https://doi.org/10.2352/ISSN.2169-2629.2017.32.203

Three-Dimensional Test Target for Illuminant Analysis

Nathan Moroney, Ingeborg Tastl, Melanie Gottwals
HP Inc. Laboratories

Abstract

Test targets, such as color charts, are often used for characterizing image capture systems. This paper describes a test target that uses three-dimensional sinusoids of varying orientation to analyze the illumination for a given scene. Specifically the magnitude of the contrast of the imaged sinusoids is used to estimate the angular orientation of the illuminant. A target consisting of sixteen sinusoids sampled every 11.25 degrees is shown to estimate the angular location of the illuminant to within 6.6 degrees in a laboratory setting. Furthermore this accuracy is achieved even with the introduction of a diffuser into the scene. This target is also used to differentiate two illuminants of differing chromaticity. In addition the target is used to differentiate the relative diffuseness of various locations on the bottom of a light booth. Finally the virtual target is also useful for understanding the illumination of 3D visualization software. The design and production of the target is described in more detail as are additional applications and use cases. The introduction of known three dimensional structure into test targets is a promising area for the design of targets for evaluation of image capture systems.

Introduction

Color Charts like the classic Macbeth ColorChecker have been used for a long time for the color characterization of image capturing systems (e.g. scanner, digital cameras) [1]. Recently, they have also been proposed for cell phone apps aiming to obtain more accurate estimations of colors in a scene (e.g. fabrics, interior design objects). In all those cases the assumption is made that the illumination of the color chart is the same as the illumination of the objects of interest. In an ideal case the illumination is diffuse. That brings up an interesting research question, namely how to determine in real time if the illumination is indeed diffuse or not. For human observers it is easy to detect strong directional light, but not so easy to detect non-uniformities the closer one gets to a diffuse illumination. Other use cases where knowing if one or several light sources are present and from which direction the light is coming from are automatic white-point detection and balancing algorithms used in camera pipelines. Augmenting an image of a real scene with computer rendered objects and aiming for a seamless high quality composition, a task that is quite common for augmented reality also requires a quite accurate estimation of the direction and intensity of the light source(s) [5]. Yet another application is shadow removal of outdoor or indoor scenes for a variety of computer vision applications [6]. Last, but not least scene reconstruction algorithms based on shape from shading techniques also require information about the lighting.

Papers in the area of illumination detection with the focus on different use cases have been published for the last 35 years, starting with the paper from Pentland in 1982 [2]. From a high level point of view they can be distinguished in ones that apply statistical image analysis of different complexity [4], or of light probes that are being put into the scenes. A light probe is an object of known 3D shape and reflectance function that is being put into a scene when the image is captured. Quite often that light probe is a sphere of either specular or matte material [3].

The goal of this paper is to expand the well-proven concept of a traditional color chart that has been used for color characterization to a version that can be used to both visually and automatically detect the direction of one illuminant in a robust way, to detect the presence of two illuminants and the respective directions and if the illumination in a scene can be categorized as diffuse or not. The fact that the chart provides a “human-readable” easy to interpret feedback also distinguishes the approach described in this paper from other approaches. Obviously the algorithm could also be deployed in form of a cell phone app.

This target also has some conceptual similarity to sundials in which a gnomon [7] or shadow casting element is used to infer the position of the sun, and hence the time of day. However this target in some ways inverts this scheme, sinusoidal tiles aligned with the direction of illumination will have the least shadows and all other tiles will have increased shadows formed. As a three dimensionally printed target the object can also be compared to other efforts to use three dimensional printing for optics [8], custom fabricated microgeometries [9] and use of self-occlusion for imaging [10]. In this case though the target is a form of macro-scale structured self-occlusion that can be used to analyze illumination. Finally known colored illuminants at fixed locations have been used to estimate the three dimensional shape of unknown objects [11]. That is lighting of known geometry and color has been used to estimate unknown surface geometries. This paper proposes making use of known surface geometries to estimation unknown illumination conditions.

Target Design

Initial versions made use of a discrete sampling of sinusoidal surfaces. The sampling and surface functions are a design consideration and this version is one that is conceptually simple and yet effective. The target is 15 by 15 by 0.7 dimensions in centimeters. The sinusoids were of amplitude 3 millimeters and frequency 5 millimeters. The sinusoids varied in their orientation and were arranged as a four by four sequence of tiles in which the angle of orientation of the sinusoids varied in steps of 11.25 degrees. The charts were 3D printed using a fused deposition modeling or FDM device. An approximately white, lower gloss material was used with no additional post processing. The chart was constructed using simple C++ function objects to directly render the corresponding STL file. Figure 1 shows an image of the STL model for the target on the top and a photograph of the actual chart as printed on the bottom of the figure. Some versions of the chart included a frame for ease of handling and inclusion of a colored frame for automatic detection and processing.

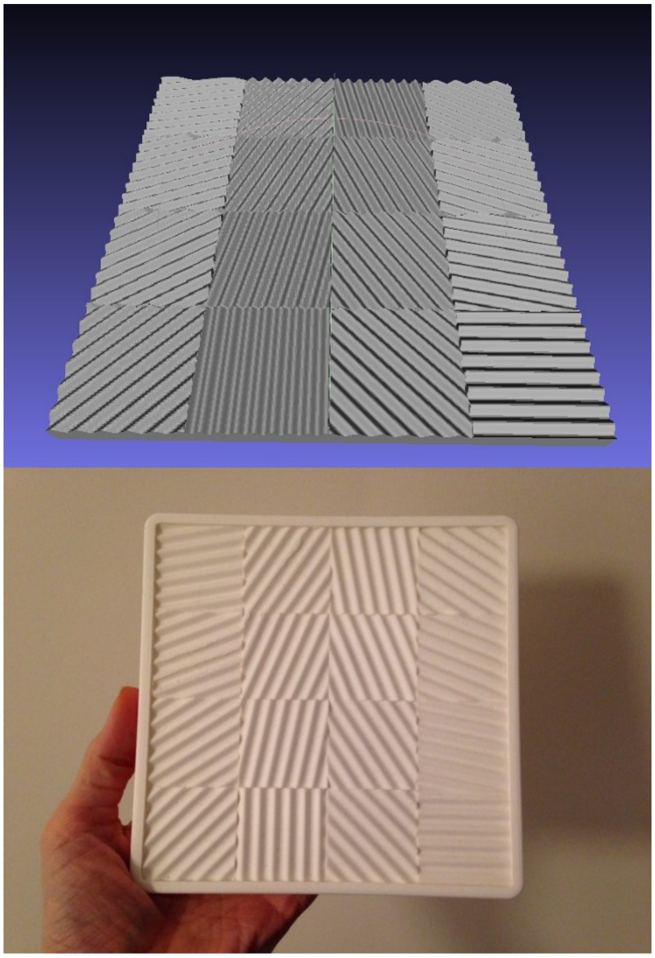

Figure 1. Rendering of the 3D test target, shown on top, and three dimensional print of the target, shown on the bottom.

Directional Illumination

As the first laboratory test of the chart, a strongly directional lighting configuration was tested. Specifically a 2D printed angular guide was attached to an optical bench and 35 cm guide wire was attached to a lamp mounted on a fixed base. The lamp used a compact fluorescent bulb with a specified color temperature of 5500K and 30 watts of power. The lamp was kept at a constant distance to the center of the chart but was rotated around the chart in 10 degree increments starting at 0 degrees and ending at 90 degrees. Figure 2 shows an example of one of the directional illumination configurations. The lamp was the only illumination used during these tests.

The analysis of the sinusoidal chart proceeded in the following steps. First the location of the chart corners and guide-wire were manually specified in terms of pixel coordinates. The angle formed by the guide wire is used as the ground truth for the performance testing. Note that previous and ongoing research makes use of automated chart detection and extraction but these algorithms will not be described in this abstract. Given the location of the corners of the charts the pixels in each of the 16 tiles of the sinusoidal chart were sampled and analyzed. A 50 pixel wide by 50 pixel high region around the center of each tile was used to sample the pixels of the imaged sinusoids. The 5th and 95th percentiles of the luminance distributions were computed for each tile. The luminance difference between the 95th and 5th percentiles was then computed for each of the tiles. Finally the tile with the minimum luminance difference between the 95th and 5th percentile was used as the estimated angular orientation of the illumination with respect to the chart. Summary statistics were then computed for all 10 angle increments and the average absolute value angular error was computed to be 6.6 degrees.

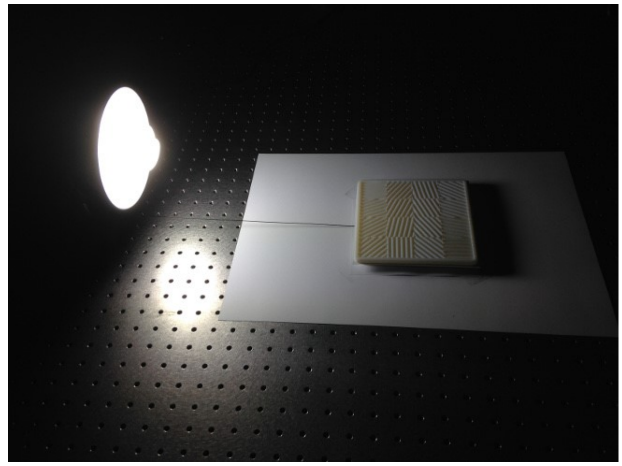

Figure 2. Test target, in center, with directional illumination, on left, connected with a black guide wire.

To visually represent all of these steps and results a visualization scheme was used in combination with the original image. Figure 3 shows one of the test images in which the manually identified corners are shown as red dots and the manually identified guide wire or ground truth are shown as blue dots. For each tile in the test chart, a bubble chart is overlaid on the image. The bubble is color coded according to the 5th, on left, and 95th, on right, luminance percentiles. The bubble radius is the square root of the luminance difference between the 95th and 5th percentiles. The smaller the bubble the lower the luminance contrast and the larger the bubble the greater the luminance contrast for the corresponding sinusoid. The tile with the minimum luminance difference is shown with a red line behind the bubble. This red line is drawn with the corresponding angle of the sinusoid for that tile in the target. Note that of the two orientations possible for this red vector, the direction is selected that is towards the side with the largest 5th and 95th percentile luminances.

An additional test of directional illuminant was performed by varying the azimuth or angle of illumination with respect to the plane of the chart. The chart was captured in steps of 10 degrees from 40 degrees to 80 degrees. For the 80 degree illumination this corresponds illuminating the chart from almost behind the camera during capture. The estimated angular orientation was accurate to ±10 degrees for azimuth angles of up to 70 degrees. At 80 degrees there was a larger error of 30 degrees but in this case the diffuseness index, which will be described in more detail in the next section, was at its lowest and in fact as the illumination approaches the angle of capture this resembles a diffuse illumination condition.

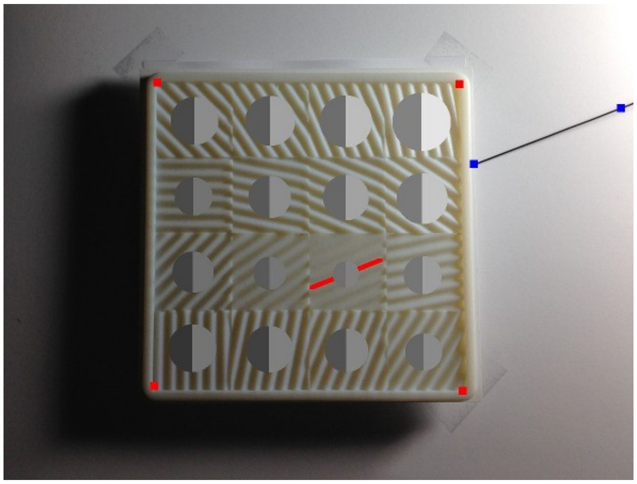

Figure 3. Visualization of experiment one and analysis of the target.

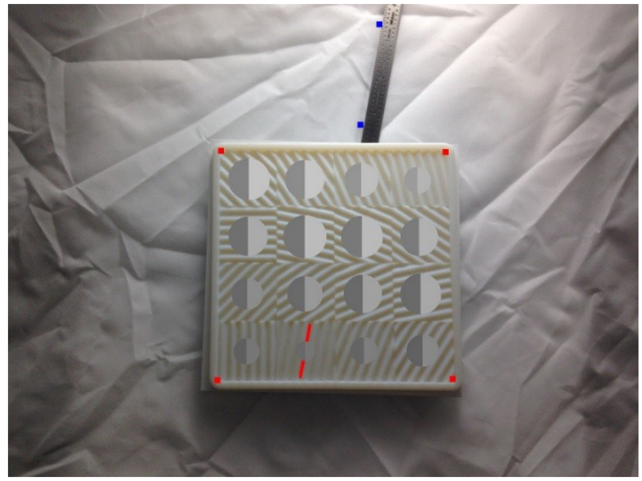

Figure 5. Visualization of experiment two and analysis of the target.

Diffuse Illumination

A second experimental test cases makes use of a diffuser to test the accuracy of the orientation estimation given more diffuse illumination. This experiment made use of identical geometry and positioning as described in the above directional illumination test case. The only difference was the introduction of a white fabric diffuser box in which the test target was placed. This arrangement is shown in Figure 4.

Figure 4. Experiment two used the same lighting and positioning shown in figure 2 but in this case the target is placed within a diffusing box and the illumination is outside of this box.

Identical sampling of 0 to 90 degrees in 10 degree increments was used. Likewise identical processing and analysis were used to estimate the angular orientation of the illumination with respect to the chart. An example analysis image for one of the test images is shown in Figure 5. The resulting average absolute value for the errors is 5.1 degrees. This error is comparable to that achieved with the strongly directional illumination. More data collection is required before performing additional statistical means testing of these results.

As an additional experimental test of the target in diffuse illumination conditions the test target was placed at the bottom of a light booth. The target was placed is various locations and photographed. In this case the illumination is quite diffuse but there is still some spatial variation in the luminance. With the target at the bottom of the booth the question is less the orientation of the light source and instead which position is most diffuse. Figure six shows images of the chart as captured in three locations in the bottom of the light booth.

To quantify "most diffuse", the following analysis is performed. First the fifth and ninety-fifth percentiles for the estimated luminance for each of the patches or tiles is computed. However instead of identifying the minimum difference in the ninety-fifth and fifth percentiles, the minimum range in the differences is computed. That is for each chart find the maximum difference in the ninety-fifth and fifth percentiles in luminance and also find the minimum difference. Given all of these differences the chart with the minimum range is then the most diffuse location. The range or difference is referred to here as the diffuseness index. This type of analysis is challenging for other lighting estimation algorithms and is of potential general use in the case where lighting or positioning of an object should be as diffuse as possible. It is also possible to use this measure as a single index of diffuseness. That is instead of comparing multiple images, a single image or video stream could be analyzed and the corresponding value reported to the user.

As ground truth data for this experiment the luminance of the corresponding locations of the light booth were measured with a Minolta CS-100. A three by three grid of measurements was taken at each position, without the chart in the scene, for a total of nine measurements. The luminance range in cd/m2 was then computed for each position. These luminance ranges were then plotted versus the diffuseness index as is shown in Figure 7. The correlation between the two metrics was 0.989 and suggests that sine chart is effective at discriminating the differences in diffuseness for different positions in the light booth.

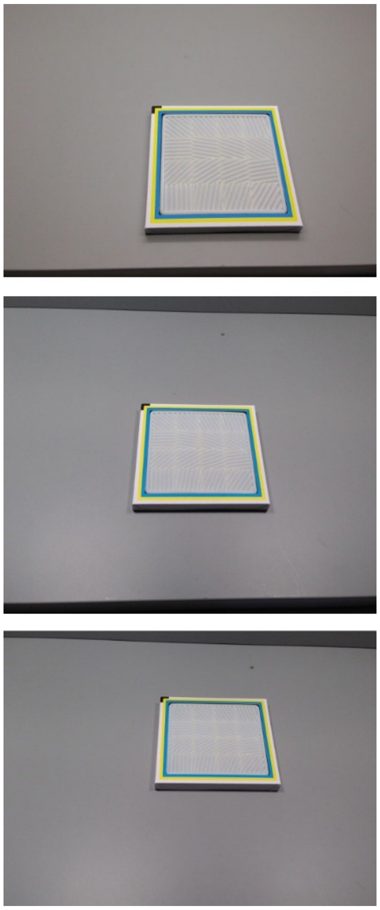

Figure 6. Identification of the most diffuse location for the target as placed in a light booth. Three different positions were tested as shown above.

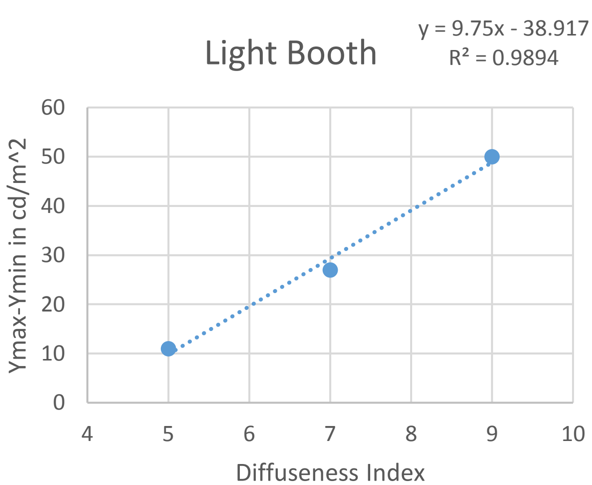

Figure 7. Computed diffuseness index based on the chart analysis versus the measured luminance range values in cd/m2. The leftmost point corresponds to the bottom image in Figure 6 while the rightmost point corresponds to the top image in Figure 6.

Additional Applications

There is also additional ongoing experimentation in the following areas. First the target is being analyzed with more than one illuminant. Figure 8 shows an example configuration with two highly chromatic illuminants in the upper left. A polar transform of the RGB chromaticities is shown in the upper right. The bottom two images are the analysis results using color separations of the image according to the two peaks of the angular histogram. Second the target is also being used to assess virtual illumination environments. For example, the model of the target can itself be used to visually infer the location nature of virtual illuminants in visualization software. Third the target detection and analysis process has been integrated with previous automated chart detection. This allows a more detailed assessment of the angular accuracy data presented in the earlier sections. Finally with the analysis fully automated it is possible to apply the target to video data and assess the performance of the chart outside of the laboratory.

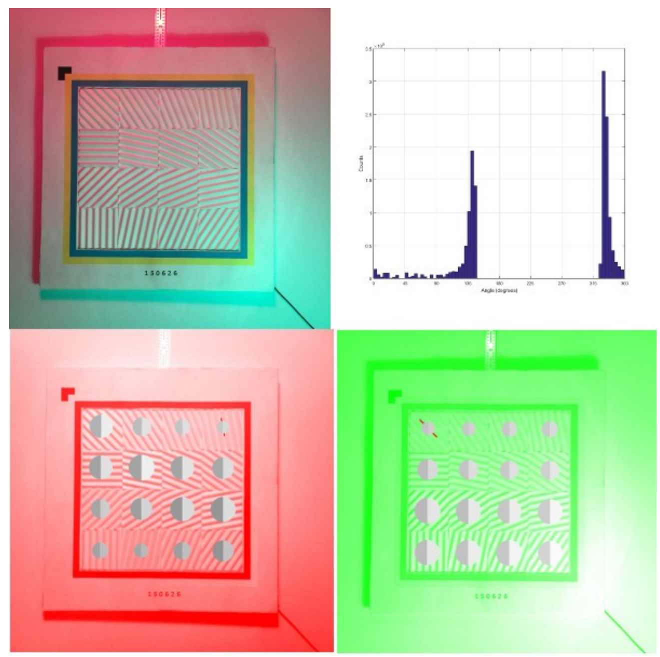

Figure 8. Test target as illuminated by two highly chromatic test illuminants (upper left) is converted to polar form using RGB chromaticities (upper right). The corresponding analysis based on color separations based on peaks in the histogram.

Conclusions

A test target consisting of three dimensional sinusoidal tiles or patches has been described and it use for the estimation of the orientation of the illuminant has been quantified. Its performance has been tested for a strongly directional and more diffuse laboratory illumination, as well as for quantifying the most diffuse location at the bottom of a light booth. The angular estimations were shown to be accurate to within 6.6 degrees and the diffuseness index was highly correlated with actual luminance measurements. Four additional areas of testing were also briefly described in the discussion section.

References

[1] McCamy, Calvin S., H. Marcus, and J. G. Davidson. "A color-rendition chart." *J. App. Photog. Eng.* 2.3 (1976): 95-99.

[2] A. P. Pentland."Finding the Illuminant direction", *J. Opt. Soc. A*. /Vol. 72, No 4 / April 1982.

[3] Y. Zhang, Y. Yang. "Multiple Illuminant Direction Detection with Application to image Synthesis", *IEEE Transaction on Pattern Analysis and Machine Intelligence*." Vol. 23, No. 8, August 2001.

[4] J. Lopez-Moreno, S. Hadap, E. Reinhard, D. Gutierrez. "Compositing Images through Light Source Detection", *Computer and Graphics*, (2010), doi:10.1016/j.cag.2010.08.004.

[5] H.E. Gerhard, L.T. Maloney. "Estimating changes in lighting direction in binocularly viewed three-dimensional scenes." *Journal of Vision*, (2010) 10(9):14, 1-22.

[6] C. Fredembach, G. Finlayson. "Detecting Multiple Illuminants", www.researchgate.net.

[7] Vincent, J., "The mathematics of sundials." *Australian Senior Mathematics Journal,* 22.1 (2008): 13.

[8] Willis, Karl, et al. "Printed optics: 3D printing of embedded optical elements for interactive devices." *Proceedings of the 25th annual ACM symposium on User interface software and technology*. ACM, (2012).

[9] Weyrich, Tim, et al. "Fabricating microgeometry for custom surface reflectance." ACM Transactions on Graphics (TOG). Vol. 28. No. 3. ACM, (2009).

[10] Alexa, Marc, and Wojciech Matusik. "Images from self-occlusion." Proceedings of the International Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging. ACM, (2011)

[11] Finlayson, Graham D., and Christopher Powell. "Shape in a Box." European Conference on Computer Vision. Springer International Publishing, (2014).

Acknowledgements

The authors would like to thank Evan Johnson for his assistance with the laboratory data collection.