE2000 threshold of 1.5. All of these use similarly-sized asymmetric neighborhood regions of about 2x0.4 degrees of visual angle for the color difference computation.

The visibility (and eventually acceptability) of chromatic and luminance nonuniformity artifacts in light pattern gradients may both be quantified with an adaptation-based color difference such as CIE DE2000. In contrast, only chromatic nonuniformities can be quantified with chromaticity-based color differences such as delta-u'v', which are used regularly in lighting characterization. There remains a need for further development of perceptual metrics for understandable quantification of light profile uniformity in lighting which can be generally applied to the light patterns of linear as well as point light sources.

References

[1] Weinert, Jonathan. *LED Lighting Explained*. Burlington: Philips Color Kinetics, 2010. Available: http://www.usa.lighting.philips.com/connect/lightinguniversity/ledexplained/uk_en/

[2] Flynn, JE. Lighting-design decisions as interventions in human visual space. In: Nasar JL, ed. *Environmental Aesthetics: Theory, Research and Applications*. New York: Cambridge University Press, 1992: 156–170.

[3] Vogels, IMLC. Atmosphere metrics. In JH Westerink, ed. *Probing experience: From assessment of user emotions and behaviour to development of products*. Dordrecht: Springer, 2008: 25-41.

[4] Seuntiens, PJH and Vogels, IMLC, "Atmosphere creation: the relation between atmosphere and light characteristics", The 6th International Conference on Design and Emotion, (2008).

[5] Custers, PJM, de Kort, YAW, IJsselsteijn, WA, and de Kruiff, ME. Lighting in retail environments: Atmosphere perception in the real world. *Lighting Res. Technol*. 42, 2010: 331-343.

[6] Energy Star. *Program Requirements for Solid State Lighting Luminaires*. v1.1, 2008. Available: https://www.energystar.gov/ia/partners/product_specs/program_reqs/SSL_prog_req_V1.1.pdf

[7] MacAdam, David. Visual Sensitivities to Color Differences in Daylight. *J. Optical Soc. Am.* 32(5), 1942: 247–274.

Author Biography

Michael Murdoch earned his PhD in Human-Technology Interaction from Eindhoven University of Technology (2013), MS in Computer Science from Rochester Institute of Technology (2006), and BS in Chemical Engineering from Cornell University (1997). His experience includes modeling and evaluating the visual quality of imaging and lighting systems as well as creating perceptually-based visualizations of architectural lighting systems. He is a member of IS&T.