Testing Effects of Color Constancy for Images displayed on CRT Devices

Holger Winnemöller, M.Sc.
University of Cape Town

CS03-03-00
September 30, 2003

Department of Computer Science
University of Cape Town
Private Bag, RONDEBOSCH
7701 South Africa
e-mail: hwinnemo@cs.uct.ac.za

Abstract

In this report we examine the effects of color constancty for images displayed on computer monitors, particularly CRT devices. The motivation for this work is based upon Edwin Land's *Retinex Theory* which attempts to explain color constancy in human perception. Using Land's theory, we can create images with a severely undersampled color space (down to a full omission of one color component) that are still appropriately interpreted by the human observer. Possible applications of this are compression schemes as well as image enhancement techniques, some of which are already known. We present a simple web-based testing scheme to verify the viability of displaying duo-chrome (e.g. extreme color undersampling) images on CRT devices in different environmental settings. Some results and possible interpretations of these results are also presented.

1 Introduction

1.1 Background

The phenomenon that humans (and many other animals) can separate the properties of a global illuminant from those of the illuminated material is called *Color Constancy*. This is important as it allows people to perceive the *true color* of a material instead of the reflected color of such. For example, in a room with red lighting, a white piece of paper will still be perceived as white, even though objectively it reflects light of a red wavelength. This is due to the fact that everything else in the room also has a red tinge and we therefore deduce the existence of a red illuminant. This global illuminant is thus subtracted from the reflections received off objects in the scene to reveal their assumed, true color. It should be noted that this is a subconscious process and needs no conscious interaction.

Another phenomenon (sometimes called *color illusion*) can be described as follows: If a small grey square (the shape is not important here) is placed upon a large green square, then it will appear tinted lightly red. Similarly, if a grey square is placed upon a large red square, it will appear tinted slightly green, i.e. exactly the opposite color of the square it is placed upon. This effect also works with cyan/yellow combinations and poses problems for theories based on the belief that the cones in the eye's retina are independently sensitive to red, blue and green wavelengths. In order to explain these color illusions, the cones' responses cannot be interpreted independently. Theories exist, which are based on antagonistic interactions between combinations of cones resulting in spectrally opposing stimulation [4], [5]. Some physiological evidence exist to support such theories.

Edwin Land [6] devised an experiment using both phenomena to suggest subjective colors, which are objectively not present. In this experiment, he used a picture slide together with a few color-filters to produce duo-*chrome* images (for an example using grey and red, see Figure 1). These images, inducing the illusion of false colors, could approximate the original image to an astonishing degree. Of course they had a heavy red tint, but due to this being interpreted as a global illuminant, its effect was largely filtered out by the visual system. This experiment and the possible explanations as to why it works are described in Land's Retinex theory (RETINa & visual

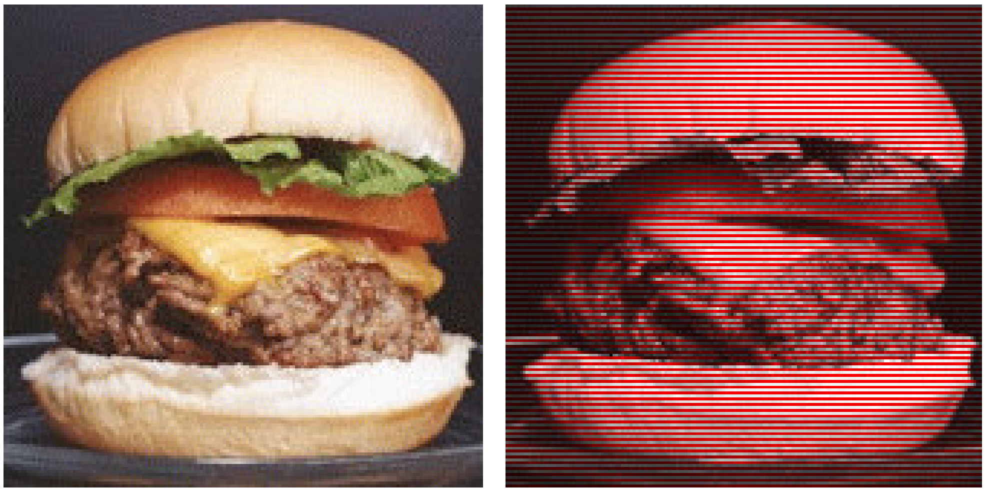

Figure 1: (Left): A full RGB image. (Right): An interlaced image, only using Red and Grayscale values. The lettuce should still appear greenish, the bun and cheese yellow and the mince brown. Note that depending on the display medium there might be more than grey and red displayed.

cortEX). This Retinex theory, even though heavily disputed at its time of inception, has recently regained interest in the research community and is used increasingly in image enhancement applications [3], [1], [2].

1.2 Purpose

As we are generally interested in using optical illusions for the benefit of computer graphics, we are intrigued by the idea of inducing perceived colors. One possible application we imagine is as follows: If full color scenery can be sufficiently approximated using only two colors in conjunction with an appropriate global illuminant, then screen-buffer and texture space requirements would drastically decrease. A less ambitious but similarly profitable approach would be to use Retinex theory as the driving mechanism for texture compression in environments with known lighting conditions. This would be possible, because the lighting conditions together with the texture information would be sufficient to predict perceived color resolution of the given textures. A quantization scheme as employed with most lossy Fourier or Wavelet compression algorithms, could then be used to improve compression ratios, while retaining the best possible image quality. Even in image compression where the scene's lighting conditions are not well known, Retinex could prove useful. In most compression schemes, RGB images are treated as comprised of three independent grey-scale images[1]. Using Retinex theory, we could model the perceived *interaction* between the color channels and use this knowledge to increase compression ratios.

The immediate purpose of our experiments is to demonstrate the feasibility of such an endeavor, if the display device is a cathode ray tube (CRT). The reason why it is important to investigate CRTs, is because most computer graphics are displayed on CRTs with red, green and blue phosphorous pigments. This means that black is produced by not firing any cathode ray cannons, while (perceived) white is produced by firing all of them (i.e. by simultaneously triggering photon emission of red, green and blue wavelengths from the phosphorous screens). Replicating Land's experiment exactly is thus not possible on a standard CRT device, because it is incapable of producing purely grey light.

Our modification of the original experiment (to work with red and green light) and the results obtained are detailed in this report. Also some detail is given on the experimental setup to be performed without supervision using a web-based interface.

2 Test Setup

In the design for our experiment, several issues are considered important:

- *Color choice* - as stated in Section 1.2, we can only use red, green or blue for our filtered sub-images (instead of red and grey as in Land's original experiment).
- *Filters vs. Interlacing* - In Land's original experiment, filters, slides and projectors were used to superimpose the desired frequency bands. As the phosphorous material in a CRT device is spatially separate,

[1] Actually, the only connection between the grey-scale images and the fact that they represent color channels is that different quantisation tables are used for the red, green and blue channels. These tables have been established empirically.

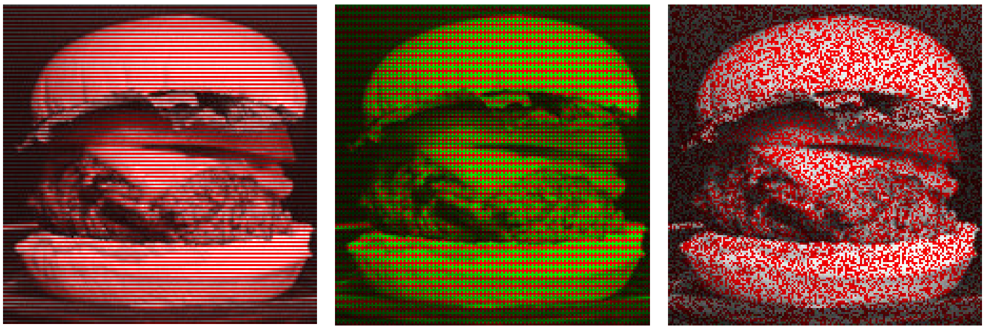

Figure 2: (Left): Image interlaced horizontally. (Center): Image interlaced horizontally and vertically. (Right): Image randomly diffused with two colors. Note that the false color effect is strongly affected by the various interlacing methods.

we cannot truly superimpose sub-images. The closest we can get to a similar effect and still use pure colors is to interlace the sub-images.

- *Real and virtual illuminant* - In Land's experiment, the real world (experimentation laboratory) would be dark, only illuminated via the slide-projector's filtered light. The images thus projected onto the wall can have their own virtual illuminant (i.e. one depicted in the slide). For our purposes, we have to expect a real illuminant to be present (people using CRT devices generally do not work in the dark). There might be an interaction between perceived real and virtual illuminant.

2.1 Design choices

- Due to the antagonistic behavior of the color illusion phenomenon (see Section 1.1) and the tricolor limitations of the CRT device, we decided to use red and green as our two input channels.

- Simply by displaying an image without any blue content on a computer monitor, we use a particular interlacing scheme, which depends on the arrangement of red, green and blue pigments on the screen, which in turn might be manufacturer-dependent. In order to control this arrangement, we manually interlace red and green elements. After experimenting with various interlacing methods (see Figure 2), we

Figure 3: (Left): Original interlaced image. The yellow border is designed to enhance the virtual illuminant effect. The same yellow is used for all other graphics on the page, as well as input elements (drop-down lists, etc.) for the same reason. (Right): A monochrome, outlined (cartoon-style) version of the image is shown below the original, where the user can select the closest choice of perceived colors.

decided to use the horizontal interlacing scheme, because its structural design appears least intrusive.

- In order to gain an insight into the effect of real/virtual illuminant interaction, we asked the user to report on the main illuminant used in the given user setting.
- Furthermore, to test for a variety of viewing conditions, we performed an uncontrolled web-based user-test. The motivation and limitations of this choice are explained in Section 2.2.

2.2 Web-Based Testing

As our initial experiment is designed not to deliver exact quantifiable results, but rather to establish qualitatively whether Land's Retinex theory works on CRT devices under different lighting and viewing conditions, we chose to perform a user-test using a web-interface, so that volunteers could partake in the experiment in their preferred working environment. Based on the success or failure of this initial experiment, further experiments could then be designed in order to pinpoint system variables and their effect on the perceived output images.

The entire experimentation interface is contained within a single webpage and divided into three parts: Instructions, Environment questions and Perception test. The interface itself employs a color scheme that supports the notion of a virtual illuminant (red + green = yellow), which is chosen and calibrated to coincide with the color of the *white* wall in the test image (Figure 3, left).

The environment questions are designed to give a very brief description of the approximate working environment of a given user. The following questions are asked, along with the given answer choices (drop-down lists):

Screen size (inches)	**Eye/Screen distance (cm)**	**Main source of illumination**
<13	<20	Sunlight Bright
13	20-40	Sunlight Middle
14	40-60	Sunlight Dim
15	60-80	Artificial (Std. Bulb)
17	80-100	Artificial (Fluorescent)
19	100-120	Monitor
>19	>120	other

Table 1: Questions and responses for user experimentation setup

For the perception part of the experiment, the user has to look at the left image in Figure 3 and then select the appropriate drop-down box in the right image of Figure 3 to indicate the perceived color of the marked region. To facilitate jumping between the two images, the user can click on either image to scroll to the other one. The available color choices for all regions are the same lest being suggestive. The color list was assembled by combining possible colors for all regions, plus adding other related colors as shown in Table 2. The user also has the option of not selecting a color (null answer) or selecting *Other*, if none of the given choices seem appropriate.

Color			
White	Beige	Black	grey
Light-grey	Dark-grey	Brown	Green
Blue	Light-Blue	Light-Green	Red
Orange	Yellow	Light-Yellow	Dark-Yellow
Other			

Table 2: Colors available for selection by the user

Figure 4: (Left): Original input image; (Center): Retinex image obtained by deleting the blue channel from the left image; (Right): Theoretically perceived image: Normalizing the global illuminant, i.e. deleting the yellow color channel.

The webpage's form-data is submitted to an evaluation script, when the user presses the form's *submit* button. The user's hashed IP address is used for security reasons and to deny multiple submissions, including malicious flooding.

3 Results

3.1 Simulation of Results

Even though we stated earlier, that we are not interested in exact quantitative results at this stage, we are interested in qualitative ones. For this reason, we discuss the results, that we expect to obtain, which are then contrasted with the actual results in Section 3.2.

Figure 5: This image was obtained by computing the color difference between the left and right images in Figure 4. Darker areas correspond to bigger differences.

Neglecting the display device issues discussed in Section 2, we can generate a suitable Retinex image, by taking an RGB image and discarding the blue (B) component (Figure 4, Left). In doing so, we obtain the center image shown in Figure 4, which should closely correspond to the retinal image (i.e. the one seen by the eye, when displayed by a monitor). According to *Color Constancy*, the yellow tint in that image is attributed to a yellow global illuminant and should be subtracted from the retinal image to obtain the perceived image. We simulate this by converting the retinal image from RGB to CMYK and deleting the yellow (Y) component, thus arriving at the theoretically perceived image (Figure 4, Right).

On first inspection we note, that the image is indeed fairly close to the original, if a bit *washed out*. Interestingly, the blue colors have survived very well, while the green colors have become blue. The pink of the skin as well as the orange of the backpack have faded, but are otherwise well retained. The latter is interesting when considering Figure 5, which is the result of creating a color difference image between the left and right images in Figure 4. As darker regions correspond to greater absolute differences, the orange backpack is the region of greatest difference between the images, yet perceptually this difference appears smaller than say the sweater draped around the leftmost person, which is perceived quite differently. This underlines the well-known fact, that L1 and L2 norms are perceptually inadequate in many practical situations.

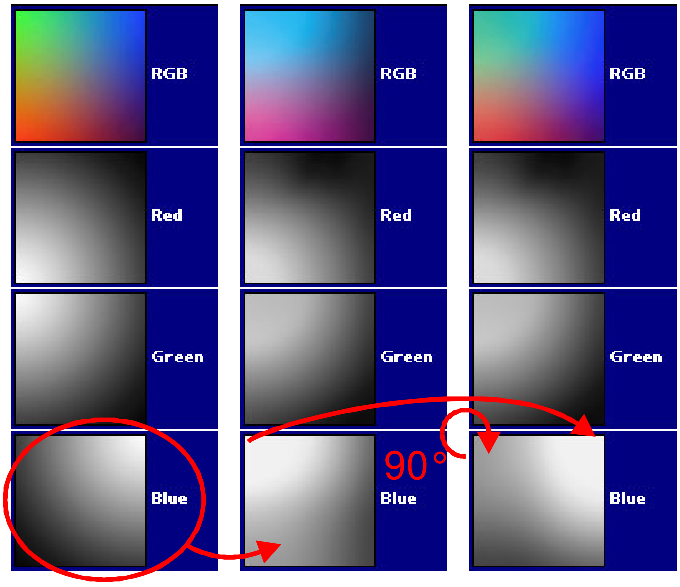

Figure 6: Color Transformations and their effects: (Left column): A homogenous RGB image and its color constituents; (Center column): After applying the retinex and color constancy operations; (Right column): After rotating the blue component of the center column clockwise by 90°.

The disappearance of green color elements and general fading can be explained by examining Figure 6. The Retinex operation (deleting the blue component) followed by the color constancy operation (deleting the yellow component) results in the middle column. We notice, that the red and green components have mostly retained their shape, but have faded slightly. The blue component on the other hand has changed drastically. The overall effect is that of mapping the complete RGB space onto a cyan/magenta-dominated space. This explains the lack of green components in the right image of Figure 4.[2]

[2]Interestingly, the shape of the blue component in the center column of Figure 6 is reminiscent of the original one in the left column. Applying a 90° clockwise transformation on the blue component produces the right column in the same Figure. Even though the colors are clearly faded as compared to the original image in the left column, the image is basically restored.

3.2 Actual Results

After investigating the theoretical results we expect to find, we now evaluate the actual web-test results. It should be noted right from the start, that this preliminary test only includes 14 subjects and is therefore statistically not significant. We can consequently only draw qualitative conclusions, if any. The complete results of the test are listed in Table 3.

Firstly, we find that viewing conditions are fairly homogenous: Most people use a 17 inch monitor at a viewing distance of 40-60 cm in Fluorescent lighting. Results for people with different viewing conditions are not markedly different (see last three columns in Table 3). Overall results for the individual items of clothing are either very good (90-100%), or very bad (30-50%). The question that arises is obviously "*What colors are perceived correctly and which ones are not?*". Interestingly, no conclusive answer can be given. On the one hand, Brown, Black and White/Light-Yellow have very high success rates. On the other hand, Green and Blue can provoke a variety of responses. While the Blue of Person 4 Pants is recognized very well, the very similar color of Person 3 Pants is not recognized so well. The Green of Person 1 Sweatshirt is identified extremely poorly, while all test-subjects detected the greenish-blue hue of Person 2 Shirt.

We deduce from this, that *color expectancy* plays a large part in color perception of difficult scenarios. What we mean by this, is that a person might have trouble deciding the color of a given material from the given reflected light information alone. In order to make a better guess the person then assumes a certain material (e.g. jeans), which has traditional colors associated with it (e.g. the indigo dye used in jeans manufacturing). Using this argument, the long pants of Person 4 are likely to be made of jeans material and therefore blue (identified well), whereas the shorts of Person 3 are not commonly made of this material (identified poorly). Similarly, the shirt of Person 2 could be seen as either Light-blue or Light-green in the original photograph (blue and green contributions are almost identical), which was decided by all test-subjects. Yet, most people opted for Light-blue, probably, because that color is more often seen on shirts like the one presented.

For red colors, the only close example available, is the orange Backpack. Many people mistook the orange for a red. This falls in line with our theo-

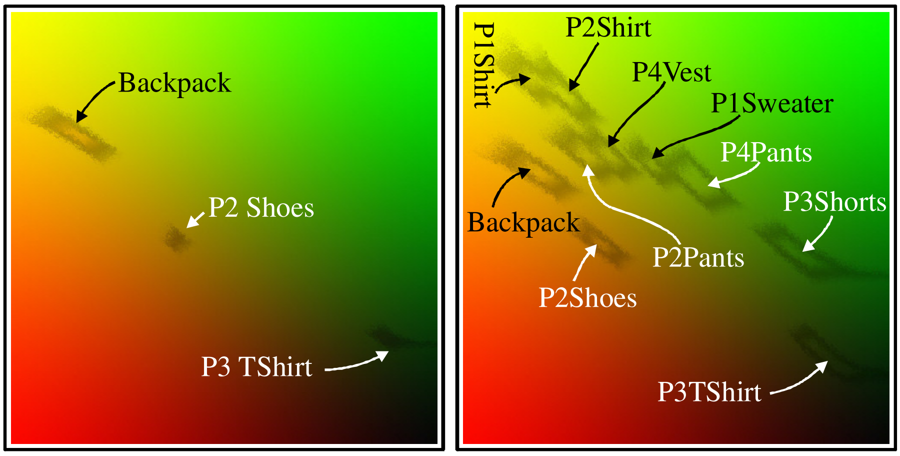

Figure 7: (Left): Gamut-matching of original image vs. retinex gamut. Only the Backpack, Person2 Shoes and Person3 T-Shirt are minimally contained within the retinex gamut, i.e. could possibly be reproduced using no blue colors. Regions have been enhanced artificially and slighlty enlarged to be visible. (Right): Gamut-matching of retinex image vs. retinex gamut. Logically, all items map into the gamut space.

retical results (Figure 4, right), which is measured as a pinkish color. Since all colors are desaturated through the retinex process, they might be perceptually boosted to natural levels, which would then transform pink into a perceived red.

4 Conclusion

Foremost, we believe that the experiment has been a success. We were able to produce varied color sensations clearly outside the gamut of the displayed image. In fact, only colors of the backpack are contained within a small portion of our chosen retinex gamut (red/green), see Figure 7, Left. If the matching algorithm is executed with fuzzy colors, then parts of Person2 Shoes and Person3 T-Shirt can be made to match, but only barely so. This means that the seven remaining items, which lie completely outside the retinex gamut can be simulated under various lighting conditions, using Retinex theory in conjunction with Color Constancy theory, on CRT display devices.

Various test-specific problems exist, most of which we discuss in Section 5.

5 Future Work

Our test merely shows that Land's Retinex experiment can be replicated in modified form on modern display devices. There are other problems that we didn't solve, like the question "What exactly influences the perception of the false colors?" (e.g. what, if any, effect do room-lighting conditions have). In order to obtain answers to these questions, we need to perform an extended test-series with controlled lighting conditions. Even more importantly, we need to include a significant number of test-subjects into our evaluation, for the results to be meaningful.

With respect to the existing test, we found that color naming is difficult (a particular branch of psychology is devoted to this topic), because the color-name association is different between individuals and not directly measurable. Developing a system, where the user can employ a *color-picker* tool would be more useful and easier to evaluate.

The issue of *expected color* perception was raised in Section 3.2. We would like to find out more about this phenomenon with respect to Retinex theory. If for example the picture of an apple were to be modified so that the apple appeared blue and this image used as the input to the retinex process, what would the result be? Are natural scenes perceived differently from artificial scenes [3] ? This question is especially important, as we are interested to see if the Retinex method can be gainfully employed for computer graphics.

The colors chosen in our experiment agreed well with the problems we anticipated (i.e. blue/green similarities), but due to the lack of other color combinations (e.g. more red tones) we may have missed interesting phenomena that we didn't expect from theory. The variety of input images thus has to be diversified and content chosen carefully.

[3] Land himself tackled this problem by using abstract images resembling the artwork of the old dutch master Piet Mondrian. The experiments thus also became known as the Land-Mondrian experiments. One important result from this research is that subjective color can be induced from abstract images, implying that expectancy is not necessary. It should be noted, on the other hand, that it does not follow that expectancy has no or little influence on perceived color.

References

[1] Barrow, H., and Tenenbaum, J. Line drawings as three-dimensional surfaces. *Artificial Intelligence 17* (1981), 75–116.

[2] Cheek, D. *True-to-Life Imaging: Retinex Image Processing.* NASA TechTracS, http://technology.nasa.gov/scripts/nls_ax.dll/w3SuccItem(202301).

[3] Funt, B., Barnard, K., Brockington, M., and Cardei, V. Luminance based multi scale retinex. In *Proceedings AIC Colour 97 Kyoto 8 Th Congress of the International Colour Association* (May 1997).

[4] Hurvich, L. *Color Vision.* Sinauer Assoc., Sunderland, Mass., 1981.

[5] Jameson, D., and Hurvich, L. Some quantitative aspects of an opponent-colors theory. i. chromatic response and spectral saturation. II. brightness, saturation and hue in normal and dichromatic vision. *Journal of the Optical Society of America 45*, 7,8 (1955).

[6] Land, E. H. The retinex theory of color vision. *Proceedings of the Royal Institute of Great Britain 47* (1974), 23–58.

A Detailed Test Results

User#	Screen Size (inch)	View Distance (cm)	Main Source of Illumination in Room	Person 1 Pants	Person 1 Sweatshirt	Person 2 Shirt	Person 3 Shirt	Person 4 Vest	Back-pack	Person 4 Pants	Person 3 Pants	Person 2 Pants	Person 2 Shoes	Answers given	Correct	Percent Correct
Correct	*n.a.*	*n.a.*	*n.a.*	*Beige/ White / Light-Yellow*	*Green*	*Light-blue/ Light-green*	*Black*	*Grey*	*Orange*	*Blue*	*Blue*	*Brown/ Beige*	*Brown/ Orange*			
1	19	20-40	Other	Beige	Dark-Gray	Light-Blue	Black	Light-Grey	Orange	Blue	Green	Brown	Brown	10	8	80.00%
2	17	40-60	Artificial (Fluorescent)	White	Light-Blue	Light-Blue	Black	Grey	Orange	Blue	Blue	Beige	Brown	10	9	90.00%
3	>19	40-60	Artificial (Fluorescent)	Beige	Green	Light-Blue	Black	Light-Grey	Orange	Blue	Blue	Beige	Brown	10	10	100.00%
4	>19	40-60	Artificial (Fluorescent)	Light-Blue	Blue	Light-Blue	Black	Undefined	Undefined	Blue	Blue	Beige	Brown	8	6	75.00%
5	17	20-40	Artificial (Fluorescent)	White	Blue	Light-Blue	Black	Light-Grey	Orange	Blue	Blue	Beige	Brown	10	9	90.00%
6	<13	20-40	Artificial (Std. Bulb)	White	Light-Blue	Light-Blue	Black	Light-Blue	Red	Blue	Light-Blue	Brown	Brown	10	7	70.00%
7	17	40-60	Artificial (Std. Bulb)	Light-Yellow	Green	Light-Green	Black	Grey	Other	Green	Green	Undefined	Brown	9	6	66.67%
8	15	40-60	Artificial (Fluorescent)	Beige	Green	Light-Green	Undefined	Black	Red	Blue	Blue	Other	Brown	9	6	66.67%
9	17	40-60	Artificial (Fluorescent)	White	Blue	Light-Blue	Black	Grey	Red	Blue	Green	Beige	Brown	10	7	70.00%
10	17	60-80	Artificial (Fluorescent)	Beige	Blue	Light-Blue	Black	Beige	Orange	Blue	Red	Beige	Brown	10	7	70.00%
11	19	40-60	Artificial (Fluorescent)	White	Blue	Light-Blue	Blue	Beige	Red	Blue	Green	Brown	Undefined	9	4	44.44%
12	17	40-60	Sunlight Dim	White	Light-Blue	Light-Blue	Black	Light-Green	Red	Blue	Blue	Beige	Brown	10	7	70.00%
13	17	40-60	Artificial (Fluorescent)	White	Green	Green	Black	Grey	Red	Blue	Green	Beige	Brown	10	8	80.00%
14	17	40-60	Artificial (Fluorescent)	Beige	Green	Light-Blue	Blue	Light-Blue	Red	Blue	Green	Brown	Brown	10	6	60.00%
			Answers Given	14	14	14	13	13	13	14	14	13	13			
			Correct	13	5	14	11	7	5	13	7	12	13			
			Percent Correct	92.86%	35.71%	100.00%	84.62%	53.85%	38.46%	92.86%	50.00%	92.31%	100.00%			

Table 3: Complete Web Test Results