

More Than Meets the Eye: An Engineering Study to Empirically Examine the Blending of Real and Virtual Color Spaces

Joseph L. Gabbard¹

Center for Human-Computer Interaction
Virginia Tech

J. Edward Swan II³

Computer Science & Engineering
Mississippi State University

Jason Zedlitz²

Industrial Systems Engineering
Virginia Tech

Woodrow W. Winchester III⁴

Industrial Systems Engineering
Virginia Tech

ABSTRACT

It is well-documented that natural lighting conditions and real-world backgrounds affect the usability of optical see-through augmented reality (AR) displays in outdoor environments. In many cases, outdoor environmental conditions can dramatically alter users' color perception of user interface elements, by for example, washing out text or icon colors. As a result, users' semantic interpretation of interface elements can be compromised, rendering interface designs useless or counter-productive – an especially critical problem in application domains where color encoding is critical, such as military or medical visualization.

In this paper, we present our experiences designing and constructing an optical AR testbed that emulates outdoor lighting conditions and allows us to measure the combined color of real-world backgrounds and virtual colors as projected through an optical see-through display. We present a formalization of color blending in AR, which supports further research on perceived color in AR displays. We describe an engineering study where we measure the color of light that reaches an optical see-through display user's eye under systematically varied virtual and real-world conditions. Our results further quantify the effect of lighting and background color on the color of virtual graphics, and specifically quantify how virtual colors change based on different real-world backgrounds.

KEYWORDS: Outdoor Augmented Reality, Optical See-through Display, User Interface Design, Color Perception.

INDEX TERMS: H.5 [Information Interfaces and Presentation]: H.5.1: Multimedia Information Systems — Artificial, Augmented, and Virtual Realities; H.5.2: User Interfaces — Ergonomics, Evaluation / Methodology, Screen Design, Style Guides

1 INTRODUCTION

Optical see-through augmented reality (AR) head-mounted displays (oHMDs) have been used in a handful of applications domains (e.g., Livingston et al. [14], Gleue and Dähne [10], Caudell and Mizell [2]) and in numerous AR-based research studies (e.g. Renkewitz et al. [17], Gabbard et al. [8, 9]). While handheld AR has understandably received much attention recently, it can be argued that head-worn AR displays will continue to be an appropriate display platform for specific applications, such as those that require hands-free viewing. And while there are a class of head-

mounted AR displays that employ video see-through technology, our research focuses on the unique perceptual challenges associated with viewing virtual graphics through an optical-see through display. For this work, a unique characteristic of optical see-through displays is the fact that users perceive light that is an optical combination of real-world and synthetically-generated light (i.e., graphics). This is not the case in handheld AR and video-based HMDs, since these technologies generally sample the real-world scene using a video camera and subsequently present the scene via some display mechanism (e.g., LCD screen) – as opposed to allowing the natural light to pass through the display to the user's eye. It is the unique blending of natural lighting (direct and/or reflected), background, and virtual graphics that defines an interesting research nexus, since there is little understanding of how to predict or characterize the light created, and viewed by the user, when these elements converge.

While not fully understood, recent work in AR has undoubtedly shown that the blending of real-world lighting, backgrounds, and virtual graphics has an effect on display usability and user performance (Livingston et al. [13], Gabbard et al. [8, 9], Thomas et al. [21]). In many cases, outdoor environmental conditions can, for example, dramatically alter users' color perception of user interface elements. A common observation in outdoor AR using optical see-through displays is that text or icon colors become altered or washed out in sunny outdoor conditions (Thomas et al. [21]). In these cases, the semantic interpretation of interface elements can be compromised, rendering software developers' user interface design choices useless or counter-productive – an especially critical problem in certain application domains where color encoding is critical (e.g., military, medical visualization, etc.).

To begin understanding and more importantly *quantifying*, how real and synthetic visual light combine in an optical see-through display, we created a formalization of color blending in oHMDs that describes the basic components of, and interaction between, light in this research space. Further, we constructed an optical testbed that emulates outdoor lighting conditions and allows us to measure the combined color of real-world backgrounds and virtual colors as projected through an optical see-through display. Our initial goals for creating the testbed were to (1) simulate natural outdoor lighting quality and brightness in a controlled indoor environment, (2) systematically vary real-world background objects, and thus systematically vary reflected, colored, light, (3) integrate an optical see-through display driven by software capable of systematically displaying virtual stimuli, (4) determine a data collection method that is accurate, reliable, and automated as much as possible, and (5) establish a research testbed flexible enough to conduct both engineering studies (our first is reported in this paper), as well as, user-based studies (which are planned).

While this work focuses on establishing a formalization of the research space and creating a testbed and collecting data to better understand how natural and synthetic spaces interact, our long-term goal is to mathematically model *how* these spaces interact. Such a model, in turn, can be used to predict, in real-time, what

¹ jgabbard@vt.edu (corresponding email)

² jzedlitz@vt.edu

³ swan@acm.org

⁴ wwwinche@vt.edu




Figure 1. The 1931 *Commission Internationale de l'Éclairage* (CIE) color space used to characterize human-visible light. The color space includes the xy chromaticity diagram; a 2D diagram depicting hue and saturation. Also shown is the Adobe RGB color gamut [1] (triangle), as well as the specific x,y coordinate commonly referred to as the D65 white point.

color will reach an oHMD user's eye, and thus can be used to drive active AR user interfaces (Gabbard et al. [8]) to ensure color-encoding are well-preserved.

In this paper, we present our experiences designing and constructing an optical testbed to measure the result of blending natural and synthetic light through an oHMD. Section 2 presents the color spaces used for this research, while Section 3 presents related work in AR and other fields. Section 4 describes our formalization of color blending in oHMDs, and characterizes the process of perceiving a colored graphic on an oHMD against a background of another color. Section 5 describes the optical testbed, including some tradeoffs we made during its design. Section 6 details an engineering study, conducted using the testbed, where we measured the color of light produced by an oHMD under systematically varied virtual and real-world conditions. We present our results and conclusions in Section 7, followed by ideas for Future Work in Section 8.

2 BACKGROUND

This work uses an oHMD driven by a typical computer graphics card. The display employs an extended color gamut indexed by the RGB color model (Foley et al. [4]); the most common color model in computer graphics. The RGB color model is an additive model, describing a color as a combination of red (R), green (G), and blue (B) primitives, which are added to black, rather than subtracted from white. Our oHMD uses the Adobe RGB [1] gamut, which is larger than the standard RGB gamut commonly used in computer displays; Adobe RGB includes more visible green and cyan colors.

To examine our real-world background colors, as well as the resulting blended colors, we employed the components of the *Commission Internationale de l'Éclairage* (CIE) color model. The 1931 CIE color model includes a CIE xyY color space that separates hue information (x and y) from luminance information (Y). The 1931 CIE model also defines a CIE xy chromaticity diagram

to graphically depict the relationship between hue (light wavelengths) and purity allowing for the specification of individual colors within an x,y grid system (see Figure 1).

The CIE chromaticity diagram exhibits a number of interesting properties, including the fact that when a line is drawn between two color points, the line passes through all possible colors created by mixing the original two color points in varying proportions (Foley et al. [4]). This mixing property can be extended to three points as well: any three color points in the diagram define a triangle that encompasses the colors produced by mixing varying proportions of the three colors (Wyszecki & Stiles [22]). Manufacturers of color monitors, televisions, color printers, and most other electronic color devices have their own specific three color points which define the complete color gamut of each device within CIE space. Similarly, the RGB color gamut can be defined within the CIE xy chromaticity diagram (Figure 1), as can the set of naturally occurring colors that fall outside of the gamut.

While the 1931 CIE model exhibits good computational properties, researchers were concerned that the xyY space was perceptually non-uniform. In 1976 the CIE model was revamped to include a 2D space defined by u' and v' , specifically designed to allow for perceptually uniform analysis of hue differences. Note that, like the 1931 xy chromaticity diagram, the $u'v'$ chromaticity diagram (also termed the CIE 1976 uniform chromaticity scale) is designed to illustrate hue difference and does not account for differences in lightness. The 1976 CIE model also includes the definition of $L^*u^*v^*$ color space, which is an approximately perceptually uniform, 3D computational color space that accounts for both chromatic (u^*v^*) and lightness (L^*) differences. We chose to use the 1976 CIE color space for our analysis so that we could calculate perceptually uniform differences in color.

The CIE color model also allows for flexible definition of “white” to account for the fact that colored surfaces can appear very different under varying lighting sources. For example, incandescent lighting produces warm, red-yellow light, while an overcast day produces cool, blue light, and this change can make the same surface appear to be a different color. This change in light color is described by *color temperature*, and is determined by comparing the light's color to the color of a known material as the material is heated. The white point is associated with a specific x,y coordinate on the 1931 CIE xy chromaticity diagram, and denotes “white” under a given color temperature. Standard white points have been developed to represent common color temperatures under standard outdoor lighting conditions, such as D50, D55, D65, and D75 (Judd et al. [12]). The most common of these color temperatures is D65, representing the white point of the color temperature 6500 Kelvin. It is often referred to as the daylight illuminant characterized by average outdoor lighting conditions.

Many other color models have been proposed (e.g., HLS, HSV, etc.) for various uses, including those that address perpetual issues and those to describe colors in terms more natural to an artist. Two examples of such models are the hue, saturation, and value (HSV) color model created by Smith [19] (sometimes referred to as HSB model, where “B” stands for brightness), and the hue, lightness, and saturation (HLS) color model proposed by Ostwald [16].

3 RELATED WORK

Chevreul [3] defines the color appearance phenomenon termed *simultaneous contrast*, where the perceived color of an area (e.g., a small colored square) can be affected by a surrounding, contrasting color. The effect is more noticeable when two examples are side by side, with identical center colors but complementary surrounding background colors (Rosotti [18]). Similar studies have produced similar results with computer monitors (e.g., Luo et al.




Figure 2. A formalization of color blending in oHMDs. Empirical studies have shown source light (L_1), background (B), and AR display light (L_3) have a direct effect on the blended light that enters a user’s eye (L_4), and thus on the perceived color (C).

[15]). While these studies do not explicitly examine the effect of blending two colors (e.g., one color on top of another), they do provide clear evidence that backgrounds (i.e., the surrounding color) have an effect on the target color (e.g., a colored square in the center of the surrounding color).

Much work has been done to develop mathematical models and equations that describe how two colors in a given color model combine to create a third. However, most of this work is conducted in a single color space, or with a common white point, and do not address, for example, how light from two qualitatively different sources (e.g., nature versus technology), much less two different color spaces and white points, can be combined and how users will perceive these colors. For example, Foley et al. [4] describe how to work with refractive and non-refractive transparency in computer graphics, by providing calculations that define how two colors, both defined in RGB space with associated transparency values, can be combined to accurately determine the resulting color. To our knowledge, this paper is one of the first published attempts to understand how color from a natural light source defined in the CIE color space interacts with the color emitted from an oHMD source defined in the RGB color space. As mentioned above, only a subset of CIE colors can be represented in RGB, it is important to note that our work is (and specifically, this initial testbed and procedure are) designed to examine CIE colors that fall both outside and inside the RGB color gamut and designed to handle conditions where the white point of each color is arbitrary.

Very little work has been done to quantify the effects of real-world background and lighting on virtual graphics. Livingston et al. [13] present a set of studies that examine the combined effect of resolution and display contrast on visual acuity, using a number of commercial AR displays. Results indicate that even modest adjustments to levels of contrast can help users achieve greater acuity. Livingston et al. [13] also describe a color matching study that indicates AR displays often distort blue regions of the color space, and that AR users perceptually distort colors to a greater extent (e.g., the color swatch a user chooses to match an AR color is “far away” from the true target color).

Renkewitz et al. [17] performed an empirical study in an outdoor AR setting to examine the effect of background textures and text color on text recognition, and to derive text size recommendations for AR designers. This study showed no effect of background and text color on legibility using word recognition tasks. However, this work was performed with a video see-through display system, where virtual text pixels effectively replace background video-captured pixels, and as such, did not examine the optical blending of background and virtual colors.

Gabbard, Swan, and Hix [9] presented a user-based outdoor AR study that examined the effect of text color and backgrounds on a text legibility task. In this study, posters were used as proxies for actual real-world backgrounds. The study was extended by Gabbard, Swan, et al. [8] to include additional independent variables and actual real-world backgrounds. Both studies used optical see-through displays and both studies found very strong evidence that lighting and real-world backgrounds affect text legibility. The independent variables used in the engineering study presented herein were directly motivated by these studies.

4 FORMALIZING COLOR BLENDING IN AR

Figure 2 characterizes the process of perceiving a computer-generated colored graphic (e.g., the letter ‘T’) on an oHMD against a background of another color. This forms a pipeline that can be specified mathematically.

In Figure 2, a light source produces light with the spectrum L_1 . The light L_1 hits a surface with background reflectance B , producing reflected light

$$L_2 = RF(L_1, B),$$

where the function RF represents how surfaces interact with reflected light. RF comes from surface reflectance physics (e.g., Stone [20], Wyszecki & Stiles [22]). In this work, we are only concerned with RF to the degree that we choose light sources L_1 and backgrounds B that are likely to be encountered in common AR applications.

The reflected light L_2 now enters the front of the AR display, where it combines with the AR display light L_3 to produce

$$L_4 = AR_D(L_3, L_2).$$

Here, L_4 is the blended light that reaches the user’s eye. AR_D is the function that characterizes the entire display system, consisting of the oHMD components, graphics card, cable that connects the display to the graphics card, physical oHMD settings such as opacity, brightness, or contrast that can be changed on the display, and so forth. We consider that the general function AR is parameterized by a particular display system; we indicate this parameterization by the subscript D in the function AR_D .

This work supports our longitudinal goal to characterize AR_D , at least for the specific display systems D located at our laboratories. AR_D can potentially be modeled with an interpolated 3D lookup table, similar to the way color management systems model gamut mappings (e.g., Stone [20]). A model of AR_D could therefore drive an active AR system.

The final step in the pipeline is when the blended light L_4 enters the user’s eye, causing them to perceive the color

$$C = HP(L_4).$$

Here C is the user’s sensation of the perceived color, and HP represents the operation of the human perceptual system.

We can compose these functions to produce the entire pipeline in the equation

$$C = HP(AR_D(L_3, RF(L_1, B))).$$

This shows that the inputs to the perceived color C are the AR display light L_3 , the light source L_1 , and the background reflectance B . Indeed, L_3 , L_1 , and B are the primary independent variables that we systematically manipulated in the engineering study described herein.

5 EXPERIMENTAL TESTBED

We designed and constructed an experimental testbed on an optical bench that emulates outdoor lighting conditions as seen through an oHMD (Figure 3). This testbed was created to evaluate the resulting combined color of a real world background and a virtual color projected through the oHMD.


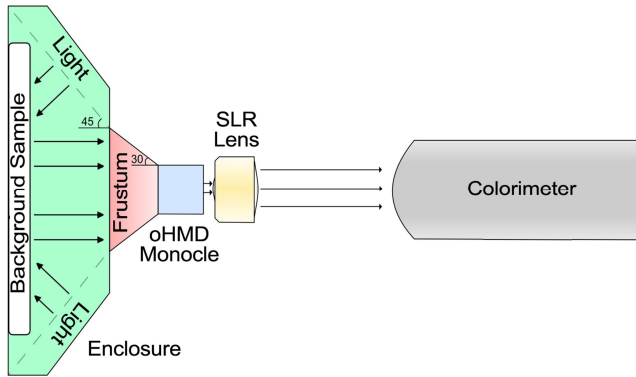


Figure 3. (Above) A top view schematic of our experimental testbed depicting an enclosure (green) containing two 15-watt lights reflecting off of real-world background samples. Indirectly reflected light exited the enclosure through the frustum (red) and into the AR oHMD monocle (blue). The AR oHMD image, designed to be focused on a user’s retina, was directed towards the colorimeter (grey) using an SLR camera lens (yellow). A desktop computer (not shown) was used to drive the AR display, as well as, to capture/log data from the colorimeter. (Below) An annotated photograph of our optical experimental testbed.

To construct the testbed we used specialized lights that reproduce the daylight standard white point of D65, and accurately project as much of the visible color spectrum as naturally exists outdoors. To assess how much of the natural visible spectrum is produced, light sources are rated using the color rendering index (CRI), which ranges from 1–100, with 100 denoting a light that can perfectly reproduce the entire daylight color spectrum. For this study, we used fluorescent lights rated at 95 CRI to simulate the D65 outdoor lighting conditions.

We sealed the lights in an enclosure constructed of black foam core board. We designed the enclosure with a hinged top to allow us to slip backgrounds in and out. To avoid direct reflections and lighting “hot spots”, we offset the lights from the background by 45 degrees, so that light did not fall directly onto the sample and reflect directly back out the aperture. To ensure we had outdoor-equivalent brightness levels at the aperture, we first established a brightness goal of 800-1000 lux, determined by measuring the amount of light reflecting off a white piece of paper while standing in the shade. This target is consistent with an overcast day, or “very dull” outdoor conditions as defined by Halsted [11]. We observed readings close to 900 lux as we measured the amount of




Figure 4. We color-matched paint to real objects to create real-world samples for use in our testbed enclosure. Specifically, we color-matched samples of brick, foliage, pavement and sidewalk. We measured both the painted sample and the actual background using a colorimeter to ensure an adequate match.

light projected through the aperture when the same white sample was illuminated inside the enclosure.

We integrated an NVIS nVisor SX oHMD into the testbed and created a small frustum to couple the oHMD monocle to the enclosure. We used a 60 degree field of view (per manufacture’s specification) to allow as much light from the sample to fall onto the front end of the monocle, and as such, illuminate and blend with the entire virtual scene.


Once we were confident that colored light was correctly entering the front of the oHMD monocle, we turned our attention to the rear of the monocle – where real and virtual blended light are typically projected onto a user’s retina, positioned directly inline with the monocle. The oHMD monocle is designed to be 23 millimeters from the eye. As expected, at this location, the resulting display projection created a footprint significantly smaller than what is required by the colorimeter sensor. Moreover, headband and other form factors prevented us from positioning the colorimeter directly in line behind the oHMD monocle. As such, we needed to address two testbed design problems: projection size and projection angle.

Initially a small piece of diffusing material was positioned at an angle behind the monocle to create a (rear) projection surface for the colorimeter to sample. A c-mount lens was used to magnify the projected light to an effect size for colorimeter measurement. This design was later discarded since the diffusing material slightly altered the color, and reduced the brightness of the light reaching the colorimeter. Instead, we used no diffusing material, and placed a 50 mm camera SLR lens at a location approximate to a user’s eye when wearing the display. The lens was pointed slightly up in the vertical plane to allow the light to be viewable above the headband of the oHMD (Figure 3).

Once assembled, we developed a calibration process to ensure that data logged from the colorimeter accurately represented what a user would experience looking through the display into the enclosure’s frustum. We used the colorimeter’s software to confirm that both our lights and display accurately produce colors using a D65 white point (even though the oHMD specification states D65). Specifically, we took colorimeter white point measurements for three experimental conditions; oHMD projecting all white graphics with enclosure lights off; oHMD off with enclosure lights on; and oHMD projecting all white graphics with enclosure lights on. In all three cases our readings were within .01 units of the true D65 point in both the x and y dimensions.

6 ENGINEERING STUDY

The engineering study aimed to measure the effect of blending natural light reflected off of real-world backgrounds with virtual light produced by an AR oHMD. We used six backgrounds and 27 virtual colors presented via the oHMD. We measured light exiting the oHMD monocle under the three experimental condi-



| | Brick | Foliage | Sidewalk | Pavement |
|-------------|-------|---------|----------|----------|
| Actual u' | 0.275 | 0.123 | 0.211 | 0.203 |
| Actual v' | 0.485 | 0.531 | 0.489 | 0.470 |
| Paint u' | 0.257 | 0.134 | 0.211 | 0.196 |
| Paint v' | 0.494 | 0.510 | 0.489 | 0.468 |

Figure 5. Using a colorimeter, we measured the actual and painted sample color for each of our four real-world backgrounds. We show the values in the table, and plotted in the CIE 1976 uniform chromaticity scale, for both actual background object (white marker) and the paint sample (black marker).

tions described in the white point calibration process. Using this approach, we examined conditions when only the backgrounds were illuminated, only the oHMD was illuminated, and when both the background and the oHMD were illuminated.

Four of the six backgrounds represent likely real-world backgrounds, and were chosen based on our previous studies (Gabbard, Swan, et al. [6, 7]): brick, foliage, pavement, and sidewalk (Figure 4). A sample of each of these objects was taken to a paint store and visually matched. The real-world objects were textured and/or too large to be scanned by paint store equipment designed to color match (a specialized colorimeter). As such, for each background, a starting paint color was first chosen by visually matching the real-world object to a specific color swatch. The paint was then manually mixed, adding small bits of various color until the resulting paint color perceptually matched the real-world object (i.e., the object and paint looked the same). We created paper samples of each background by applying the paper to various media. We first tried regular laser printer paper, but found it to be too thin and glossy, which lead to fluctuations in colorimeter readings. We then applied paint to white poster board, but found that the slick surface did not absorb enough paint to create a consistent color across the sample. Finally, we used a slightly textured, cold press watercolor paper, which provided a matte finish, and enabled us to have consistent light and color reflections across all background conditions, with no specular highlights.

To quantify how closely our paint samples matched the real-world objects, we took each object and sample outdoor and measured their color using a colorimeter. The real-world objects were too big to fit into the enclosure, so we measured their color outdoors in the shade where natural lighting levels closely matched the capability of our testbed. The colorimeter measures color in CIE xyY , however we converted these values to CIE $L^*u^*v^*$ color

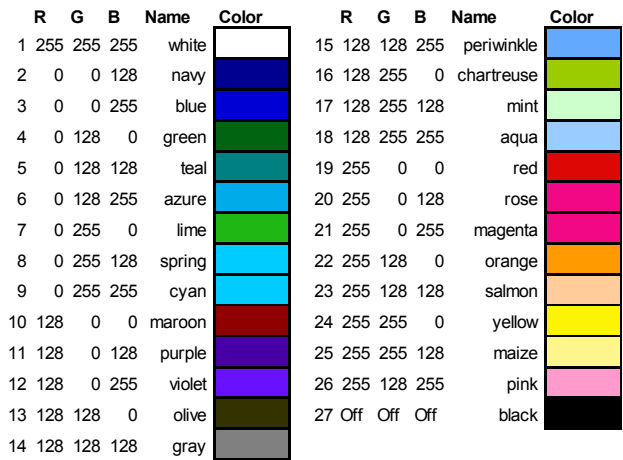


Figure 6. The 27 tested colors, defined by every combination of red (R), green (G), and blue (B), at the levels of 0, 128, and 255. For the color black, the oHMD was turned off. Note that the colors shown here, as well as elsewhere in this paper, are just representative of the real-world experience of seeing that color.

space in order to make perceptually uniform comparisons between data points. Mapping these points onto the CIE 1976 uniform chromaticity scale (Figure 5) indicates that the paint samples closely matched the actual samples.

We used two additional backgrounds: a *no-background* condition where the lights were turned off, and a *white* condition where we placed a white sheet of paper in our apparatus. There were thus six different backgrounds in the study: *brick, foliage, sidewalk, pavement, white, and no-background*.

Figure 6 shows the colors that we projected in the oHMD against these backgrounds. We chose 27 combinations of fully saturated (256), half saturated (128) and desaturated (0) values for red, green, and blue. For the color black (0,0,0) we turned the oHMD off, so the black color condition measured the background color as seen through the oHMD’s display optics.

6.1 Collecting and Processing the Data

We created a PowerPoint presentation with each of the projected colors to flood the screen of the oHMD. On average, our colorimeter collected approximately 3 readings per second, although the rate varied according to the color that was being measured. Our rig displayed each color for two minutes, and we continuously recorded the data during that time. Our colorimeter reported colors in CIE 1936 xyY format. Due to the variable amount of time per reading, we collected between 9416 and 10,294 readings per background, for a total of $N = 59,328$ data points for the six backgrounds.

We next transformed the data into the CIE 1976 u^*v^* and $L^*u^*v^*$ color spaces, using the formulas found in Foley et al. [5]. There is a direct linear relationship between xyY and u^*v^* , but converting from xyY to $L^*u^*v^*$ requires that the xyY colors be normalized relative to a *white point*. For our experimental setup, this white point represents the brightest possible color (the color with the largest energy density). Our experimental setup yielded three different white points: (1) when only the oHMD was illuminated (the no-background condition), and the oHMD was presenting the color white (255,255,255), (2) when the oHMD was off (the color black) against the white background, and (3) when the oHMD presented the color white against the white background. Our collected data contained xyY values for each of these conditions; we used these values for our three white points. We used white point (1) when only the oHMD was illuminated, white point




Figure 7. The colored bars are the mean color distances, averaged over all colors, between all pairs of background conditions, calculated in CIE 1976 $L^*u^*v^*$ space. The grey bars are the proportion of each distance that comes from changes in lightness (L^*). The different colors group qualitatively similar background pairs.

(2) when only the background was illuminated, and white point (3) when both were illuminated; we only used these white points to convert xyY values into $L^*u^*v^*$ values.

Our next goal was to reduce the data points to one repetition per background by color condition, giving $N = 161$ data points⁵. Before calculating means, we examined the noise in the data. We did this by calculating the metric

$$x_{M\%} = \text{MAX}(\text{SEM}(x)) / \text{MAX}(\text{MEAN}(x)) * 100\%,$$

where $\text{MEAN}(x)$ is the average of the repetitions within each background by color cell, $\text{SEM}(x)$ is the standard error of the mean, and $\text{MAX}()$ finds the maximum over all $N = 161$ background by color combinations. This metric reports the largest possible noise value over all experimental cells, and normalizes the noise value to take into account the varying magnitude of the different values.

We found very little noise in the measured xyY values: $x_{M\%} = 0.0156\%$, $y_{M\%} = 0.0155\%$, and $Y_{M\%} = 0.0160\%$. We examined the same metric for the color values transformed to the CIE 1976 u^*v^* and $L^*u^*v^*$ color spaces, and again found very little noise: $u^*_{M\%} = 0.0274\%$, $v^*_{M\%} = 0.0124\%$, $L^*_{M\%} = 0.0100\%$, $u^*_{M\%} = 0.0225\%$, and $v^*_{M\%} = 0.0199\%$. Based on this analysis, we calculated the mean xyY , u^*v^* , $L^*u^*v^*$ color values for the $N = 161$ background by color combinations, and used these mean values for the rest of the analysis.

7 RESULTS & CONCLUSIONS

Because CIE 1976 $L^*u^*v^*$ space is an approximately perceptually linear 3D space, we calculated 3D Euclidean norm distances in this space. Our motivation was to examine how changing from one background to another affected the perceptual measurement of each displayed color. For each of our 27 colors, we calculated the 3D Euclidean distance for each pair of backgrounds (for backgrounds paired with the *no-background* condition we calculated 26 distances). Figure 7 shows the mean distances (with standard error bars) for each of the resulting 15 background pairs; the

⁵ There were 5 (background: brick, foliage, pavement, sidewalk, white) \times 27 (color) + 1 (no-background) \times 26 (all colors except black) = 161 experimental cells; the *no-background*, *black* condition yields no light energy and was not included.

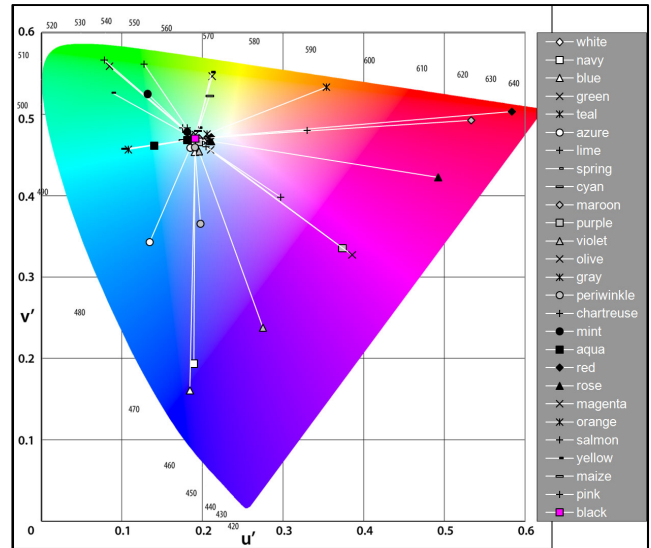


Figure 8. Chromatic changes between the *no-background* condition (outer values) and the *white* background (inner values). The pink square represents the white background color as seen through the oHMD optics.




Figure 9. Chromatic changes between the *foliage* background (outer values) and the *white* background (inner values). The pink squares represent the background colors as seen through the oHMD optics.

y -axis in this graph is expressed in 3D $L^*u^*v^*$ units. The colored bars show the total distance in $L^*u^*v^*$ space. The grey bars show the proportion of this distance that is due to lightness (the length of L^* projected onto the vector $L^*u^*v^*$).

The blue bars in Figure 7 shows that the largest perceptual distances occurred when the *no-background* condition was paired with each of the other background conditions. The grey bars indicate that only a small proportion of these distances were due to lightness changes; the percentages of the blue bars due to lightness changes ranged from 5% (*no-background*, *sidewalk*) to 13% (*no-background*, *pavement*). Instead, most of the perceptual distance is due to chromatic changes. Figure 8 examines the effect of these changes for the *no-background*, *white* condition; this figure plots the u^*v^* coordinates for each color against the CIE




Figure 10. Chromatic changes between the *brick* background (right-hand values) and the *foliage* background (left-hand values). The pink squares represent the background colors as seen through the oHMD optics.




Figure 11. Chromatic changes between the *brick* background (outer values) and the *sidewalk* background (inner values). The pink squares represent the background colors as seen through the oHMD optics.

1976 chromaticity diagram, with a line connecting colors seen against the no-background condition (outer values) to the same color seen against the white background condition (inner values). For example, the rightmost data point (a black diamond (◆)) represents the color red in the no-background condition; under the white background this color has shifted leftwards and is part of the central grouping of colors. This indicates that the color red is quite vibrant and saturated in the no-background condition, but becomes desaturated against white, and may even no longer appear to be red (verifying this last conjecture would require a human color judgment). The other colors are affected in a similar way. The pink square in Figure 8 is the color of the white background seen through the oHMD optics (the oHMD is turned off

for the color black); note that all of the inner colors are clustered around this white background hue.

This analysis raises two potential concerns: (1) will the annotated graphics themselves become less legible against the bright white background? And, (2) will the hues in the center become less perceptually discriminable, meaning that fewer hues are available as a UI design component? Figure 8 suggests that both effects may occur, but human judgments are required to validate both concerns. The other backgrounds paired with the no-background condition (the other blue bars in Figure 7) show a similar star pattern, although the central grouping is not as tight. These suggests that the oHMD colors seen against these different backgrounds will be less saturated, but will still be perceived as their original hues, and hence these usability concerns likely do not apply to the same degree (although human judgments are again required to validate this conjecture).

The red bars in Figure 7 show that the next largest perceptual distances occurred when each background was paired with the white background. Here the grey bars indicate that the major proportion of these distances were due to lightness changes; the percentages of the red bars due to lightness changes ranged from 66% (brick, white) to 75% (pavement, white). Figure 9 shows a representative result from this series; it examines the *foliage*, *white* background pair. The outer color values are the colors against the foliage background. Note that this background is brighter than the no-background condition in Figure 8, and hence the colors are shifted towards the central white point, although the different hues should still be discriminable (how discriminable would require human color judgments). The inner color values are against the white background; these are the same as the inner colors in Figure 8. The pink square is the color of the background itself as seen through the oHMD optics; note that the background changes from a greenish foliage color to white. The same usability concerns expressed above exist for all of the backgrounds paired with the white background.

The green bars in Figure 7 show relatively small perceptual distances. Here the grey bars indicate that almost none of these distances were due to lightness changes; the percentages due to lightness ranged from 0.1% (brick, pavement) to 4% (foliage, pavement). In these cases the lightness of the backgrounds remained constant, and the oHMD colors were primarily affected by the shift in background color. Figure 10 shows a representative result from this series; it examines the *brick*, *foliage* background pair. The right-hand values are colors seen against the brick background, while the left-hand values are seen against the foliage background. The pink squares show the reddish brick background hue and the greenish foliage background hue. Unlike the star shapes seen earlier, this figure shows a linear hue shift, in this case a shift between the red and green lobes of the chromaticity diagram. The primary usability concern from such a linear shift would be situations where hues change; for example in Figure 10 oranges shift into yellows, yellows shift into greens, and purples shift into blues. However, Figure 10 suggests that there are a variety of hues available for UI components against both backgrounds. Other $u'v'$ chromaticity diagrams from the green bars in Figure 7 are similar to Figure 10.

The orange bars in Figure 7 show perceptual distances between the brick, foliage, and pavement backgrounds and the sidewalk background. Although the orange bars are similar in overall magnitude to the green bars, the sidewalk is a relatively bright background and hence about half of the perceptual distances were due to lightness changes; these percentages ranged from 39% (brick, sidewalk) to 46% (pavement, sidewalk). Figure 11 shows a representative result from this series; it examines the *brick*, *sidewalk* background pair. Here the outer values are colors seen against the brick background, while the inner values are seen against the si-

dewalk background. The pink squares show the reddish brick background hue shifting to the tan-colored sidewalk background. This figure shows a combination of the star shape seen in Figures 8 and 9 and the linear shift seen in Figure 10; here the colors from the red, purple, and blue lobes show a larger degree of hue shifting than those from the cyan to green lobe. Therefore, we could expect a combination of the usability concerns expressed above. The other u^*v^* chromaticity diagrams from the orange bars in Figure 7 are similar to Figure 11.

8 FUTURE WORK

We intend to perform additional analysis on the data collected and reported on herein, to for example, carefully explore the contribution of, and interactions between, each of the $L^*u^*v^*$ components. From this analysis we hope to better understand the perceptual underpinnings of virtual colors presented against real world backgrounds. We also plan to conduct a user-based perceptual color-matching study using these same independent variables to compare how users' perception matches this data.

We also intend to expand the functionality of the testbed to, for example, allow us to experiment with actual real-world objects instead of painted surfaces. Moreover, we are considering a mobile version of the testbed to use outdoors in truly natural outdoor lighting conditions. In addition, we are considering a follow-on study to include varied or dynamic white points to better understand how colors perception is effected when users move around in real-world outdoor environments. Once we understand how colors shift under these different conditions, we plan to use this knowledge to construct adaptive algorithms to account for these shifts.

Based on the results of these studies (and others), we hope to create a predictive model from the formalization of color blending presented herein. Subsequently, we hope to apply this predictive model to a real-time active system that compensates for changes in viewing conditions to both optimize legibility as well as preserve color encoding.

REFERENCES

- [1] ADOBE, "Adobe RGB (1998) Color Image Encoding", Accessed: December 2009, <<http://www.adobe.com/digitalimag/pdfs/Adobe-rgb1998.pdf>>.
- [2] CAUDELL TP and MIZELL DW, "Augmented Reality: An Application of Heads-Up Display Technology to Manual Manufacturing Processes", Proc. 25th Hawaii International Conference on System Sciences (HICSS), 1992, Vol. 652, pp. 659–669.
- [3] CHEVREUL ME, *The Principles of Harmony & Contrast of Colors*, Faber Birren, 1839.
- [4] FOLEY J, PHILLIPS R, HUGHES J, VAN DAM A, and FEINER S, *Introduction to Computer Graphics*, Addison-Wesley Longman Publishing, 1994.
- [5] FOLEY JD, VAN DAM A, FEINER SK, and HUGHES JF, *Computer Graphics: Principles and Practice in C*, 2nd ed., Addison-Wesley Professional, 1995.
- [6] GABBARD JL and SWAN II JE, "Usability Engineering for Augmented Reality: Employing User-Based Studies to Inform Design", *IEEE Transactions on Visualization and Computer Graphics*, 14(3), 2008, pp. 513–525.
- [7] GABBARD JL, SWAN II JE, HIX D, LUCAS J, and GUPTA D, "An Empirical User-based Study of Text Drawing Styles and Outdoor Background Textures for Augmented Reality", *IEEE Virtual Reality 2005*, Bonn, Germany, 2005, pp. 11–18.
- [8] GABBARD JL, SWAN II JE, HIX D, SI-JUNG K, and FITCH G, "Active Text Drawing Styles for Outdoor Augmented Reality: A User-Based Study and Design Implications", Proc. *IEEE Virtual Reality 2007*, Charlotte, NC, USA, 2007, pp. 35–42.
- [9] GABBARD JL, SWAN II JE, and HIX D, "The Effects of Text Drawing Styles, Background Textures, and Natural Lighting on Text Legibility in Outdoor Augmented Reality", *Presence: Teleoperators & Virtual Environments*, 15(1), 2006, pp. 16–32.
- [10] GLEUE T and DÄHNE P, "Design and Implementation of a Mobile Device for Outdoor Augmented Reality in the Archeoguide Project", *Conference on Virtual Reality, Archeology, and Cultural Heritage*, Glyfada, Greece, 2001, pp. 161–168.
- [11] HALSTED C, "Brightness, Luminance, and Confusion", *Society for Information Display*, March, 1993.
- [12] JUDD DB, MACADAM DL, WYSZECKI GT, BUDDE HW, CONDIT HR, HENDERSON ST, and SIMONDS JL, "Spectral Distribution of Typical Daylight as a Function of Correlated Color Temperature", *Journal of the Optical Society of America*, 54(8), 1964, pp. 1031–1036.
- [13] LIVINGSTON MA, BARROW JH, and SIBLEY CM, "Quantification of Contrast Sensitivity and Color Perception using Head-worn Augmented Reality Displays", *IEEE Virtual Reality 2009*, Lafayette, LA, USA, 2009, pp. 115–122.
- [14] LIVINGSTON MA, ROSENBLUM LJ, JULIER SJ, BROWN D, BAILLOT Y, SWAN II JE, GABBARD JL, and HIX D, "An Augmented Reality System for Military Operations in Urban Terrain", Proc. *Interservice / Industry Training, Simulation, & Education Conference (IIITSEC '02)*, Orlando, FL, USA, 2002.
- [15] LUO MR, GAO XW, and SCRIVENER SAR, "Quantifying Colour Appearance. Part V. Simultaneous Contrast", *Color Research & Applications*, 20(1), 1995, pp. 18–28.
- [16] OSTWALD W, *Colour Science*, Winsor & Winsor, London, UK, 1931.
- [17] RENKEWITZ H, KINDER V, BRANDT M, and ALEXANDER T, "Optimal Font Size for Head-Mounted-Displays in Outdoor Applications", Proc. *IEEE Information Visualization (InfoVis 2008)*, Columbus, OH, USA, 2008, pp. 503–508.
- [18] ROSOTTI H, *Why the World Isn't Grey*, Princeton University Press, Princeton, NJ, 1985.
- [19] SMITH AR, "Color Gamut Transform Pairs", Proc. *Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1978, pp. 12–19.
- [20] STONE MC, *A Field Guide to Digital Color*, A K Peters, Natick, Massachusetts, 2003.
- [21] THOMAS B, CLOSE B, DONOGHUE J, SQUIRES J, BONDI PD, and PIEKARSKI W, "First Person Indoor/Outdoor Augmented Reality Application: ARQuake", *Personal and Ubiquitous Computing*, 6(1), 2002, pp. 75–86.
- [22] WYSZECKI G and STILES WS, *Color Science*, 2nd ed., John Wiley & Sons, New York, NY, 1982.