The Influence of Speed and Amplitude on Visibility and Perceived Subtlety of Dynamic Light

Michael J. Murdoch, Dragan Sekulovski, Pieter Seuntiëns; Philips Research Europe; Eindhoven, The Netherlands

Abstract

LED lighting can provide pleasant temporal dynamics, though the understanding of what constitutes pleasant has so far remained largely in the experiential domain of lighting designers. Only recently has research begun to uncover what affects preferences for dynamic lighting. Based on the insight that visible, yet subtle, dynamic lighting is desirable, an experiment was conducted which studied the influence of the speed and amplitude of periodic temporal color transitions on the thresholds for visibility and perceived subtlety. Using RGB LED lights to illuminate a wall in a living room setting, stimuli with periodic color transitions were presented to 40 observers, half of whom evaluated whether the dynamic transitions were visible, and half of whom evaluated whether the dynamics were subtle. The observers' data were fit using psychometric surfaces over the independent variables speed and amplitude of transitions. The 2D contours at the middle of the height of these surfaces were taken as thresholds. Experimental confidence was estimated via parametric bootstrapping, showing significant effects of the location and direction in color space in which the transitions were made, further strengthening the need for a temporal difference based color space. Interestingly, a distinct difference between the threshold for visibility and the threshold for subtlety can be seen, allowing for the creation of transitions that are both visibly and yet subtly dynamic. The results indicated that observers are more sensitive to hue changes than to chroma changes, and observers are more sensitive to hue changes in the orange region than to those in the blue region. Additionally, questionnaire responses provided insight into interpretations of the word subtle, understood to be pleasant, gentle, and smooth, which are also terms used to describe desired characteristics of dynamic lighting.

Introduction

New lighting technologies, notably LED, enable inexpensive ways to create lighting effects that are colorful and temporally dynamic. Solid state lighting systems can enable full control of both intensity and color, whereas traditional lighting systems typically change only the intensity of the illumination produced (for most of them only at two levels – on and off). Furthermore, the speed at which these changes can be controlled in solid state lighting is very fast, enabling the creation of fully dynamic atmospheres. A Dutch study [1] showed that people who were familiar with the Philips Living Colors, a LED based decorative lamp, liked the dynamic demo mode of the device and considered the addition of controlled dynamic light effects a valuable feature. The appeal of dynamic lighting is not surprising as light in nature is typically dynamic: sunlight changes in color temperature and intensity over the day, clouds attenuate it periodically, and wind-blown leaves provide ever-changing patterns of shadow and light. The subtle dynamics of natural light are familiar and pleasant.

In some applications, professional lighting designers tailor dynamic solutions for specific purposes or customers, for example a retail store or hotel lobby. In other applications, such as for home use, a lighting designer is unlikely to be involved, so a dynamic lighting solution must be pre-packaged. Thus, for the maker of consumer dynamic lighting solutions, much must be known about the perception of dynamic color effects and the preferences of the end user. However, there is little research in technical literature on the topic.

Recent work by Hartog et al. [2, 3] studied preferences for the speed, direction, and amplitude of periodic color changes. The quantative findings show that in general small amplitudes of change as well as slow light effects were seen as more attractive, showing preference for subtle dynamics. In an interview afterwards, 72% of the participants indicated that they want dynamic lighting in their living room. One interesting finding in her work was that people overwhelmingly like the idea of dynamic lighting, but the dynamics in most of the presented stimuli were seen as too obtrusive or too fast. Inspired by this, the main question of the experiment presented in this work was: what makes a temporal light transition appear subtle?

Experiment

The specific goal of the experiment was to understand the influence of speed and amplitude of periodic temporal transitions on the thresholds for visibility and perceived subtlety. Of special interest was whether these thresholds are distinct, which would mean that it is possible to generate visibly dynamic transitions that are subtle. Recognizing that existing color spaces (i.e. CIELAB) are not perceptually uniform for temporal color differences, two additional factors, the location and the direction of the transition in color space, were added to the experiment.

Laboratory Setup

The experiment was conducted in a laboratory arranged in a living room setting, with a sofa, chair, coffee table, and television. The dynamic lighting effects were provided by a pair of LED luminaires, one on each side of the television. The luminaires were controlled via a network interface by a Java program. Each luminaire was a tall vertical cylinder with two sets of RGB LEDs mounted on the ends of a solid cylindrical PMMA light mixing chamber. The chamber had a patterned diffuse coating on one side to couple the light out toward the wall. The result of this was a bright, oblong, soft-edged patch of light on the wall. Test participants sat on the sofa, from which vantage point they could see the two light patches on the wall, separated by 40 degrees of visual angle and each at a distance of 3m. They were asked to fixate on a marked spot at the center of the left light patch while evaluating the stimuli. The light patches themselves were very soft-edged, about 10x20 degrees in size. The room was otherwise lit by a single warm white (2700K) narrow-beam (12°) halogen spot light above the sofa, which provided 110 lux illumination at one meter height above the floor in the region of the sofa but no direct illumination on the LED-lit wall. The arrangement can be seen in the photo in **Figure 1**.

***Figure 1**: Photograph of the experimental setup, showing the LED luminaires' light patches on the left wall, the sofa where the observers sat at the right, and the laptop PC that was used for their input on the table. At the rear of the room was a small table where the experimenters were seated.*

Stimuli

The LED luminaires' RGB primaries can reproduce a very wide color gamut, and they were measured and modeled in order to accurately display colors. Spectral measurements were made of the light pattern reflected on the wall, and while the reflectance of the wall's white paint was not explicitly measured, its contribution is included in these measurements. The spectral data were converted to CIELAB LCh, using a reference white of D65 at 30cd/m^2, which was the maximum achievable luminance at this chromaticity. Assuming that the observers were adapted to this particular color is probably not correct, but determining the actual adaptation point was impossible, and given the fact that CIELAB is only approximate in this temporal application, unnecessarily precise. Note that CIELAB, which is reasonably perceptually uniform for spatial color differences, is very non-uniform for temporal color differences. However, there is no temporal equivalent, and the familiarity of CIELAB for describing colors and differences makes it useful for this work.

Based on Hartog's findings that dynamic colors varying in luminance were generally not preferred, the dynamic stimuli were made to vary in the CIELAB chroma and hue directions. For a given stimulus, two color endpoints were selected, equidistant from a base color point in the axis of interest, for example chroma, and the LEDs were made to smoothly vary periodically between the endpoints. The color transitions were substantially linear (had a fixed ΔE/sec), with exception near the endpoints where the rate of change was decelerated exponentially to avoid a visible "bounce." Details of the deceleration were taken from the work of van Beurden [4, 5], as were guidelines for ensuring smoothness by addressing the LEDs at a frequency high enough to avoid visible steps. An example of such a transition for a stimulus with changing chroma is shown in **Figure 2**.

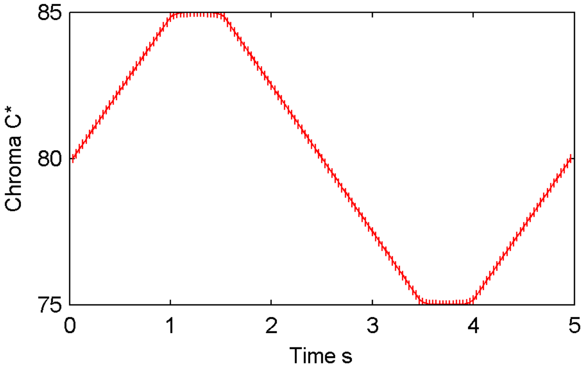

***Figure 2**: A plot of CIELAB chroma as a function of time for a dynamic light stimulus. The stimulus has a base chroma of 80, amplitude of 10 ΔE, and speed of 5 ΔE/sec in the linear portions. The discrete points show the chroma values at time steps corresponding to 30 Hz addressing, which is visibly smooth.*

Two base colors were selected, a deep blue (LCh 50, 50, 275) and an orange (LCh 60, 80, 50), and each was varied in the chroma and hue directions. The blue and orange colors were chosen because of their usefulness in atmosphere creation for living rooms, as described by Seuntiens [6]. The stimuli varied in the speed of the transitions, i.e., the slope of the linear portion, and in the amplitude, or the ΔE distance between the endpoint colors. In pilot experiments, it was seen that both speed and amplitude affect visibility and subtlety, and that their effects are correlated, but not fully dependent. For example, at very low speeds the transition remains not visibly dynamic regardless of the amplitude. Likewise, at low amplitudes, all speeds produce a transition that is not visibly dynamic. However, in intermediate speed/amplitude combinations, the threshold can be reached by varying either, i.e. they are dependent. Rough estimates of the thresholds were made by 2 observers, and these were used in the selection of parameter combinations for the full experiment.

Experiment Design and Method

The experiment used a within-subjects design including four factors: base color (blue or orange), transition direction (chroma or hue), speed, and amplitude. For each of the four base color and direction combinations, 25 periodic color transition stimuli were created, varied over a grid in speed and amplitude. The parameter values of the stimuli are shown in **Figure 3**, blue dots indicating subtlety stimuli and red dots indicating visibility. Between-subjects, two different dependent variables were measured, visibility and subtlety. Conceptually, two sub-experiments were conducted simultaneously: one looking at visibility threshold, and the other looking at subtlety threshold. Thus, in the whole experiment, there were 200 different stimuli, though each observer only saw the 100 for the dependent variable (visibility or subtlety) he or she was assigned.

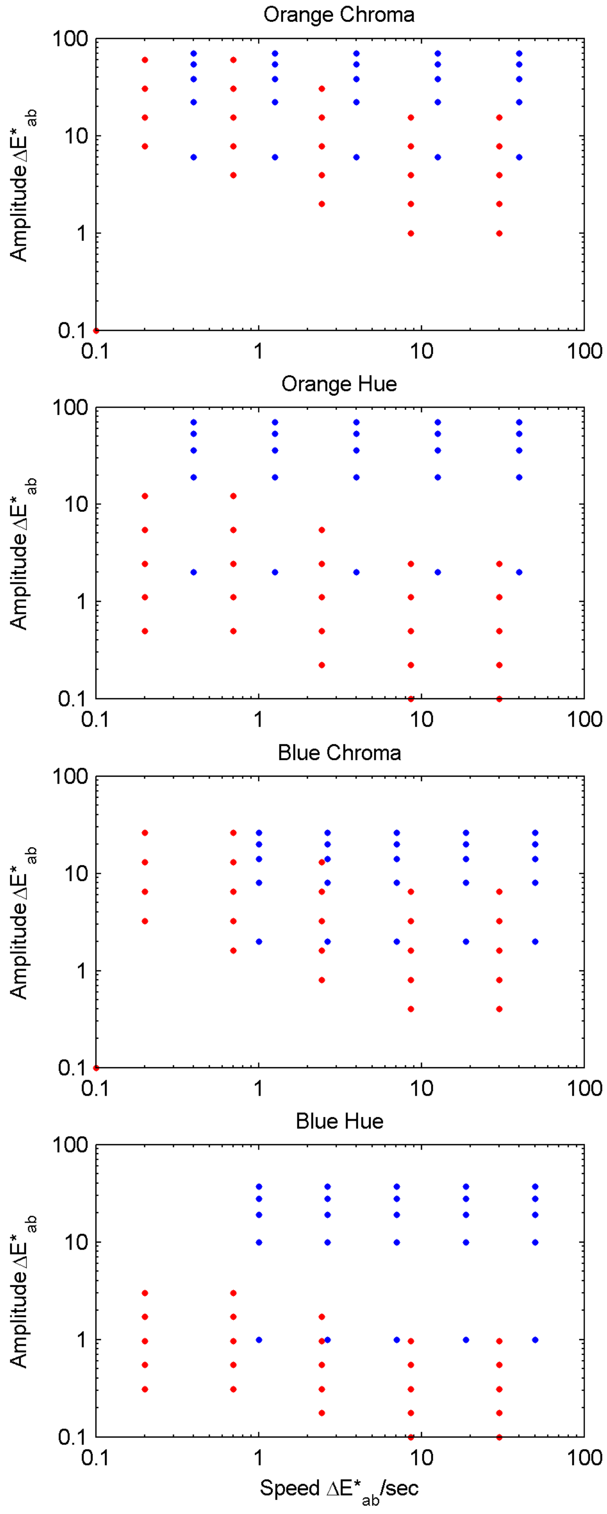

***Figure 3**: Diagrams showing the (speed, amplitude) coordinates of the 200 stimuli used in the experiment. The four plots correspond to orange chroma, orange hue, blue chroma, and blue hue variations as labeled. Within each plot, the red plots indicate stimuli used to measure visibility thresholds, and the blue points indicate stimuli used to measure subtlety thresholds. Placement of each of these grids of stimuli was chosen based on the 2 pilot observers' thresholds.*

After 4 trial stimuli showing high and low visibility (or subtlety) samples, the observer evaluated 100 stimuli, presented in random order. The observers' task was to indicate, with a keyboard press indicating yes or no, whether they perceived the stimulus as changing; or, for the subtlety response, whether they perceived the stimulus as subtle. Between stimuli, the luminaires showed static, D65 neutral light patches for 3 seconds. There was no time limit on the evaluations of the stimuli, though in the instructions observers were advised that some transitions were slow, so that waiting for up to 20-30 seconds may be necessary. In actuality, the average time to evaluate a stimulus was 20 seconds, for an average experiment duration of 33 minutes.

At the end of the observations, each participant answered a few open written questions about their experience, whether they found the task difficult, etc. The questionnaire for the subtlety participants included two additional questions about their interpretation of the word subtle.

Observers

40 observers participated in the experiment, half responding for visibility and the other half for subtlety. The visibility participants were 7 females and 13 males with an average age of 28 years. The subtlety participants were 6 females and 14 males, averaging 27 years of age. All were either students or regular employees sampled from the working population of Philips Research Eindhoven, and all were confirmed to have normal color vision by way of the Ishihara tests for color deficiency [7]. They received no additional compensation for their participation.

Analysis of Results

Psychometric Surfaces

For each of the 200 stimuli, the experiment yielded binary responses for each of the 20 observers, the mean of which was taken to be the probability that an observer in the population would find that stimulus either visible or subtle. In many experiments, such data can be used to fit a psychometric curve, yielding a threshold for the response variable on a one-dimensional scale of the independent variable. In this experiment, however, there are two independent variables, speed and amplitude, which affect the response variable together, so the psychometric curve must be generalized to a psychometric surface. This implies that the threshold is not a single point halfway up the curve, but the locus of points describing a curve in the plane halfway up the surface. Just as a 1D psychometric curve can be fit with a variety of sigmoid functions, such as a cumulative normal distribution function or a logistic function, the 2D psychometric surface can be constrained in different ways. In the domain of academic testing, where it is described as a multidimensional item response surface, Reckase [8] outlines the use of a 2D logistic function – actually the product of two orthogonal logistic functions – but this method constrains the curvature of the threshold locus to a fairly tight bend. Because in the present experiment there seems no justification for such a constraint, a bivariate cumulative normal distribution function was used to fit the psychometric surfaces. This function offers the familiar sigmoid shape in two orthogonal directions, each with a mean and variance, as well as a covariance parameter that controls the blend between the orthogonal directions and thus also the curvature of the resulting threshold locus. All five parameters were optimized to fit the probability data corresponding to the 25 stimuli in each response/base color/direction combination, resulting in 8 psychometric surfaces. An example of the fit with the resulting threshold curve is given in **Figure 4**.

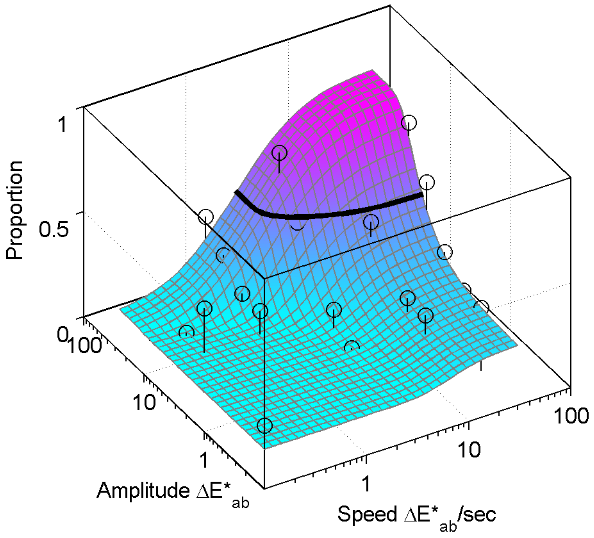

***Figure 4**: Example of one psychometric surface fit for the visibility of blue chroma transitions. Speed and amplitude are shown on the x and y axes, and the proportion of positive responses is shown on the z axis. Experimental data are shown as black circles, the bivariate cumulative normal distribution function fit to those points is shown as the surface, and the threshold halfway up the surface is shown as the thick black line.*

The minimum and maximum values were specified for each fitted surface to allow for some incorrect observer responses. The maximum was set at 0.95 in all cases, and the minimum was set to 0.05 for all of the subtlety cases, 0.3 for the orange visibility cases and 0.2 for the blue visibility cases. These last two values were used based on the actual proportions observed in a "null" case of (0.1, 0.1) speed and amplitude, which was assumed to appear static for everyone. For each psychometric surface, a threshold locus was drawn as the intersection of the surface and a horizontal plane halfway between its minimum and maximum values. Unfortunately, the surface fit for visibility of blue hue changes was unsuccessful because the 25 proportion values were all low and trendlessly noisy. None of the stimuli were reliably visible enough to yield a threshold; presumably this was due to poor choice in stimuli parameters based on incomplete piloting with too few viewers. For the other seven cases, good fits of the psychometric surfaces were obtained.

Confidence Estimates

Parametric bootstrapping is one of a group of Monte-Carlo methods used to estimate the confidence intervals of the fitted parameters. The bootstrap in general is used in cases where a known distribution cannot easily be fit, or where the final result is computed in a complex way. In this experiment, the resulting thresholds are traces in a two-dimensional space computed using a non linear fitting procedure. Simulated observer response sets were sampled 2000 times from a binomial distribution with the number of observers and the estimated probability of visibility or subtlety from the experiment as parameters. For each bootstrap sample a new set of psychometric surfaces was fit, resulting in a new set of threshold traces. These traces were parameterized, and at a series of points along their lengths, the distribution of samples was sorted, allowing 95% confidence intervals to be estimated. The result is a confidence region for each threshold, and these are shown along with the measured thresholds in **Figure 5**. Where these confidence regions do not overlap, there is high confidence that the measured thresholds are significantly different.

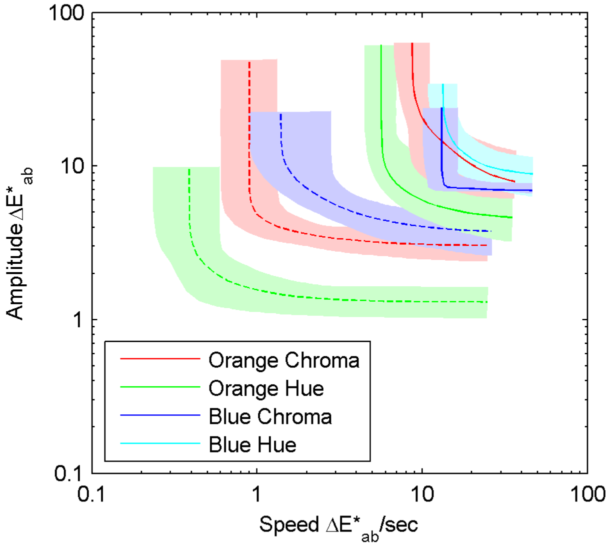

***Figure 5**: Visibility (dashed) and subtlety (solid) thresholds over a range of transition speeds and amplitudes. The shaded regions show the estimated 95% confidence intervals around each threshold. Note that the threshold for blue hue visibility is missing due to low visibility for all corresponding stimuli.*

Questionnaire Analysis

The first question answered by all participants was, "Did you find the task difficult? Why?" 19 of 20 visibility participants said it was difficult, as did half of the subtlety participants. People cited reasons like difficulty focusing on the light patch for the length of the experiment and feeling like their eyes were adapting and changing with time. The perceived difficulty may explain some of the noise measured in the responses, especially for low-visibility stimuli.

The subtlety participants were asked to provide a short description of what subtle means to them, as well as to list some synonyms for the word subtle. Answers to both questions were roughly grouped into three categories: the first (most common) included phrases like "soft and smooth," "gradual," and "change is barely noticeable," seemingly descriptive terms; the second included "pleasant and comfortable" and "gentle," perhaps describing the feelings associated with subtlety; and the third included different phrasings of "slow," explicitly describing the speed.

All participants were asked to "describe, on the basis of what you have seen, what you think is pleasant for dynamic lighting (e.g. in terms of color, variability of color, variability of speed)?"

Some of the most common responses were “slow changes,” “smooth transitions,” and “colors in the same range.” These responses overlap heavily with the descriptions of subtlety, reinforcing the idea that subtle dynamic lighting is also pleasant dynamic lighting.

Conclusions

The experiment shows that there is a clear, measurable subtlety threshold for dynamic lighting transitions which is significantly higher in speed and amplitude than the visibility threshold for chroma and hue variations. This finding provides useful boundaries for periodic color transitions that will be visibly dynamic while remaining subtle. The threshold depends on the location in color space and the direction of change: observers are more sensitive to hue changes than to chroma changes, and they are more sensitive to hue changes in the orange region than to hue changes in the blue region. This latter distinction may be due to color name boundaries, specifically the smaller color name regions around orange. The authors observed that for the same size (same ΔE amplitude) hue transition, an orange stimulus easily crossed into yellow or pink, while a blue stimulus remained blue. Judging from the confidence intervals, there appears to be no significant effect of base color on the thresholds for transitions in the chroma direction.

Observers seem to share a common understanding of the term subtle, and while preference was explicitly not measured, it is clear from the observers’ positive words, such as “pleasant” and “comfortable,” that subtle is a good characteristic for dynamic lighting. Further, the features observers suggested for pleasant dynamic lighting largely overlap with those of subtle dynamic lighting.

The limitations of the experiment show opportunities for future work. First, the use of a color space designed for spatial color differences is not optimum. CIELAB is not perceptually uniform for temporal color changes, meaning that results obtained in one area of color space would not be expected to be universal. Hence, the thresholds determined only for two base colors and two transition directions offer only a preview of the overall perception of dynamic lighting. Further, the thresholds were determined for a very critical situation, with the observer fixating on an illuminated patch in central vision. In normal activities, dynamic lighting is expected to be a background element, likely in peripheral vision, and thus the thresholds may in practice be affected by the observer’s visual attention and the differences in temporal sensitivity of peripheral versus central vision.

Acknowledgements

The authors extend thanks to Eindhoven University of Technology students Jente de Pee, Jesús Muñoz Alcántara, Monica Perusquía Hernández, and Yvonne Limpens for their careful and enthusiastic help with experimental execution.

References

[1] (Accessed 2007). Philips Living Color. *Tijdschrift voor Marketing*. Available: http://www.marketing-online.nl/topofflop/case0708.2007.html

[2] M. Hartog, "Dynamic Coloured Lighting," MS Thesis, Industrial Design Engineering, Delft University of Technology, Delft (2010).

[3] D. Sekulovski, P. Seuntiens, and M. Hartog, "Measuring Dynamic Lighting Atmospheres," Proc. AIC, Zurich (2011).

[4] M. H. P. H. van Beurden, "Smoothness and flicker perception of temporal color transitions," MS Thesis, Human Technology Interaction, Eindhoven University of Technology, Eindhoven (2007).

[5] D. Sekulovski, I. Vogels, M. van Beurden, and R. Clout, "Smoothness and Flicker Perception of Temporal Color Transitions," Proc. 15th CIC, Albuquerque (2007).

[6] P. J. H. Seuntiens and I. M. L. C. Vogels, "Atmosphere creation: the relation between atmosphere and light characterisitics," Proc. 6th Conference on Design & Emotion, Hong Kong (2008).

[7] S. Ishihara, Ishihara's Tests for Colour Deficiency, 24 Plates ed. (Kanehara & Co, Ltd., Tokyo, 1999).

[8] M. D. Reckase, "Multidimensional item response theory," in Handbook of Statistics. vol. 26, (Elsevier, Amsterdam, 2006) pp. 607-642.

Author Biography

Michael Murdoch earned his MS in Computer Science from Rochester Institute of Technology (2005) and his BS in Chemical Engineering from Cornell University (1997). He has experience modeling and evaluating system image quality of photographic and display technologies. Since 2008, he has worked at Philips Research in Eindhoven, The Netherlands, on visual perception for displays and lighting. He is a member of IS&T.