Supplementary Material: A computational model for color assimilation illusions and color constancy

Oguzhan Ulucan, Diclehan Ulucan, and Marc Ebner

University of Greifswald, Institute of Mathematics and Computer Science
17489 Greifswald, Germany
{oguzhan.ulucan,diclehan.ulucan,marc.ebner}@uni-greifswald.de

1 Overview

In this supplementary material, we provide the following:

- A visual investigation of computing the local space average color at each scale rather than computing it only on the finest scale and then carrying the estimations into scale-space
- The impact of our multiresolution color constancy strategy on color illusions
- Example images of the color illusion set that we utilize
- A brief introduction of the color constancy benchmarks and the error metrics
- Additional results for color illusions
- Visual results on color constancy benchmarks
- A simple illustration of the multiresolution color constancy strategy

2 Visual Investigation for the Scale-Space Computations

As we mentioned in our paper (Sec. 3.2), one can compute the local space average color only at the finest scale and then carry this estimation into scale-space rather than computing it at each scale separately. However, such an approach causes degradation in the local estimates, especially at the coarser scales. As shown in Fig. 1, the locality at the coarser scales is distorted when we carry the estimation at the finest scale into scale-space. However, this problem is mostly avoided when we compute the estimations at each scale which benefits not only color assimilation illusions but also color constancy. For instance, as demonstrated on the bottom row of Fig. 1, the local estimates at the coarsest scale are significantly distorted when we compute the local space average color only at the finest scale while details, *i.e.*, the sharp changes caused by the illumination falling onto the objects such as brighter regions and shadows, are greatly preserved when we compute the estimates at each scale separately.

3 The Impact of Our Strategy on Color Illusions

In our paper (Sec. 3.2), we mentioned that our proposed strategy allows us to mimic our sensation on color illusions, while also performing color constancy. As we stated, if we do not utilize our strategy and directly scale the pixel intensities of the input image

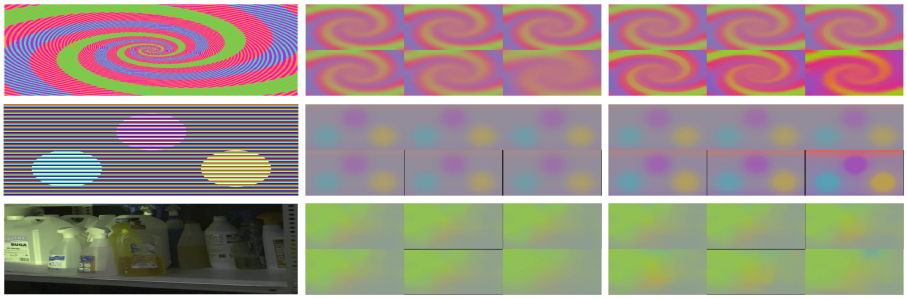

Fig. 1: A demonstration for the scale-space computations. (Left-to-right) Input image, estimation computed only at the finest scale and carried into scale-space, and estimations computed at each scale separately. In the estimations, the left-most image on the top row corresponds to the estimations at the finest scale, while the right-most image on the bottom row represents the estimations at the coarsest scale. The details are better preserved when the estimations are computed at each scale.

Fig. 2: Example visual comparison on color assimilation illusion by using multi-illuminant color constancy approaches and the proposed method. (Left-to-right) Input illusion, and input's target region, pixel-wise estimations, and estimated target region by GI, pixel-wise estimations, and the estimated target region by proposed approach. Since multi-illuminant color constancy algorithms, such as GI [16], focus on mimicking our ability of recognizing the (shaded) reflectance of the objects in the scenes with a color cast, they produce satisfying results for color constancy. Yet, their strategy in discounting the illuminant is not sufficient if we want to go one step further and mimic our sensation on illusions as well. Our approach can be deceived by the illusions while it also discounts the effects of the lighting conditions in the scenes with a color cast.

according to their estimates only on a single scale, we would be able to discount the effects of the light source and perform to color constancy, but we would not be able to reproduce our sensation on color illusions with a single pipeline. For instance, as shown in Fig. 2, a multi-illuminant algorithm, *i.e.*, GI [16], fails to produce satisfying results for color illusions, but it can present visually pleasing output images for color constancy (the color constancy results of GI are presented in Fig. 7). On the other hand, our proposed technique can produce outputs for both color assimilation illusions and color constancy.

4 Example Color Assimilation Illusions from our Set

While obtaining a set of color assimilation illusions, we consider illusions with various shapes of objects, inducer frequencies/thicknesses, and colors. As briefly stated in the main paper, the context has a crucial importance while controlling the illusion sensation.

Fig. 3: Example images [5] showing that the characteristics of the inducers affect our perception of color illusions. The images in the first two columns hardly evoke an illusion sensation due to their inducer's low frequency of occurrence and the inducer's large thickness. On the other hand, the images in the last two columns can easily deceive us due to their high inducer frequency and small thickness. We recommend zooming in on the image and fixate our focal vision on each image individually. We suggest the former since viewing the image from a distance can create the false impression of a strong illusion, while we suggest the latter since our blurred peripheral vision may amplify slight color illusions.

Hence, both the inducer's frequency and its thickness affect the sensation of illusions. As seen in Fig. 3, these two features are related to each other, *i.e.*, in order to increase the inducer's frequency of occurrence, we need to decrease its thickness. As shown in the first two columns of Fig. 3, the illusion sensation is very weak, while in the last two columns, the colors of the disks and stripes are severely affected by the context. We would like to stress that while investigating the color illusions in Fig. 3, it is important to zoom in on the image and fixate our focal vision on each image separately. We recommend the former since viewing the image as small-sized or from a distance will cause the image to appear as if the frequency of occurrence of the inducer is adequate to evoke an illusion sensation even though there is no (or a weak) illusion effect present in the image. We recommend the latter since our peripheral vision is blurred, *i.e.*, locally disordered [12] which may ease our perception of color illusions, e.g., even a slight illusion sensation can be perceived as a strong illusion effect due to the shift of the target region's colors towards that of its neighboring pixels.

In Fig. 4, we present several examples from the images that we consider during our investigation. We either created images ourselves or used images which are courtesy of Michael Bach, Akiyoshi Kitaoka, and David Novick.

5 Additional Results on Color Illusions

Due to the page constraint in the main paper, we demonstrated limited results on distinct color assimilation illusions. In this material, we provide further results in Fig. 5. We would like to note that all images are in sRGB format.

6 Utilized Color Constancy Benchmarks

We benchmarked our algorithm on the following datasets; the Multiple Illuminant and Multiple Object (MIMO) dataset [7], the Mixed-Illumination Test Set [3], the Rendered WB Dataset (Set 2) [4], and the rendered version of the Cube+ [4].

The MIMO dataset contains scenes captured under non-uniform illumination conditions and their ground truths, and it is one of the notable multi-illuminant datasets since numerous methods have already been tested on this benchmark [9]. The MIMO dataset

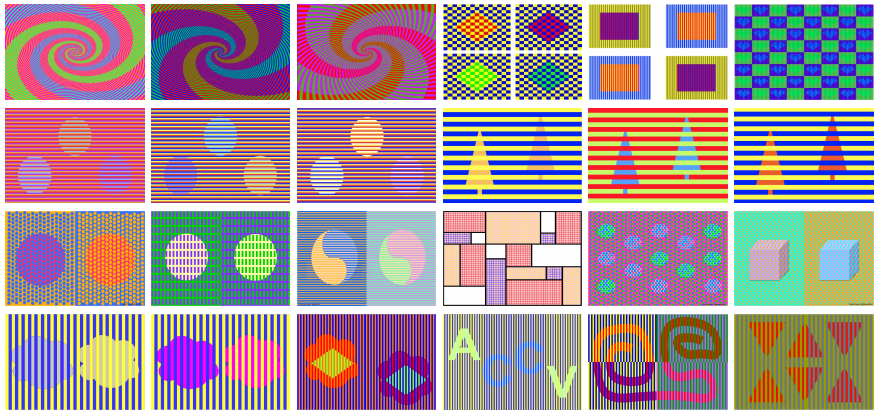

Fig. 4: Example color assimilation illusions with various colored shapes, and inducer frequencies/thicknesses from our set.

consists of two sets, namely, the Real-World set and the Laboratory set. While the former contains 20 images, the latter includes 50 scenes. The Mixed-Illumination Test Set consists of 30 different synthetic scenes, and each scene is rendered with 5 different color temperatures to generate a total of 150 images. The Rendered WB Dataset (Set 2) contains $2,881$ scenes with their corresponding ground truths. The images are captured with 4 different mobile phones and a DSLR camera. The rendered version of the Cube+ dataset is based on the linear raw images in the Cube+ dataset [6]. It contains $1,707$ unique scenes, and each scene is rendered with 6 different illumination conditions to obtain a total of $10,242$ rendered images. Example images for datasets can be seen in Fig. 6.

7 Error Metrics

In order to statistically investigate the performance of our method in the field of color constancy, we utilize well-known quality metrics, namely, the angular error and the $\triangle$E 2000 (CIEDE2000) [14,18]. The angular error (ε) between the ground truth illuminant (L_{gt}) and the estimated light source (L_{est}) can be computed as follows:

$$\varepsilon(L_{gt}, L_{est}) = cos^{-1}\left(\frac{L_{gt}L_{est}}{\|L_{gt}\|\|L_{est}\|}\right). \tag{1}$$

The $\triangle$E 2000 metric analyzes the color difference between two images, where a score closer to 0 corresponds to a better result [14,18]. Additionally, outcomes less than 1 indicate that the color difference is not perceivable to observers, while results in the range $[1, 4)$ might also be unnoticeable [11,15].

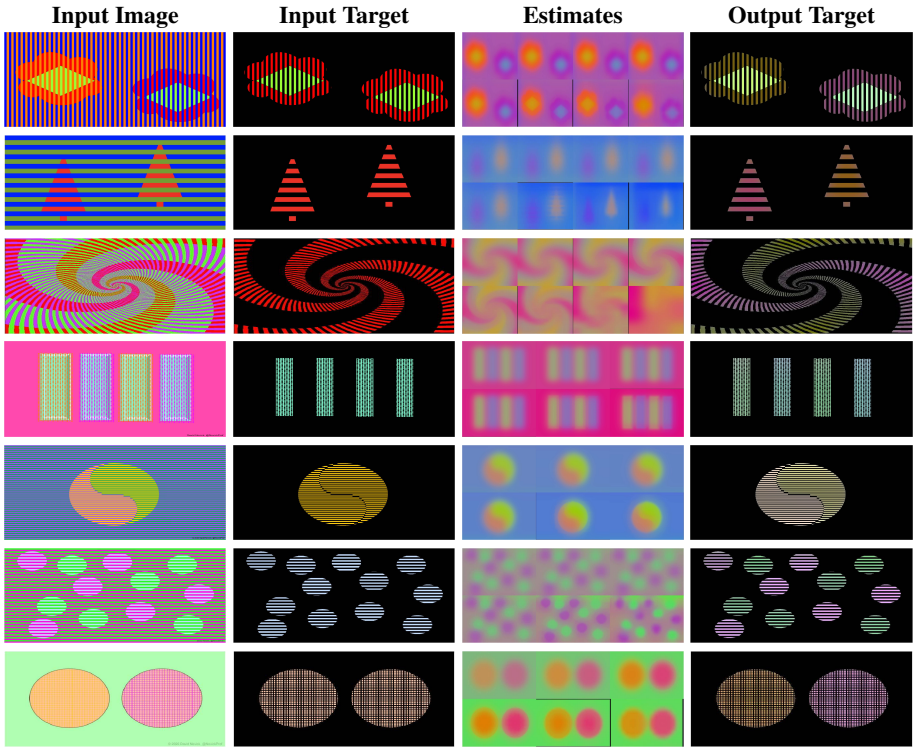

Fig. 5: Additional results of the proposed approach on color illusions. (Left-to-right) The input image, the input's target region, the estimations at each scale, and the target area of our output. The results of our algorithm are demonstrated in the output target. Our algorithm is deceived by various color illusions having different inducer frequencies/thicknesses, colors, and shapes. It is important to note that, the colors in the output target may seem darker than our perception in the input image due to the black background in the output target. The image in the first row is created by ourselves, while the remaining images are courtesy of Michael Bach, Akiyoshi Kitaoka, and David Novick, respectively.

Fig. 6: Random examples from the benchmarks. (Left-to-right) The MIMO dataset, the Mixed-Illumination Test Set, the Rendered WB Dataset (Set 2), and the rendered version of the Cube+ dataset.

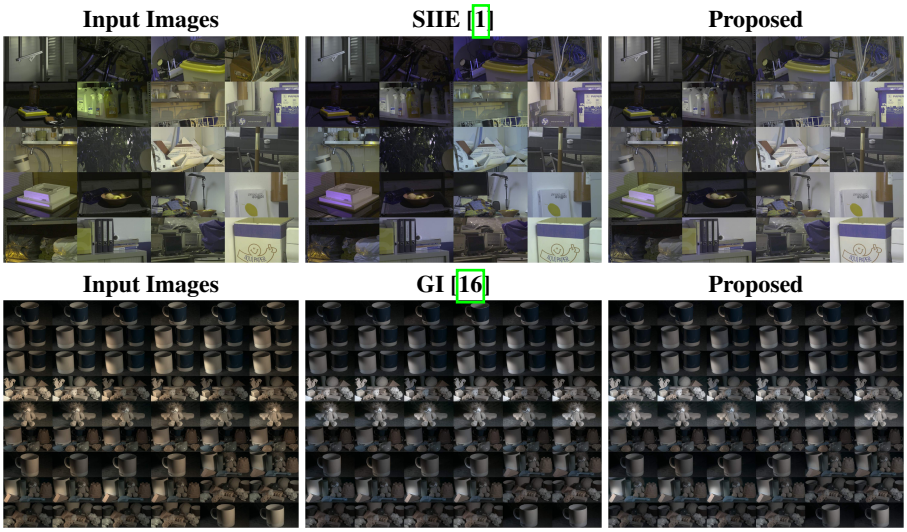

Fig. 7: Visual results of the MIMO dataset. (Left-to-right) Top row: The input images of the Real-World set, the results of SIEE [1], and our results. Bottom-row: The input images of the Laboratory set, the results of GI [16], and our results.

8 Additional Results on Color Constancy

In the main paper, due to the page limitations, we only reported statistical results on 4 different color constancy benchmarks. In this material, we demonstrate visual comparisons with several color constancy methods.

In Fig. 7, we provide visual comparisons for both the Real-World set and the Laboratory set of the MIMO dataset. On average, our algorithm presents the best statistical results on the Real-World set which is also supported by the visual outcomes. In a few of the images it is seen that specular regions and shadows slightly challenge our algorithm which is not surprising since it is known that such regions may be troublesome for color constancy methods. The image on the third row first column of the Real-World set in Fig. 7 can be given as an example for this case.

In Fig. 8, we provide visual results on the Mixed-Illumination Test Set [3], the Rendered WB Dataset (Set 2) [4], and the rendered version of the Cube+ [4]. Even under strong illumination conditions our algorithm provides visually pleasing outcomes.

9 Illustration of Multiresolution Color Constancy

We present a simple illustration of our approach in Fig. 9.

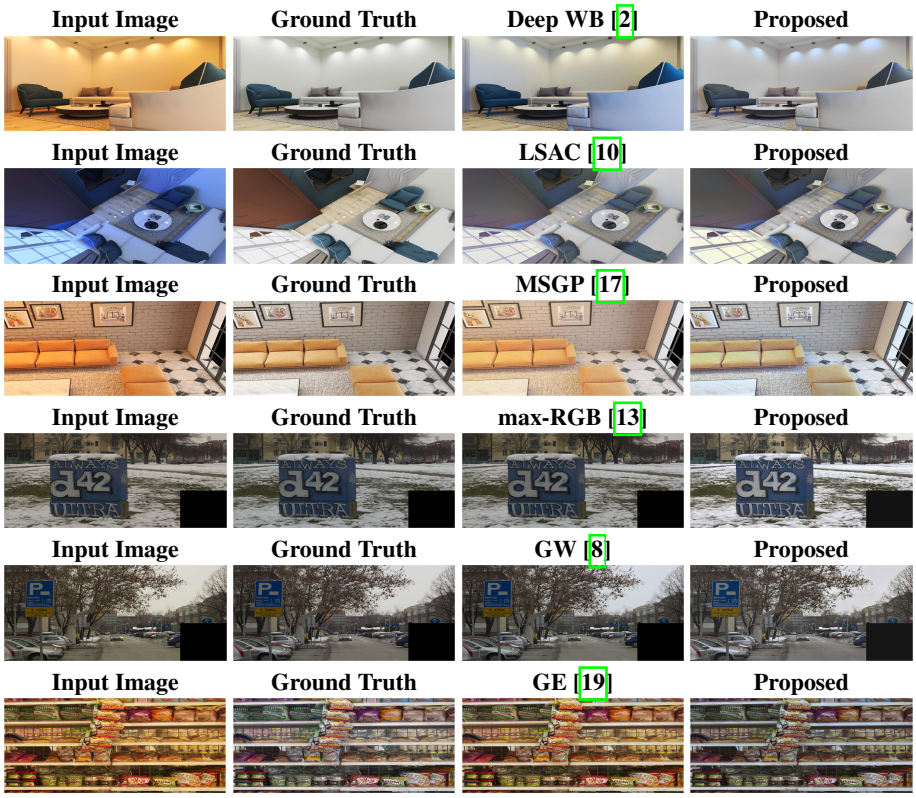

Fig. 8: Random visual results of indoor and outdoor scenes of different benchmarks [3, 4].

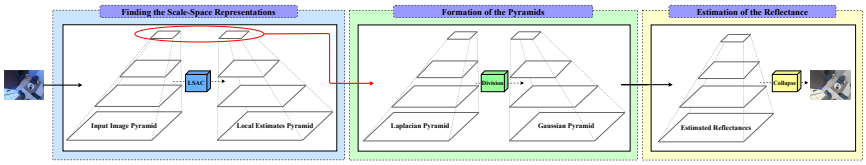

Fig. 9: A simple illustration of our multiresolution color constancy strategy. We carry the input image into scale-space where we compute its local estimates via the modified version of the local space average color method. Then, for each level of the pyramids, we form the Laplacian pyramid of the input image and the Gaussian pyramid of the local estimates. Subsequently, we estimate the (shaded) reflectance at each scale by dividing the Laplacian of the input image into the Gaussian of the local estimates. Afterwards, we collapse the resulting pyramid to obtain a reflectance estimate for the corresponding level. We repeat this procedure for all levels, and lastly collapse the resulting pyramid to obtain a single reflectance estimate for the input image.

References

1. Afifi, M., Brown, M.S.: Sensor-independent illumination estimation for DNN models. In: Brit. Mach. Vision Conf. BMVA Press (2019)

2. Afifi, M., Brown, M.S.: Deep white-balance editing. In: IEEE/CVF Conf. Comput. Vis. Pattern Recog. pp. 1397–1406 (2020)
3. Afifi, M., Brubaker, M.A., Brown, M.S.: Auto white-balance correction for mixed-illuminant scenes. In: IEEE/CVF Winter Conf. Appl. Comput. Vision. pp. 1210–1219 (2022)
4. Afifi, M., Price, B., Cohen, S., Brown, M.S.: When color constancy goes wrong: Correcting improperly white-balanced images. In: IEEE/CVF Conf. Comput. Vis. Pattern Recog. pp. 1535–1544 (2019)
5. Bach, M.: Color assimilation illusions (Last accessed: 03-10-2024), michaelbach.de/ot
6. Banić, N., Lončarić, S.: Unsupervised learning for color constancy. In: Int. Joint Conf. Comput. Vision Imag. Comput. Graph. Theory Appl. vol. 4, pp. 181–188. INSTICC (2018)
7. Beigpour, S., Riess, C., Van De Weijer, J., Angelopoulou, E.: Multi-illuminant estimation with conditional random fields. IEEE Trans. Image Process. **23**, 83–96 (2013)
8. Buchsbaum, G.: A spatial processor model for object colour perception. J. Franklin Inst. **310**, 1–26 (1980)
9. Buzzelli, M., Zini, S., Bianco, S., Ciocca, G., Schettini, R., Tchobanou, M.K.: Analysis of biases in automatic white balance datasets and methods. Color Res. Appl. **48**(1), 40–62 (2023)
10. Ebner, M.: A parallel algorithm for color constancy. J. Parallel Distrib. Comput. **64**, 79–88 (2004)
11. Ebner, M.: Color Constancy, 1st ed. Wiley Publishing, ISBN: 0470058299 (2007)
12. Koenderink, J.J., van Doorn, A.J.: Blur and disorder. J. Visual Communication Image Representation **11**(2), 237–244 (2000)
13. Land, E.H., McCann, J.J.: Lightness and retinex theory. J. Opt. Soc. America A **61**(1), 1–11 (1971)
14. Luo, M.R., Cui, G., Rigg, B.: The development of the CIE 2000 colour-difference formula: CIEDE2000. Color Res. Appl. **26**, 340–350 (2001)
15. Poynton, C.: Digital video and HD: Algorithms and Interfaces. Elsevier, – (2012)
16. Qian, Y., Kamarainen, J.K., Nikkanen, J., Matas, J.: On finding gray pixels. In: IEEE/CVF Conf. Comput. Vis. Pattern Recog. pp. 8062–8070 (2019)
17. Qian, Y., Pertuz, S., Nikkanen, J., Kämäräinen, J.K., Matas, J.: Revisiting gray pixel for statistical illumination estimation. In: Int. Joint Conf. Comput. Vision Imag. Comput. Graph. Theory Appl. vol. 4, pp. 36–46. INSTICC (2019)
18. Sharma, G., Wu, W., Dalal, E.N.: The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. Color Res. Appl. **30**, 21–30 (2005)
19. Van De Weijer, J., Gevers, T., Gijsenij, A.: Edge-based color constancy. IEEE Trans. Image Process. **16**, 2207–2214 (2007)