To appear in *Journal of Imaging Science and Technology*, 64(5): 000000-1000000-8, 2020

This is a preprint version and not identical to the published version. For the version of record, see JIST.

A Flying Grey Ball Multi-illuminant Image Dataset for Colour Research

Hoda Aghaei and Brian Funt; Simon Fraser University, Vancouver, B.C., Canada

Abstract

For research in the field of illumination estimation and colour constancy there is a need for ground truth measurement of the illumination colour at many locations within multi-illuminant scenes. A practical approach to obtaining such ground truth illumination data is presented here. The proposed method involves using a drone to carry a grey ball of known percent surface spectral reflectance throughout a scene while photographing it frequently during the flight using a calibrated camera. The captured images are then post-processed. In the post-processing step, machine vision techniques are used to detect the grey ball within each frame. The camera RGB of light reflected from the grey ball provides a measure of the illumination colour at that location. In total, the dataset contains 30 scenes with 100 illumination measurements on average per scene. The dataset is available for download free of charge.

Introduction

A technique is proposed for measuring the chromaticity of the incident illumination at many locations within both indoor and outdoor real world scenes, with the goal of creating a useful database for research on illumination estimation and colour constancy. The strategy employs a drone to carry a calibrated grey ball throughout a scene while it is simultaneously being photographed by a calibrated camera at a fixed location. Using the proposed technique, a database of 30 scenes with 100 to 150 images per scene with accompanying ground truth illumination measurements is constructed and made publicly available.

Figure 5 shows the complexity of the illumination in some of the typical scenes from the database. For example, the image in the second column, second row (from the top) includes both direct sun and light filtered through tree leaves creating strong shadows. The image in column one, row two includes both indoor light and light entering from a window. Column two, row six is an example of multiple indoor lights. Clearly, a single measurement of the scene illumination using a grey card or a Macbeth ColorChecker will not suffice as a measure of the scene illumination in these cases.

Among the many colour constancy approaches such as the works by Buchsbaum [1], Land [2], Van De Weijer and Gevers [3], Finlayson and Trezzi [4], Barron and Tsai [5], perhaps the most common approach involves two steps: first, the illumination chromaticity (which is assumed to be constant throughout the scene) is estimated; and second, the estimated illumination is used to adjust the image colours to appear as similar as possible to how they would be under some standard chosen 'white' light. However, the illumination chromaticity is rarely constant throughout a scene. Recently, methods have been proposed for illumination estimation for multi-illuminant scenes such as the works by Bleier et al. [6], Zhu and Funt [7], Gao et al. [8], Bianco et al. [9], Beigpour et al. [10]. To evaluate the performance of these methods, the ground truth illumination chromaticity at every, or at least many, image locations is required over a significant number of scene. This paper describes the construction and use of just such a dataset. It overcomes the problem of needing to have many calibration targets placed throughout the scene by using a drone to carry a single target to many locations.

This work aims to overcome these difficulties by using a new technique that involves moving the illumination measurement target within the scene with the help of a drone. This idea facilitates the process of taking multiple pictures in order to capture the illumination's chromaticity at multiple positions within each scene. The resultant dataset can be used to develop illumination estimation methods for real-world, multi-illuminated scenes.

Related Work

There are many large image datasets, such as COCO, which are available for use in a variety of image processing tasks. However, few image datasets satisfy all the necessary criteria required for colour constancy and illumination estimation projects. One such factor is that images should be in RAW format, which is before any gamma adjustment or white balancing methods are applied. While some datasets, such as RAISE, may provide images in RAW format, they do not provide any information about the scene illumination.

The most important and challenging feature required of an image dataset for colour constancy and illumination estimation is that it should provide the ground-truth value of illumination chromaticity for each surface in the scene, preferably at every pixel. As it is a particularly challenging task, simplifying assumptions are usually made about scenes. These assumptions affect the approaches and instruments that are used to measure the illumination chromaticity.

The most common assumption is that the illumination chromaticity is uniform across each scene. Therefore, only one measurement of the illumination chromaticity would suffice for the entire scene. This measurement could be captured by measuring the RGB of a standard calibrated grey card. An ideal grey surface reflectance ($S(\lambda)_{visible} = 0.5$) will have $R = G = B$ under ideal white light ($E(\lambda)_{visible} = constant$), and thus any deviation from this represents the illumination's chromaticity. There are several image datasets available based on the uniform-illumination assumption, including those by Gehler et al. [11], Shi and Funt [12], Hemrit et al. [13], Cheng et al. [14], Banić and Lončarić [15], Ciurea and Funt [16].

The first version of the Colour Checker dataset was released by Gehler et al. [11]. It consists of 568 images provided in both RAW and nonlinear gamma adjusted format. A Macbeth colour chart was placed in each scene to measure the illumination chromaticity. This dataset has been widely used in colour constancy research due to its high-quality images and large number of uncorrelated images. In 2011, Shi and Funt [12] provided a reprocessed version in linear, 16-bit TIFF format derived directly from the original RAW image dataset. However, there was some ambiguity in the instructions provided with the dataset that led to inconsistency in how researchers interpreted the black level for a subset of the images. To overcome these problems, new ground-truth data has been recommended by Hemrit et al. [13] as the new standard.

The NUS dataset by Cheng et al. [14] contains 1,736 images from 8 different cameras with a colour checker in each image. As the images are taken with different cameras, this is helpful for carrying out analysis independent of the cameras. Information is provided in Matlab metadata file format (*.mat*) for each image, including the coordinates of the colour checker, masks for the colour patches on the colour checker, and the black level of each camera.

The Cube image dataset by Banić and Lončarić [15] consists of 1,365 outdoor images taken with a Canon EOS 550D. For this dataset, a measurement object named SpyderCube is placed at the lower corner of each scene. There is also an extension to the dataset that adds 342 images, including indoor and night-time outdoor images named The Cube+ image dataset [17]. The SpyderCube calibration object has the advantage that it includes faces at different angles and the faces include different shades of grey, including 18 percent grey, 96 percent white, 4 percent black, and a hole which is close to absolute black. These different shades allow users to adjust and correct the exposure and the black level of the images. Additionally, in scenes with multiple illuminants it is useful to have two faces as a measure of the illumination.

For the Simon Fraser University (SFU) grey ball image dataset Ciurea and Funt [16], a grey ball attached to a stick so that it is held in the camera's field of view is used to measure the chromaticity of illumination. This dataset is one of the largest image colour constancy datasets with approximately 11,000 images but there is some correlation between the images. The NTSC RGB (not sRGB) of the ball are provide. However, the values are recorded after post-processing, so they are not linear RAW values and are effectively uncalibrated.

The uniform illumination assumption does not hold for most real-world scenes. The challenge is how to efficiently measure the illumination throughout a scene. Some approaches to this problem include those by Bleier et al. [6], Bell et al. [18], Hao et al. [19], Gijsenij et al. [20], Nascimento et al. [21]

The image dataset by Bell et al. [18] includes more than 5,000 indoor images that have been annotated using crowd-sourcing. The annotations provide ground-truth estimates of surface reflectance properties based on human judgment. During the crowd sourcing, some sample points are selected across the images. Following this, users are asked to give their opinion about a pair of sample points as to which point has darker reflectance value. The dataset is quite large and provides beneficial information about surfaces in scenes. However, as the authors point out the annotation ignores the effect of coloured illumination on what users perceived as the colour of the surface. Another limitation of the work is that the users only give their judgments on the reflectance's intensity and not its chromaticity. This particularly raises issues when there is potentially a comparison between two points with distinctly different colours. In this case, it is a difficult task to have an accurate opinion about the comparison of the brightness of the points.

Another image dataset by Bleier et al. [6] was generated under laboratory conditions and contains 408 images. There are 4 different scenes, mostly composed of objects with diffuse surface materials, with 17 illumination conditions for each scene. The images are captured under different exposures and shutter times in RAW format. Following this, the images were converted to linear images by *DCRAW* without any white balancing. To measure the ground truth, the entire scene was painted with grey paint and re-photographed. Although it is a creative method, it is not a practical approach for real-world scenes, and the images in the dataset are of limited variety.

Hao et al. [19] use ray-traced computer graphics to render full-spectrum, physically correct as possible, photorealistic images. Since images are synthesized, the ground truth is available at every pixel. The data set provides 1,000 stereo pairs. The renderings are limited to indoor scenes.

Gijsenij et al. [20] placed several grey balls in each scene in order to measure the illumination chromaticity at several scene locations. Images were captured with a Sigma SD10 camera with a Foveon X3 sensor with a spatial resolution of 384x256. The camera's white balancing mode was set to the overcast setting for the entire dataset. The images are captured in RAW format and converted to linear sRGB for experimentation purposes. The dataset consist of 59 images under laboratory settings and 9 outdoor images, so the number of real-world images is still quite limited.

Nascimento et al. [21] use a similar approach but capture hyperspectral rural and urban scenes with grey spheres embedded in each. Their goal is to analyse spatial variations in local illumination. The spectral information of the local illumination in each scene is extracted from which a total of 1,904 chromaticity coordinates, and correlated colour temperatures, (CCTs) are derived. In each scene there are multiple spheres and from each sphere, 17 sample points are uniformly selected.

As can be seen from these projects, it is clear that capturing the ground-truth value for illumination chromaticity in real world scenes is a challenging task. Hence, the grey-ball-on-a-drone method is proposed as a possible solution. A camera at a fixed location and orientation is used to photograph the grey ball on the drone as it is flown through the scene. The images are then processed using machine vision techniques to locate the grey ball in each frame, record its RGB, adjust for the percent surface spectral reflectance of the ball, and thereby obtain the chromaticity of the light incident upon the the ball at each location within the scene.

Method and Experiment

Hardware Setup

The camera used to capture images is a Nikon D700 DSLR with a Nikon 50mm 1:1.4G lens. The spectral properties of the camera's sensor using this lens are shown in Figure 1. The shutter is activated by a remote control in 'serial shots' mode. The exposure is automatic, and then the exposure value is recorded separately. Images are saved in RAW format, which is before any white balancing or gamma correction is applied. The images are therefore linear and require demosaicing.

A light plastic ball 13 cm diameter was painted uniformly with a high-value-low-pressure (HVLP) sprayer, using Munsell N5 Neutral Matt Grey Vinyl Latex Emulsion Paint. The spectral reflectance of the ball is measured under standard daylight simulator, D65. The measurement then is divided by the spectrum of the D65 simulator. The final result is shown in Figure 2. The spectrum is measured using the Photo Research SpectraScan PR650 spectroradiometer. As shown in Figure 2, the spectral reflectance of the grey paint is effectively flat across the visible wavelength range.

A drone is used to carry the grey ball around the scene and specifically near its surfaces. Figure 3 shows the drone and the grey ball attached to it. The drone is a CineBee 4K Whoop, model F006661RX. Short battery life combined with the additional weight of the ball means that each flight lasts only 10 minutes. The propellers are guarded, which allows the drone to be flown next to surfaces without the risk of damage. The drone is flown by a pilot with the basic certificate required under Canadian regulations for drone flight safety.

Image Processing

After gathering the image data, the ground truth chromaticity data is extracted as described in the following subsections.

Black Level Correction

Many cameras have a black level greater than zero. The black level of the Nikon D700 camera was measured by taking several exposures with the lens cap on and averaging the pixel values of the images. The black level for camera is effectively zero at 0.8 out of 4096.

Ball and Drone Detection

To measure the chromaticity of grey ball in an image, it needs to be accurately located, and given the large number of images an automatic method is needed. Simple background subtraction was tested but failed too often. As a result the $Mask\ R-CNN$ [22] method is used. This requires a large training and validation image dataset. A synthetic dataset of 800 images was created by inserting drone images on top of background images randomly chosen from the COCO dataset. Images of all 30 scenes without a grey ball in them were also included in the training set as background images.

The $Mask\ R-CNN$ is trained for 300 epochs, 100 steps in each iteration, with a batch size of 1. The training requires approximately 3 hours on an NVIDIA GeForce GTX 1080. Finally, the trained model is applied to the actual dataset. (Figure 4) shows a sample of the results in which the bounding box correctly includes the entire ball and drone.

The ball still needs to be isolated within the bounding box. The Hough-circle transform [23] applied to the image area inside the bounding box reliably detects as shown in Figure 4. To do so, a range of possible values for the radius of the circle is searched to find the circle with the highest confidence value.

Results

Dataset properties

The dataset contains images of 30 scenes, both indoor and outdoor, captured during the day in summer July and August in Vancouver. For each scene there are 150 images on average. The final version of the image dataset is approximately 60 GB. The images are freely available from (URL included here when published). They are in three formats including: NEF (Nikon RAW format), linear Tiff and non-linear PNG. The scenes are located at Simon Fraser University's Burnaby and Downtown campuses, and at Confederation Park. There are a variety of lighting conditions including direct sunlight, daylight in shadow, daylight in cloudy weather, artificial indoor lights, and reflections from colourful surfaces.

RAW images in the dataset are first demosaiced using DCRAW. Measurements of illumination chromaticities for each scene are provided in two ways. The first format is as a NumPy dictionary. NumPy [24] is a library for the Python programming language for processing multi-dimensional arrays. For each image, the data includes the coordinates of the detected grey ball, its radius, and the RAW RGB values from the grey ball in the camera's native RGB colour space. In all cases, the location of the grey ball was confirmed by visual inspection and adjusted manually when necessary.

The second format is as an image, an 'illumination map,' within which the illumination chromaticity information is embedded. The illumination map for a scene consists of a black ($R = G = B = 0$) background with the images of the detected grey balls placed on it. To account for the fact that the grey ball is not a perfect grey, each RGB in the map is divided by the computed RGB of the grey ball under the ideal white light. This value can be calculated by multiplying the camera sensitivity for each channel with the spectral reflectance of the grey ball, which are both already measured.

A 15x15 uniform averaging filter is applied to images of scenes prior to cropping out the grey ball so in order to average out the effects of the ball's slight surface roughness. Finally, the RGB in the illumination map are converted to rg-chromaticity values by normalizing each pixel with its $R+G+B$ value. The illumination map is stored as an 8-bit PNG image. An example is shown in the Figure 10.

Illumination Chromaticity Distribution

The goal of this project is to obtain the true chromaticity of the incident illumination at many locations in images of typical scenes. The question naturally arises as to how much variation in the chromaticity actually occurs. This is evaluated in terms of the angular difference. Each rg-chromaticity 2-tuple is converted to an rgb 3-tuple using $r+g+b = 1$. The angle between the given grey ball illumination chromaticity measurement and the chromaticity of the grey ball under ideal white illumination is then computed for each ball location in the scene.

The average angle in degrees, along with the corresponding standard deviation, is computed for each scene. Based on the results shown in 6 the significant standard deviation in the angles found for most scenes indicates that they contain a significant variation in illumination chromaticity.

The next issue is whether or not the scene illuminations represent a reasonable sample of range of illumination chromaticities that are likely to be found in real world scenes. Plots of the measured illumination chromaticites are shown in Figure 7, Figure 8 and Figure 9. Each illumination chromaticity found in a scene is represented by a dot on the plot. The first chromaticity plot (in camera rg-chromaticity space) corresponds to the scene "Uncle Fatih's Pizza" at SFU). The chromaticities are primarily located along the Planckian locus and cover a range of colour temperatures. Figures 8 and 9 separately plot the chromaticities for found in indoor and outdoor scenes.

Optimal Single-Illuminant Estimate

Another measure of the variation in the illumination across the scene is provided by analysing the optimal single-illuminant estimate of the illumination. This involves solving for the illuminant chromaticity that minimizes the angular error across all the ground truth illumination measurements. Hao et al. [19] describe the Oracle method for solving this optimization problem. The Oracle method is based on complete information about the ground truth illumination so it is not an illumination estimation method. However, it does tell us what the minimum error any possible single-illuminant method can obtain. It also provides a target for potential multi-illuminant methods to surpass. (Note to Reviewers: The final paper will include Oracle results. Unfortunately, due to COVID all university labs have been closed so we are unable to include the Oracle result at this time but it is straightforward to compute once we regain lab access.)

Discussion

Applying the Dataset

The drone flight path for each scene is specifically planned so that each distinct surface will be sampled by the drone flying near it at least once. Due to the limitations of the drone's battery life and the large size of the RAW images exhausting the camera's memory capacity, the number of captured frames for each scene is usually limited to 150. However, this is expected to be sufficient for making comparisons between illumination estimation methods. Note that the camera's position and orientation remains fixed during the drone's flight. An image of the scene without the drone is also saved.

Using this data set, an illumination estimation algorithm can be evaluated by applying its algorithm to the drone-free image and its estimate of the illumination can be compared to the ground truth measurements at the corresponding image locations. This method has the possible disadvantage that the grey ball only flies near the surfaces but is not actually on the surfaces. This is unlikely to make a significant difference but in an extreme case the ball could be fully lit while the surface directly behind it is in shadow (not the ball's shadow). However, there are no such cases in the current dataset.

An alternative is to predict the chromaticity of the grey ball itself in each of the separate images. So as long as the algorithm does not explicitly depend on finding the ball or specifically detecting grey areas then this method has the advantage that the ground truth measurement is precisely at the same 3D location as the ball in the image.

Future Work

Although the chromaticity plots in Figs. 7, 9 and 8 show that the dataset contains a wide variety of illumination chromaticities, expanding the dataset further to include many more scenes would be desirable, and would make it more useful for training Machine Learning methods. The range of illumination chromaticities could also be increased by recording scenes at sunrise/sunset, during different seasons, and different parts of the world.

Summary

In most case, the illumination varies in its chromaticity throughout a scene due to effects such as direct sunlight versus shadow, the filtering of the light by plant leaves, indoor light versus daylight through a window, lights of different correlated colour temperature within a room, and by direct light versus light reflected off coloured surfaces. Investigations focused on illumination estimation and colour constancy have been hamper by a lack of measured scene illumination data. This project demonstrates that using a drone to carry a grey ball calibration target around a scene while simultaneously photographing it from a fixed location is an effective method of collection information about the chromaticity of the incident illumination at many different locations. A database of 30 real world scenes with illumination measurements at 100 or more locations in each is provided for download. Analysis of the data shows that the database includes many different illumination chromaticities and that there is a significant variation in the illumination across most of the scenes.

Acknowledgments

We thank the proprietors of Subway, Uncle Fatih's Pizza and A&W Canada (all at locations adjoining the Simon Fraser University Burnaby Campus) for permission to fly the drone on their premises, and Image Engineering GmbH & Co. KG for

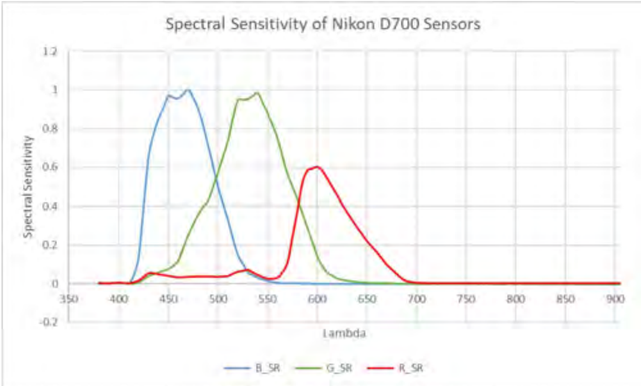

Figure 1. *The spectral sensitivity functions of the Nikon D700 used in these experiments were generously measured by Image Engineering GmbH & Co. KG using their camSPECS device*

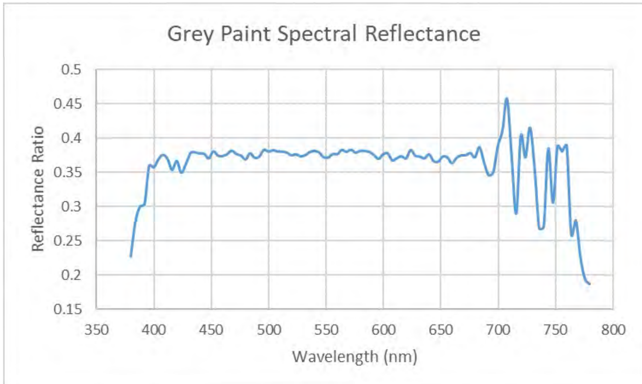

Figure 2. *The spectral reflectance of Munsell Neutral 5 grey paint as measured from the painted ball using a SpectraScan PR650*

Figure 3. *Right: CineBee 4K Whoop drone is used in the experiment to carry the grey ball around. The camera that comes with the drone is detached to reduce weight. Left: The drone carries the grey ball throughout each scene.*

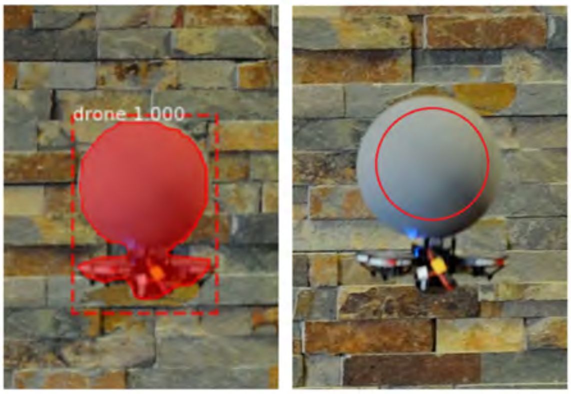

Figure 4. *Left: The trained model detects the drone in the scene and then determines a bounding box and mask for the ball attached to a drone. The mask area is shown in red. Right: The red circle represents the area used for illumination colour measurement. The circle is detected using the Hough-circle transform in the given bounding box.*

contributing the measurements of the camera's spectral sensitivity functions. Funding by the Natural Sciences and Engineering Research Council of Canada is also gratefully acknowledged.

References

[1] Gershon Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin institute*, 310(1): 1–26, 1980.

[2] Edwin H Land. The retinex theory of color vision. *Scientific american*, 237(6):108–129, 1977.

[3] Joost Van De Weijer and Theo Gevers. Color constancy based on the grey-edge hypothesis. In *IEEE International Conference on Image Processing 2005*, volume 2, pages II–722. IEEE, 2005.

[4] Graham D Finlayson and Elisabetta Trezzi. Shades of gray and colour constancy. In *Color and Imaging Conference*, volume 2004, pages 37–41. Society for Imaging Science and Technology, 2004.

[5] Jonathan T Barron and Yun-Ta Tsai. Fast fourier color constancy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 886–894, 2017.

[6] Michael Bleier, Christian Riess, Shida Beigpour, Eva Eibenberger, Elli Angelopoulou, Tobias Tröger, and André Kaup. Color constancy and non-uniform illumination: Can existing algorithms work? In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 774–781. IEEE, 2011.

[7] Ligeng Zhu and Brian Funt. Colorizing color images. *Electronic Imaging*, 2018(14):1–6, 2018.

[8] Shao-Bing Gao, Yan-Ze Ren, Ming Zhang, and Yong-Jie Li. Combining bottom-up and top-down visual mechanisms for color constancy under varying illumination. *IEEE Transactions on Image Processing*, 28(9):4387–4400, 2019.

[9] Simone Bianco, Claudio Cusano, and Raimondo Schettini. Single and multiple illuminant estimation using convolutional neural networks. *IEEE Transactions on Image Processing*, 26(9):4347–4362, 2017.

[10] Shida Beigpour, Christian Riess, Joost Van De Weijer, and Elli Angelopoulou. Multi-illuminant estimation with conditional random fields. *IEEE Transactions on Image Processing*, 23(1):83–96, 2013.

[11] Peter Vincent Gehler, Carsten Rother, Andrew Blake, Tom Minka, and Toby Sharp. Bayesian color constancy revisited. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.

[12] Lilong Shi and Brian Funt. Re-processed version of the gehler color constancy dataset of 568 images. *http://www.cs.sfu.ca/~color/data/*, 2000.

[13] Ghalia Hemrit, Graham David Finlayson, Arjan Gijsenij, Peter Vincent Gehler, Simone Bianco, Mark Drew, Brian Funt, and Lilong Shi. Providing a single ground-truth for illuminant estimation for the colorchecker dataset. *IEEE transactions on pattern analysis and machine intelligence*, 2019.

[14] Dongliang Cheng, Dilip K Prasad, and Michael S Brown. Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution. *JOSA A*, 31(5):1049–1058, 2014.

[15] Nikola Banić and Sven Lončarić. Unsupervised learning for color constancy. *arXiv preprint arXiv:1712.00436*, 2017.

[16] Florian Ciurea and Brian Funt. A large image database for color constancy research. In *Color and Imaging Conference*, volume 2003, pages 160–164. Society for Imaging Science and Technology, 2003.

[17] Color constancy the cube+ dataset. `https://ipg.fer.hr/ipg/resources/color_constancy`.

[18] Sean Bell, Kavita Bala, and Noah Snavely. Intrinsic images in the wild. *ACM Transactions on Graphics (TOG)*, 33(4): 159, 2014.

[19] Xiangpeng Hao, Brian Funt, and Hanxiao Jiang. Evaluating colour constancy on the new mist dataset of multi-illuminant scenes. In *Color and Imaging Conference*, volume 2019, pages 108–113. Society for Imaging Science and Technology, 2019.

[20] Arjan Gijsenij, Rui Lu, and Theo Gevers. Color constancy for multiple light sources. *IEEE Transactions on Image Processing*, 21(2):697–707, 2011.

[21] Sérgio MC Nascimento, Kinjiro Amano, and David H Foster. Spatial distributions of local illumination color in natural scenes. *Vision Research*, 120:39–44, 2016.

[22] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.

[23] HK Yuen, John Princen, John Illingworth, and Josef Kittler. Comparative study of hough transform methods for circle finding. *Image and vision computing*, 8(1):71–77, 1990.

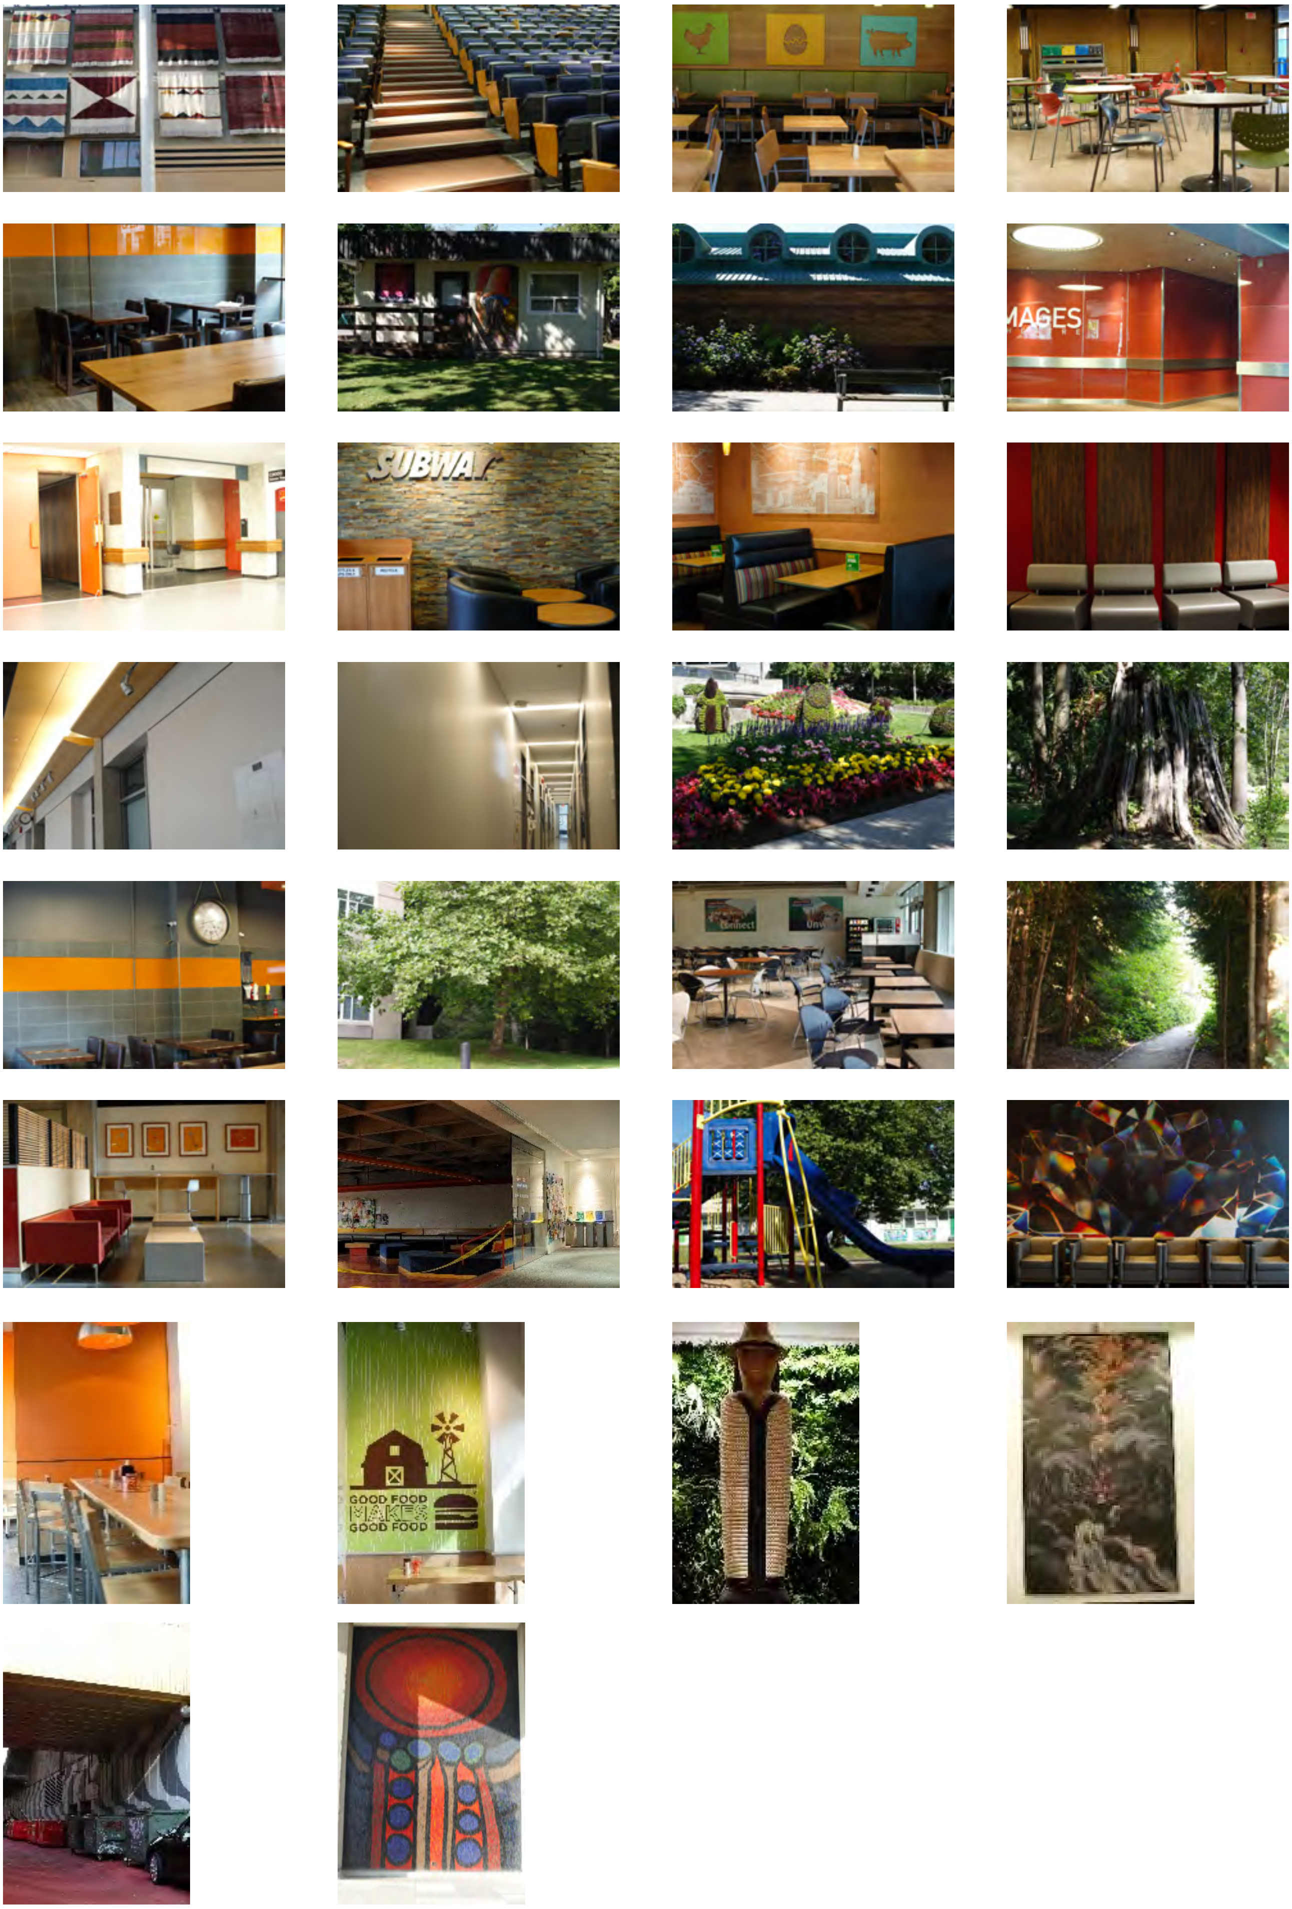

Figure 5. *Scenes of the image dataset (images are in nonlinear format for visualization purposes)*

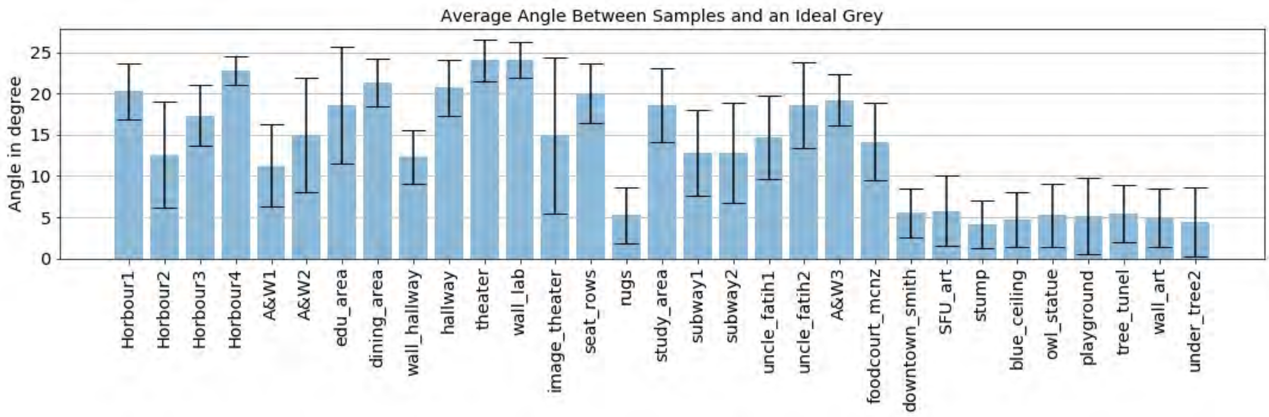

Figure 6. *The average of the angles between illumination chromaticity measurements and an ideal grey is computed for each scene. Each bar represents the average angle in degrees for the given scene. The black bars indicate the standard deviation.*

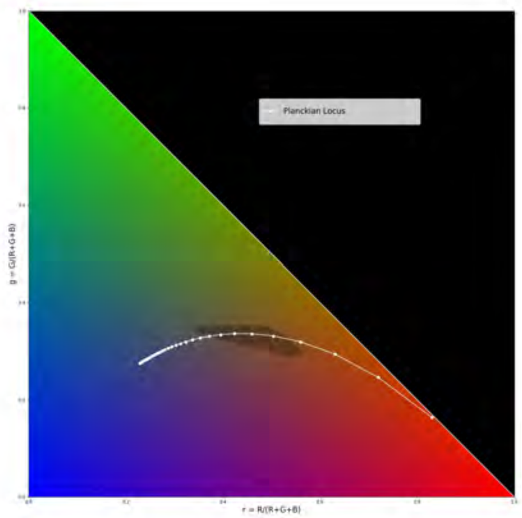

Figure 7. *Plot of the rg-chromaticity illumination measurements from a single scene (Uncle Fatih's Pizza at SFU). Each dark dot represents an illumination chromaticity measured from the grey ball at a different scene location.*

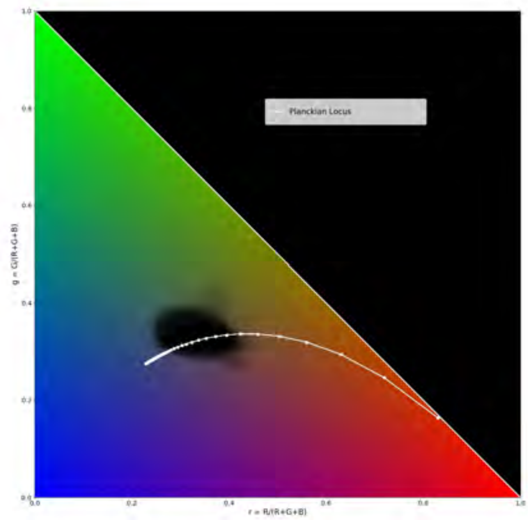

Figure 8. *Plot of the rg-chromaticity values for the illumination measurements from all the outdoor scenes in the dataset. Each dark dot represents an illumination chromaticity measured from the grey ball at a different scene location.*

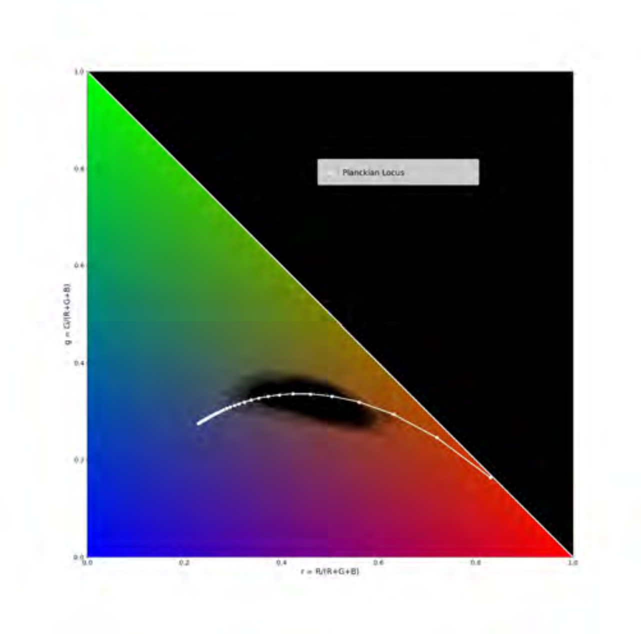

Figure 9. *Plot of the rg-chromaticity values for the illumination measurements from all the indoor scenes. Each dark dot represents an illumination chromaticity measured from the grey ball at a different scene location.*

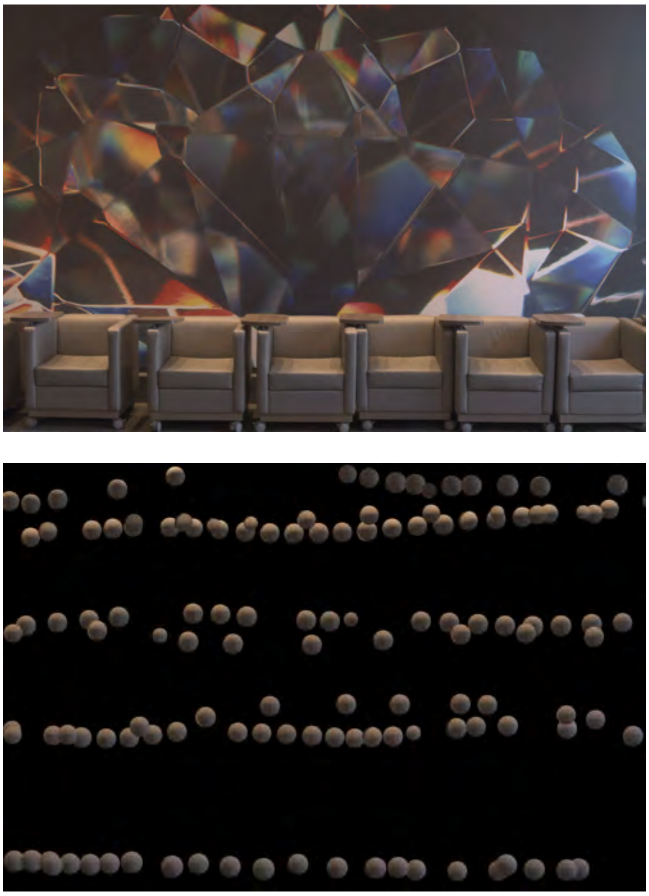

Figure 10. *Left: input image of a scene containing chairs in front of a painting in a hallway. Right: illumination map showing a superimposition of the chromaticities of all the detected grey balls on a black background. The normalized rg-chromaticity value of the grey ball in each frame of a scene is placed in the exact position that the grey ball is detected. Gamma correction is applied to the images here for visualization.*

[24] Travis E Oliphant. *A guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.