Research Article Vol. 37, No. 11 / November 2020 / Journal of the Optical Society of America A 1721

Journal of the Optical Society of America A
OPTICS, IMAGE SCIENCE, AND VISION

ARC: Angle-Retaining Chromaticity diagram for color constancy error analysis

MARCO BUZZELLI,* SIMONE BIANCO, AND RAIMONDO SCHETTINI

Department of Informatics, Systems, and Communication, University of Milano–Bicocca, 20126 Milan, Italy
**Corresponding author: marco.buzzelli@unimib.it*

Received 27 May 2020; revised 5 August 2020; accepted 13 September 2020; posted 14 September 2020 (Doc. ID 398692); published 8 October 2020

Color constancy algorithms are typically evaluated with a statistical analysis of the recovery angular error and the reproduction angular error between the estimated and ground truth illuminants. Such analysis provides information about only the magnitude of the errors, and not about their chromatic properties. We propose an Angle-Retaining Chromaticity diagram (ARC) for the visual analysis of the estimated illuminants and the corresponding errors. We provide both quantitative and qualitative proof of the superiority of ARC in preserving angular distances compared to other chromaticity diagrams, making it possible to quantify the reproduction and recovery errors in terms of Euclidean distances on a plane. We present two case studies for the application of the ARC diagram in the visualization of the ground truth illuminants of color constancy datasets, and the visual analysis of error distributions of color constancy algorithms.

https://doi.org/10.1364/JOSAA.398692

1. INTRODUCTION

Color constancy is the ability of the human visual system to perceive consistent object colors under different illumination conditions [1]. The exact mechanisms behind such phenomenon have been studied extensively [2,3], and human beings are generally assumed by vision scientists to have evolved this ability in assistance to object recognition [4].

Digital camera sensors do not naturally possess color constancy capabilities, and therefore *computational* color constancy algorithms are employed in the early stages of the image acquisition pipeline [5] to emulate human color constancy, and to eventually produce consistent object appearances under different illuminations. Computational color constancy, from now on referred to as "color constancy" for simplicity, is usually composed of two steps: the first one estimates the color of the scene illuminant from the analysis of image data; the second one corrects the image based on this estimate to generate a new rendition of the scene, as if it was taken under a fixed reference light source, such as standard daylight conditions. A widely adopted model for the color constancy mechanism is based on the von Kries transform [6], which describes the correction of the tristimulus values [e.g., red, green, and blue (RGB) for a digital image] with a diagonal matrix.

In order to evaluate the illuminant estimation part of computational color constancy, the error between a ground truth illuminant $U = (u_R, u_G, u_B)$ and an estimated illuminant $V = (v_R, v_G, v_B)$ can be computed with the recovery angular error [7,8] in RGB space,

$$\mathrm{err}_{\mathrm{rec}} = \arccos\left(\frac{U \cdot V}{|U||V|}\right) = \arccos\left(\frac{\sum_i u_i v_i}{\sqrt{\sum_i u_i^2}\sqrt{\sum_i v_i^2}}\right). \quad \textbf{(1)}$$

The reproduction angular error [9] has been proposed as an alternative evaluation of color constancy algorithms. It is computed by comparing a perfect white surface (1, 1, 1) with the reproduction of a white surface U (corresponding to the ground truth illuminant) corrected using the estimated illuminant V,

$$\mathrm{err}_{\mathrm{rep}} = \arccos\left(\frac{\frac{U}{V} \cdot (1, 1, 1)}{\left|\frac{U}{V}\right| \sqrt{3}}\right) = \arccos\left(\frac{\sum_i \frac{u_i}{v_i}}{\sqrt{\sum_i \frac{u_i^2}{v_i^2}}\sqrt{3}}\right). \quad \textbf{(2)}$$

Color constancy solutions, evaluated through either the recovery or the reproduction errors on a properly annotated dataset, are usually compared through simple summary statistics such as the mean, maximum, and median [10]. Alternatively, more advanced statistical tools can be employed, such as the Wilcoxon signed-rank test [11] as suggested by Hordley *et al.* [12], to compare in a pairwise fashion the whole error distributions of color constancy algorithms. The assessment through a single-valued metric such as the angular error is extremely valuable for a straightforward comparison of different methods, and we believe it should be maintained as the principal means of comparison due to its widespread adoption. By definition, however, angular metrics focus on the scale of the error, and ignore its direction (i.e., the chromaticity component of the error itself) [8]. We show with a concrete example in Section 6 that two methods for illuminant estimation that are considered statistically equivalent in terms of advanced analysis on the reproduction error can in fact hide significantly different error distributions, when the chromaticity information is considered. This suggests that whenever two or more solutions are found to be equivalent in a comparison based on the widely established angular error, further analysis can be conducted to reveal deeper insights.

On top of quantitative comparison, the visual inspection of error distributions can be exploited to highlight hidden characteristics of color constancy solutions, or to provide a more intuitive understanding of the error distribution itself. The output of color constancy methods is a 3D vector; nonetheless, illuminant information is normally analyzed in a bidimensional chromaticity diagram, since the magnitude of the illuminant is not relevant for color correction. Differently from what happens with error measures, for chromaticity diagrams there is not one universally accepted standard representation. We argue that the specific choice of chromaticity is highly critical, and show that the most commonly adopted solutions introduce distortions of the angular distances, which are responsible for unwanted biases in visual inspection of error distributions, as we will show through qualitative and quantitative analysis. We will also evaluate some alternative and less-used solutions, which present a better representation only for RGB triplets that are very close to the neutral axis, although we argue that highly chromatic lights are often more relevant in error analysis, and thus should not be neglected.

To overcome all these limitations, we define and present ARC: an Angle-Retaining Chromaticity diagram that introduces the least possible amount of distortions across the whole range of RGB values; in terms of angular distance preservation, we show that the RGB reproduction error directly maps into the Euclidean distance in ARC, and that the recovery error is highly preserved. We derive both the definition of RGB to ARC transformation, as well as the numerically stable inversion back into the RGB domain. We also present two applications of the proposed chromaticity diagram, which can be used to compare distributions of illuminant estimation datasets, as well as error distributions of existing methods for color constancy.

2. CHROMATICITY DIAGRAMS FOR COLOR CONSTANCY

In the framework of color constancy, different chromaticity diagrams have been employed either as a visualization tool, or as a working space for color-related processing.

Ratio chromaticity is the simplest solution for dimensionality reduction of RGB data. It consists in normalizing two of the components for a third one (for example, $\frac{R}{G}$ and $\frac{B}{G}$). As the normalization component approaches zero, however, the ratio quickly diverges toward infinity. This representation was used to describe the 2017 INTEL-TUT dataset for camera-invariant color constancy [13], and to build chromaticity histograms for illuminant estimation [14].

In order to partially constrain the co-domain of ratio chromaticity, the logarithm of the ratios has been introduced [15] (also known by the name of *uv chromaticity*, or log-chrominance [16]) and used as an intermediate data representation for color constancy algorithms [17,18], or for presentation purposes [19]. This formulation, however, heavily dilates the distances for input values that are close to the pure RGB colors, as we will show in Section 4.

A common solution is the *rg chromaticity* diagram, which divides two components (namely, R and G) by the sum of the three values. It has been employed as a working space for illuminant estimation [20,21], as a data visualization diagram [17,22,23], and to present annotated datasets Color Checker [24] and Cube+ [25] (in its rb variant). However, rg chromaticity is heavily skewed on the red–green axis, thus also introducing hue-specific distortions.

Alternative representations for chromaticity information are the so-called Maxwell triangle (i.e., the projection of RGB colors onto a plane perpendicular to the neutral axis), and the hue/saturation components from the hue, saturation, value (HSV) color space. These solutions are not commonly associated with the color constancy domain, although they do present valuable properties, as we will show through proper experimentation.

3. ANGLE-RETAINING CHROMATICITY DIAGRAM

The recovery angular error measures the 2D angle between two color vectors in 3D RGB color space. As such, it implies that all RGB vectors lying on the same ray from the origin are considered equivalent. We exploit the inherent low-dimensionality of the angular error to define a bidimensional chromaticity space, which we called Angle-Retaining Chromaticity (ARC), such that:

1. The angular distance between any RGB vector and the neutral axis is maintained as Euclidean distance in ARC [this distance corresponds to the reproduction error as defined in Eq. (2)].
2. The angular distance between any two RGB vectors is highly correlated to the Euclidean distance between the corresponding points in ARC [this distance corresponds to the recovery error as defined in Eq. (1)].

Following this definition, ARC is expressed in polar coordinates, such that:

- The **azimuth** (α_A) corresponds to the direction of the original RGB vector with respect to the neutral gray axis. This component is related to the *hue* of the color.
- The **radius** (α_R) corresponds to the angular error between the original RGB vector and the neutral gray axis. This component is loosely related to the *saturation* of the color.

The conversion for a generic RGB vector to the corresponding point in ARC diagram is depicted in Fig. 1. The Maxwell chromaticity triangle is also reported for reference.

Given an input RGB vector $P = (\rho_R, \rho_G, \rho_B)$, ARC polar coordinates $A = (\alpha_A, \alpha_R)$ are computed as follows:

$$\alpha_A = \text{arctan2}\left(\sqrt{3}(\rho_G - \rho_B), 2\rho_R - \rho_G - \rho_B\right), \quad \textbf{(3)}$$

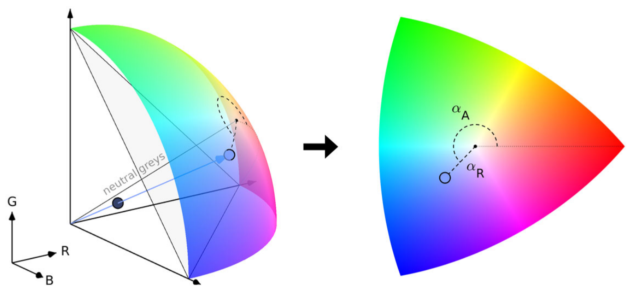

Fig. 1. Visualization of the conversion from a point in RGB color space to the corresponding point in ARC diagram.

$$\alpha_R = \arccos\left(\frac{\rho_R+\rho_G+\rho_B}{\sqrt{3}\sqrt{\rho_R^2+\rho_G^2+\rho_B^2}}\right). \tag{4}$$

With the first argument of arctan2 referring to the vertical axis. The polar coordinates system (α_A, α_R) can also be converted to Cartesian coordinates (α_X, α_Y), in order to facilitate their management, and the measurement of Euclidean distances,

$$\alpha_X = \frac{\arccos\left(\frac{\rho_R+\rho_G+\rho_B}{\sqrt{3}\sqrt{\rho_R^2+\rho_G^2+\rho_B^2}}\right)}{\sqrt{(2\rho_R-\rho_G-\rho_B)^2+3(\rho_G-\rho_B)^2}}(2\rho_R-\rho_G-\rho_B), \tag{5}$$

$$\alpha_Y = \frac{\arccos\left(\frac{\rho_R+\rho_G+\rho_B}{\sqrt{3}\sqrt{\rho_R^2+\rho_G^2+\rho_B^2}}\right)}{\sqrt{(2\rho_R-\rho_G-\rho_B)^2+3(\rho_G-\rho_B)^2}}\sqrt{3}(\rho_G-\rho_B). \tag{6}$$

When expressed in Cartesian coordinates (α_X, α_Y), the range of ARC is

$$-\arccos\left(\sqrt{\frac{2}{3}}\right) < \alpha_X < \arccos\left(\frac{1}{\sqrt{3}}\right), \tag{7}$$

$$-\frac{\sqrt{3}}{2}\arccos\left(\frac{1}{\sqrt{3}}\right) < \alpha_Y < \frac{\sqrt{3}}{2}\arccos\left(\frac{1}{\sqrt{3}}\right). \tag{8}$$

Conversely, any given ARC point $A = (\alpha_A, \alpha_R)$ expressed in polar coordinates can be converted in RGB space as a ray starting from the origin. Using either ρ_G or ρ_B as the independent variable (for numerical stability), the other components can be obtained as

$$\rho_R = \begin{cases} \frac{3\,\mathrm{sign}(c)\sqrt{(c^2-c+1)d}+(c^2-c-2)d-(c^2-c+1)}{(c^2+2c+1)d-(c^2-c+1)}\rho_G, & \text{if } \alpha_A > 0 \\ \frac{2\sqrt{d}+1}{1-\sqrt{d}}\rho_G, & \text{if } \alpha_A = 0 \\ \frac{-3\,\mathrm{sign}(c)\sqrt{(c^2-c+1)d}+(c^2-c-2)d-(c^2-c+1)}{-3\sqrt{c^2(c^2-c+1)d}-(2c^2+c-1)d-(c^2-c+1)}\rho_B, & \text{if } \alpha_A < 0 \end{cases}, \tag{9}$$

$$\rho_G = \begin{cases} \text{(independent variable)}, & \text{if } \alpha_A > 0 \\ \text{(independent variable)}, & \text{if } \alpha_A = 0 \\ \frac{(c^2+2c+1)d-(c^2-c+1)}{-3\sqrt{c^2(c^2-c+1)d}-(2c^2+c-1)d-(c^2-c+1)}\rho_B, & \text{if } \alpha_A < 0 \end{cases}, \tag{10}$$

$$\rho_B = \begin{cases} \frac{3\sqrt{c^2(c^2-c+1)d}-(2c^2+c-1)d-(c^2-c+1)}{(c^2+2c+1)d-(c^2-c+1)}\rho_G, & \text{if } \alpha_A > 0 \\ \rho_G, & \text{if } \alpha_A = 0 \\ \text{(independent variable)}, & \text{if } \alpha_A < 0 \end{cases}, \tag{11}$$

where c and d are computed, respectively, from α_A and α_R as

$$c = f_c(\alpha_A) = \frac{2\sqrt{3}}{\sqrt{3}-3\cot(\alpha_A)}, \tag{12}$$

$$d = f_d(\alpha_R) = \frac{\tan(\alpha_R)^2}{2}. \tag{13}$$

The full derivation of the conversion between RGB and ARC is provided in Section 5. An official implementation is made available at the project web page: http://www.ivl.disco.unimib.it/activities/arc/.

4. ARC GEOMETRIC INTERPRETATION AND PROPERTIES

The transformation between any given RGB point and its counterpart in ARC can be geometrically described by the following process:

1. We project the RGB point, following a ray from the origin, onto the surface of the octant of a sphere with center in (0, 0, 0) and radius $\frac{4}{\pi}\arccos(\frac{1}{\sqrt{3}})$.
   Fixing the radius of the sphere is important in order to provide a 1:1 scale between RGB angular error in degrees, and ARC Euclidean distance.
2. We flatten this surface by following an equidistant projection [26], which preserves the great-circle distances with respect to the point of neutral grays.

The reader may refer to Fig. 1 for a visual depiction of this procedure.

Plotting the entire RGB color space into ARC produces a gamut whose shape resembles a Reuleaux triangle [27]. Differently from a Reuleaux triangle, though, the boundary of the ARC gamut is not a curve of constant width. For the purpose of angular distance preservation, constant width would be a desirable property. It is, however, impossible to impose for all pairs of points at the same time: as a direct consequence of the *Theorema Egregium* [28,29], in fact, any mapping between a curved surface and a planar one, necessarily introduces distortions since the two surfaces are not respectively isomorphic. We choose to constrain the preservation of distances with respect to the white, in virtue of its relevance for the domain of color constancy. It follows that, given any general pair of RGB vectors U and V, their recovery angular error is not necessarily preserved in terms of Euclidean distance in ARC. However, thanks to the limited extent of the curved surface, and thanks to the full preservation of the distances with respect to the white, there is still a high correlation between angular distance in RGB and Euclidean distance in ARC. This correlation is, in fact, higher than what is observed with other commonly used chromaticity diagrams, as demonstrated with the following experiments. As a result, we effectively map 3D angular distances (errors) into segments on a plane. The segments' length can be visually

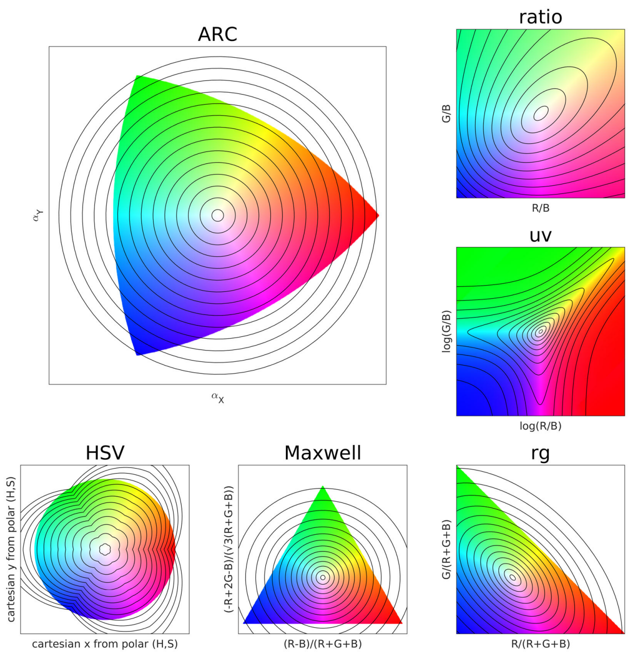

Fig. 2. Neighborhoods of the RGB neutral axis with increasing angles of 4° steps, mapped into different chromaticity diagrams. ARC correctly produces equidistant circles.

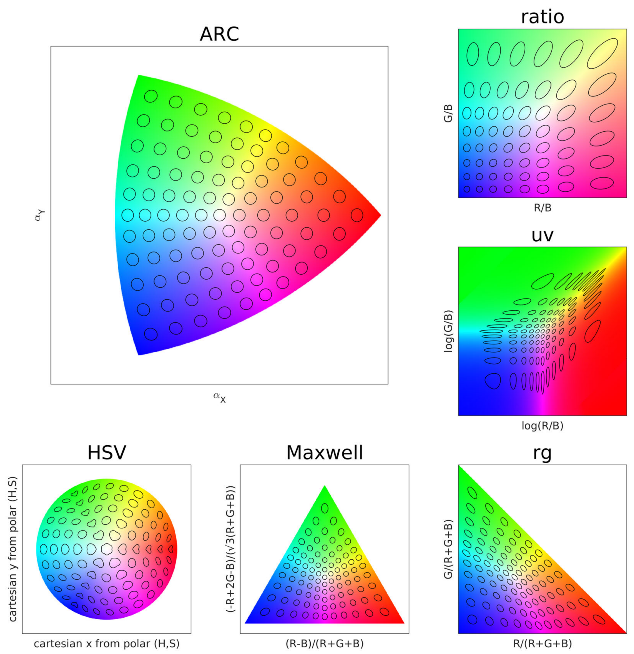

Fig. 3. Local RGB neighborhoods of 2° mapped into different chromaticity diagrams. ARC produces the best solution in terms of eccentricity and homogeneity of the neighborhoods.

appreciated and estimated by a human observer, as supported by Stevens's power law [30,31]. Furthermore, each segment is embedded in a hue-oriented space, which allows for a visual estimation of the error direction as well.

For the sake of completeness, in the following, we extend our evaluation of the diagrams described in Section 2 (ratio uv, and rg chromaticity) to the Maxwell color triangle, which is obtained by projecting the input values onto a plane perpendicular to the neutral axis, and to the hue/saturation pair from HSV brought into Cartesian coordinates. First, we analyze how the angular distance between any given RGB vector and the reference white (1, 1, 1) is affected by the transformation into chromaticity coordinates. We do so by defining neighborhoods of the neutral axis at different angles with a 4° step. The result is visually presented in Fig. 2.

A chromaticity diagram that preserves angular distances with a reference white would display equidistant concentric circles, while any type of distortion would result in circles being spaced unevenly, and/or in irregular shapes. By construction, ARC perfectly preserves the angular distance with respect to the neutral axis [corresponding to the (0, 0) "white" point in ARC]. As a consequence, it is possible to visualize the *reproduction* angular error between two illuminants U and V by computing their ratio $\frac{U}{V}$ as per Eq. (2), and plotting the resulting illuminant into ARC. Conversely, all other analyzed chromaticity diagrams introduce some form of distortion. Among these, the best alternative is the Maxwell triangle, which presents hue-invariant distortions that depend only on the distance from the center. All the remaining diagrams exhibit both hue-specific and saturation-specific distortions.

A similar experiment can be performed to investigate the local degree of distortion between pairs of colors. We sample 91 colors in RGB space, and define for each one a neighborhood of 2° diameter. We transform once again these neighborhoods into different chromaticity representations, and show the result in Fig. 3.

If the chromaticity transformation was to perfectly maintain the local angular distance as Euclidean distance, the resulting plot would display perfect circles of the same size. Conversely, different circle sizes correspond to regions in the chromaticity having higher (or lower) local angular distance with respect to other regions, while different shapes mean that specific directions in hue or saturation have a different impact on representing angular distances. ARC shows the best behavior in terms of local angle-retaining properties: the sizes of the neighborhoods are highly consistent, and the shapes get only slightly distorted for very saturated combinations of RGB triplets (corresponding to highly chromatic lights in the context of color constancy). Please note that the uneven distancing between neighborhoods is to be attributed to the sampling strategy, and not to a property of the chromaticity itself, as demonstrated in Fig. 2. Another well-performing diagram is the Maxwell triangle, which shows similar properties, but with a visibly stronger distortion for triplets close to the RGB axes. As a consequence, the *recovery* angular error between two illuminants U and V can be better approximated and visualized as the Euclidean distance between the corresponding points in ARC.

In order to offer a quantitative assessment of the angle-preservation capability of different chromaticity diagrams, we provide two evaluations. First, we compute fitting ellipses for the local neighborhoods of Fig. 3, and report in Table 1 the mean eccentricity and the coefficient of variation for the area. A lower eccentricity means that the neighborhoods are more similar to a circle, and a lower coefficient of variation for the areas means that the neighborhoods have a similar size.

Table 1. Fitting Ellipse Properties for Local Neighborhoods in Different Chromaticity Diagrams, the Lower the Better

Chromaticity	Eccentricity (mean)	Area (Coefficient of Variation)
ARC	**0.2516**	**0.0280**
ratio	0.7921	2.7600
uv	0.8864	1.0510
rg	0.8139	0.3092
Maxwell	0.4232	0.3091
HS (HSV)	0.7211	0.3864

We then compute the Pearson correlation coefficient [32] between the angle of color pairs randomly sampled in RGB, and the Euclidean distance of the corresponding chromaticities. This is shown in Table 2: both with respect to color pairs in which one is the reference white, and with arbitrary color pairs.

These quantitative evaluations empirically show the superiority of ARC with respect to other chromaticity diagrams. Conversely, the commonly used ratio chromaticity and uv chromaticity are shown to possess the least angle-retaining properties.

5. FULL DERIVATION OF THE TRANSFORMATION BETWEEN RGB AND ARC

In this section, we define the steps that lead to the transformation from an input RGB vector $P = (\rho_R, \rho_G, \rho_B)$ to ARC polar coordinates $A = (\alpha_A, \alpha_R)$.

A. RGB to ARC Transformation

The **radius** α_R [Eq. (4)] is computed as the recovery error, defined in Eq. (1), between the input RGB vector $P = (\rho_R, \rho_G, \rho_B)$ and a neutral such as $N = (1, 1, 1)$,

Table 2. Linear Correlation between Angular Error of Pairs of Points, and Euclidean Distance of the Corresponding Points in Different Chromaticity Diagrams, the Higher the Better

Chromaticity	Pair with White (Correlation)	Arbitrary Pairs (Correlation)
ARC	**1.0000**	**0.9996**
ratio	0.0157	0.0067
uv	0.7861	0.7291
rg	0.9200	0.9162
Maxwell	0.9922	0.9874
HS (HSV)	0.9531	0.9630

$$\alpha_R = \arccos\left(\frac{(\rho_R, \rho_G, \rho_B)\cdot(1, 1, 1)}{|(\rho_R, \rho_G, \rho_B)||(1, 1, 1)|}\right) = \arccos\left(\frac{\rho_R + \rho_G + \rho_B}{\sqrt{3}\sqrt{\rho_R^2 + \rho_G^2 + \rho_B^2}}\right). \quad (14)$$

The **azimuth** α_A [Eq. (3)] is defined as the angle formed by the projection of the input RGB vector P onto a plane perpendicular to the neutral gray axis. The reference line for this angle is given by the projection of the red hue (i.e., $\alpha_A = 0$), and the origin is the projection of the neutral axis itself. The azimuth component is obtained by first repositioning the reference frame with respect to the neutral axis. To do so, the initial step is to rotate the original RGB space with a transformation matrix M_R so that the white vector $N = (1, 1, 1)$ transforms into $N' = (0, 0, \sqrt{3})$ (i.e., the neutral axis becomes vertical in the new reference frame). We use Rodrigues' rotation formula [33] to rotate points around an arbitrary axis, defined through a unit vector ω, by an arbitrary amount θ,

$$M_R = \begin{bmatrix} \cos\theta + \omega_x^2(1-\cos\theta) & \omega_x\omega_y(1-\cos\theta) - \omega_z\sin\theta & \omega_x\omega_z(1-\cos\theta) + \omega_y\sin\theta \\ \omega_y\omega_x(1-\cos\theta) + \omega_z\sin\theta & \cos\theta + \omega_y^2(1-\cos\theta) & \omega_y\omega_z(1-\cos\theta) - \omega_x\sin\theta \\ \omega_z\omega_x(1-\cos\theta) - \omega_y\sin\theta & \omega_z\omega_y(1-\cos\theta) + \omega_x\sin\theta & \cos\theta + \omega_z^2(1-\cos\theta) \end{bmatrix}. \quad (15)$$

Let θ be the 2D angle in 3D space between the white N and N', following once again Eq. (1),

$$\theta = \arccos\left(\frac{(1, 1, 1)\cdot(0, 0, \sqrt{3})}{|(1, 1, 1)||(0, 0, \sqrt{3})|}\right) = \arccos\left(\frac{1}{\sqrt{3}}\right). \quad (16)$$

Let ω be defined as the point in coordinates $(1, -1, 0)$ normalized to have unitary norm, $(\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}}, 0)$. Equation (15), thus, becomes

$$M_R = \begin{bmatrix} \frac{1}{2} + \frac{1}{2\sqrt{3}} & \frac{1}{2\sqrt{3}} - \frac{1}{2} & -\frac{1}{\sqrt{3}} \\ \frac{1}{2\sqrt{3}} - \frac{1}{2} & \frac{1}{2} + \frac{1}{2\sqrt{3}} & -\frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} \end{bmatrix}. \quad (17)$$

We introduce a second rotation matrix M_P that operates around the new vertical axis (the neutral gray axis) by an angle $\gamma = 15°$, to eventually obtain the green vertex of ARC in the top-left corner of the chromaticity diagram, as it is common for chromaticity diagrams,

$$M_P = \begin{bmatrix} \cos\gamma & -\sin\gamma & 0 \\ \sin\gamma & \cos\gamma & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \frac{1+\sqrt{3}}{2\sqrt{2}} & -\frac{\sqrt{3}-1}{2\sqrt{2}} & 0 \\ \frac{\sqrt{3}-1}{2\sqrt{2}} & \frac{1+\sqrt{3}}{2\sqrt{2}} & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (18)$$

The RGB point P' in the new reference frame is, therefore,

$$\begin{bmatrix}\rho_X\\ \rho_Y\\ \rho_Z\end{bmatrix} = M_P M_R \begin{bmatrix}\rho_R\\ \rho_G\\ \rho_B\end{bmatrix} = \begin{bmatrix}\sqrt{\frac{2}{3}} & -\frac{1}{\sqrt{6}} & -\frac{1}{\sqrt{6}}\\ 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}}\\ \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}}\end{bmatrix}\begin{bmatrix}\rho_R\\ \rho_G\\ \rho_B\end{bmatrix} = \begin{bmatrix}\frac{2\rho_R-\rho_G-\rho_B}{\sqrt{6}}\\ \frac{\rho_G-\rho_B}{\sqrt{2}}\\ \frac{\rho_R+\rho_G+\rho_B}{\sqrt{3}}\end{bmatrix}. \quad (19)$$

The projection of the input RGB point P onto a plane perpendicular to the neutral axis (such as $z = 1$ in the new reference frame) implicitly eliminates one dimension. To this end, let l be the line passing through $O = (0, 0, 0)$ and $P' = (\rho_X, \rho_Y, \rho_Z)$, defined as $x = t\rho_X$, $y = t\rho_Y$, $z = t\rho_Z$. The intersection between line l and plane $z = 1$ leads to $t = \frac{1}{\rho_Z}$; therefore,

$$x = \frac{\rho_X}{\rho_Z} = \frac{\sqrt{2}(2\rho_R - \rho_G - \rho_B)}{2(\rho_R + \rho_G + \rho_B)}, \quad (20)$$

$$y = \frac{\rho_Y}{\rho_Z} = \frac{\sqrt{6}(\rho_G - \rho_B)}{2(\rho_R + \rho_G + \rho_B)}. \quad (21)$$

Finally, the planar angle formed by the ray passing through $(0, 0)$ and (x, y) with respect to a reference "horizontal" orientation is

$$\begin{aligned}\alpha_A &= \operatorname{arctan2}(y, x)\\ &= \operatorname{arctan2}\left(\frac{\sqrt{6}(\rho_G - \rho_B)}{2(\rho_R + \rho_G + \rho_B)}, \frac{\sqrt{2}(2\rho_R - \rho_G - \rho_B)}{2(\rho_R + \rho_G + \rho_B)}\right)\\ &= \operatorname{arctan2}\left(\sqrt{3}(\rho_G - \rho_B), 2\rho_R - \rho_G - \rho_B\right).\end{aligned} \quad (22)$$

B. ARC to RGB Transformation

Inverting Eq. (3) for the azimuth component produces a half-plane in the RGB space that hinges on the neutral axis, with an orientation depending on the value of α_A. Inverting Eq. (4) for the radius component results in an infinite cone with its vertex in the RGB origin, and its axis corresponding to the neutral axis. The aperture of this cone is directly related to the value of α_R. As the axes of these two surfaces coincide, their resulting intersection is a line in the RGB space, which corresponds to the whole set of RGB points that would produce the initial ARC point $A = (\alpha_A, \alpha_R)$. The whole procedure is illustrated in Fig. 4.

The α_A-originated half-plane must lay on the following plane:

$$\rho_B = c\rho_R + (1 - c)\rho_G, \quad (23)$$

where coefficients c and $(1 - c)$ ensure that the plane rotates around the equal-coordinate (neutral) axis. The precise relationship between c and α_A can be inferred by backtracking from Eqs. (22) to (19). According to Eq. (22),

$$x = y\cot(\alpha_A). \quad (24)$$

By definition of Eq. (19), then

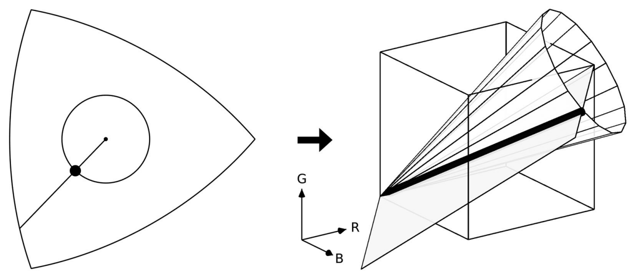

Fig. 4. Visual representation of the conversion between a point in ARC and the corresponding generating line of RGB points. This line is the intersection between an α_A-originated half-plane and an α_R-originated cone.

$$\begin{aligned}x &= \sqrt{\frac{2}{3}}\rho_R - \frac{1}{\sqrt{6}}\rho_G - \frac{1}{\sqrt{6}}\rho_B,\\ y &= \frac{1}{\sqrt{2}}\rho_G - \frac{1}{\sqrt{2}}\rho_B.\end{aligned} \quad (25)$$

This leads to

$$\sqrt{\frac{2}{3}}\rho_R - \frac{1}{\sqrt{6}}\rho_G - \frac{1}{\sqrt{6}}\rho_B = \left(\frac{1}{\sqrt{2}}\rho_G - \frac{1}{\sqrt{2}}\rho_B\right)\cot(\alpha_A). \quad (26)$$

Rearranging Eq. (26) in the explicit form of the plane [as in Eq. (23)], gives the definition of c as provided in Eq. (12). The precise half of the plane where the ARC point belongs (depicted in Fig. 4) depends on the sign of α_A, as defined in the following.

The α_R-originated cone must follow the equation

$$\rho_B^2 - \rho_G\rho_B - \rho_R\rho_B + \rho_G^2 - \rho_R\rho_G + \rho_R^2 = d(\rho_B + \rho_G + \rho_R)^2. \quad (27)$$

This parametrization, which ensures that the cone axis is aligned with the equal-coordinate (neutral RGB) axis, can be obtained by applying the Rodrigues rotation matrix from Eq. (17) to the general equation for a vertical axis infinite cone,

$$\rho_R^2 + \rho_G^2 = r^2\rho_B^2, \quad (28)$$

where $r = \sqrt{2d}$. The angle between the cone surface and its axis is equal to $\arctan(r)$, and, by the given definition of ARC, it is also equal to α_R,

$$\alpha_R = \arctan(r) = \arctan(\sqrt{2d}). \quad (29)$$

Solving Eq. (29) for d leads to Eq. (13).

The intersection between the half-plane given by Eq. (23) and the positive nappe of the infinite cone given by Eq. (27) leads to a ray in RGB space, starting from the origin, defined by the following equation:

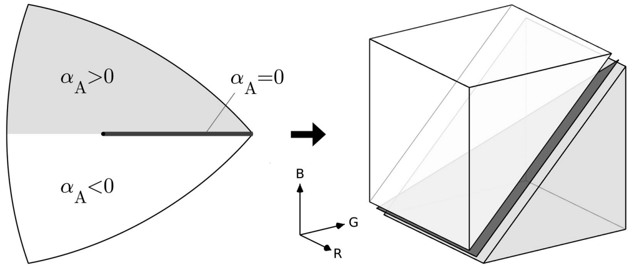

Fig. 5. Three different cases for ARC to RGB inversion that ensure numerical stability, as a function of the sign of α_A.

$$\frac{\rho_R}{3\mathrm{sign}(\alpha_A)\mathrm{sign}(c)\sqrt{(c^2-c+1)d}+(c^2-c-2)d-(c^2-c+1)}$$
$$=\frac{\rho_G}{(c^2+2c+1)d-(c^2-c+1)}$$
$$=\frac{\rho_B}{3\mathrm{sign}(\alpha_A)\mathrm{sign}(c)\sqrt{(c^2-c+1)d}-(2c^2+c-1)d-(c^2-c+1)}. \quad \textbf{(30)}$$

This relation is, once again, visualized in Fig. 4. In order to ensure numerical stability, when $\alpha_A > 0$, we express the RGB coordinates as a function of ρ_G, and when $\alpha_A < 0$, we express the RGB coordinates as a function of ρ_B. As can be observed in Fig. 5, this allows us to avoid defining the inversion as a function of a variable that may approach zero, as it happens with ρ_G when $\alpha_A < 0$, and with ρ_B when $\alpha_A > 0$.

A special case occurs when $\alpha_A = 0$. In this situation, the α_A-originated half-plane in Eq. (3) is described as $\rho_B = \rho_G$. Replacing ρ_B with ρ_G in Eq. (27) and solving for ρ_R gives the special case defined in Eq. (9) for $\alpha_A = 0$.

6. ARC CASE STUDIES

The visual comparison of two RGB illuminants (such as a ground truth triplet and an estimation triplet) can be performed in ARC in terms of approximate recovery error by plotting them individually and then considering their Euclidean distance. Alternatively, they can be compared in terms of exact reproduction error by plotting their ratio, and then considering the distance of the resulting point with the diagram origin.

ARC can also be exploited to visualize and compare entire color distributions, such as the set of ground truth illuminants in color constancy datasets, as well as the error distributions of algorithms for color constancy on a given dataset. This application is described in the following sections.

A. Distribution Analysis of Dataset Illuminants

In order to compare existing datasets on equal ground, it is necessary to consider the different spectral sensitivities of the sensors involved in their acquisition, i.e., to map the dataset illuminants into a device-independent color space, before converting them into chromaticity. Gao *et al.* [34] developed an approach to discount the camera spectral sensitivity in the estimation of illuminants, which relies on calibrated images and the availability of reflectance spectra. Banić *et al.* [25] proved that, under the assumption of similar content distributions, two datasets acquired with different sensors can be brought into comparable RGB color spaces through a pair of von Kries transformations [6], whose coefficients are computed as the median per channel of the illuminants in each dataset. In the following visualizations, we adopt the approach devised by Banić *et al.*, as we do not have access to spectral sensitivity functions of the involved cameras.

The distributions of three popular datasets (ColorChecker [35], NUS [36], and Cube+ [25]) are shown in Fig. 6, in conjunction with the CIE series D illuminants [37] from D40 to D150 as a guiding reference. The first dataset is the ColorChecker by Gehler *et al.* [35], which was acquired using two professional cameras (Canon EOS-1DS and CANON EOS 5D). It is composed of 568 images, each including a 24-patch Macbeth Color checker target for the ground truth estimation as recommended by Hemrit *et al.* [24]. The illuminant distribution of this dataset, depicted in Fig. 6(a), presents roughly three clusters of illuminants. Two clusters are in common between the two cameras: one gathered around the neutral lights, and one more shifted toward lower correlated color temperatures (CCT). A third cluster is exclusive to the Canon EOS-1DS camera, used to acquire images in scenes with higher CCT. These characteristics are less evident in the corresponding rg chromaticity diagram of Fig. 6(d), whose local distortions produce a compact, and thus less-discernible, visualization of the underlying distribution.

The second dataset is the National University of Singapore (NUS) by Cheng *et al.* [36], which is composed of 1853 images, shot with a total nine different cameras [reported in Fig. 6(b)], and also using a 24-patch Macbeth Color Checker target for ground truth estimation. In general, there is a high overlap among different cameras. With respect to the ColorChecker dataset, NUS presents a distribution that is more widely spread in a direction perpendicular to the CIE D illuminants. Furthermore, it reaches more into the red area of the chromaticity diagram, and less into the blue area. The two clusters are characterized by a different density, which can be clearly appreciated from the ARC representation of Fig. 6(b), and much less evident in the rg chromaticity diagram of Fig. 6(e).

The third dataset is the Cube+ by Banić *et al.* [25], which is an extension of the Cube dataset. It is composed of a total of 1707 images acquired with a Canon EOS 550D camera. The adopted color target is a SpyderCube calibration tool, with its two neutral 18% gray faces used to determine the ground truth

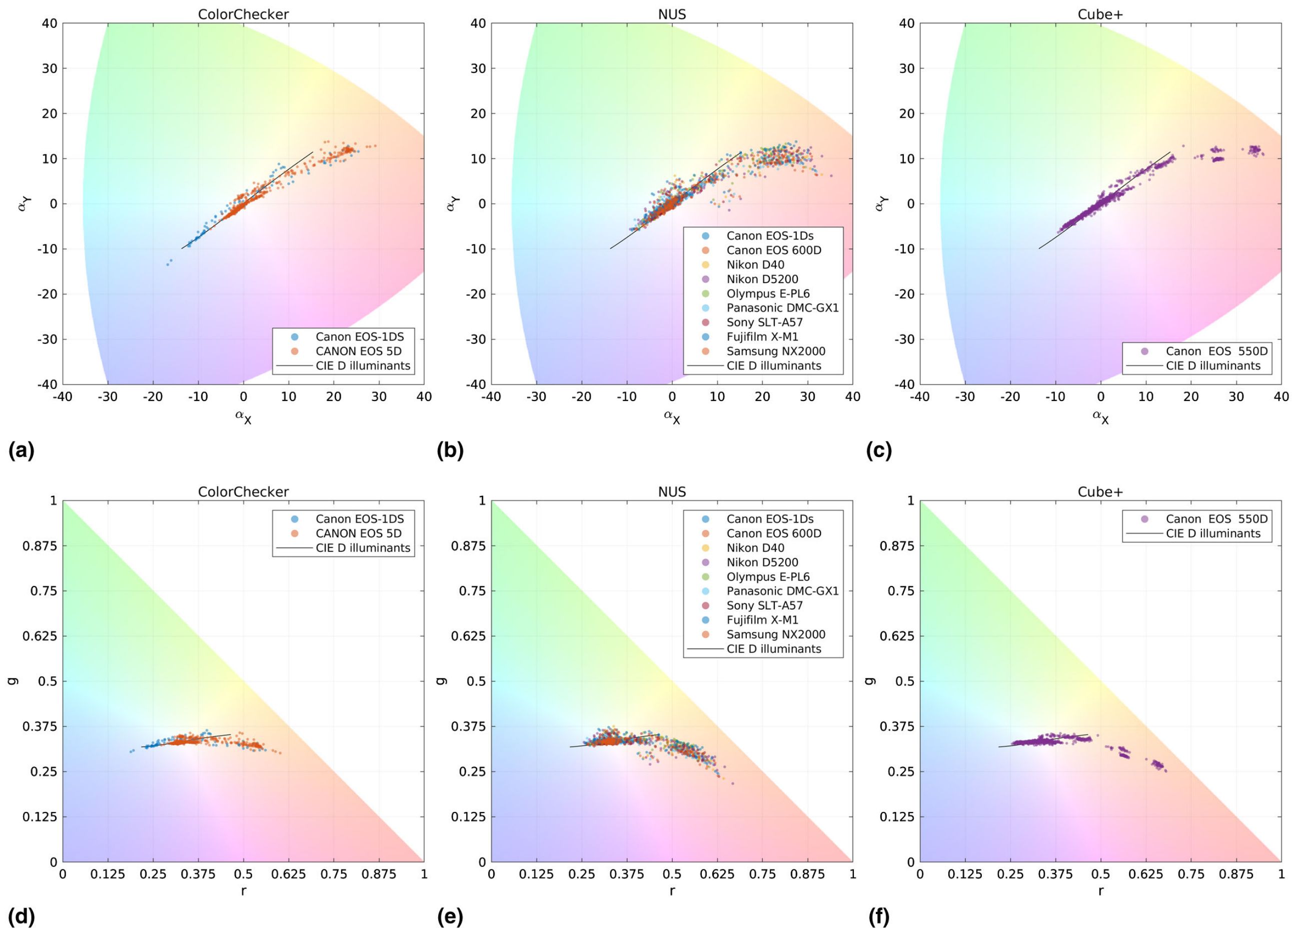

Fig. 6. Illuminant distributions for popular color constancy datasets ColorChecker [35], NUS [36], and Cube+ [25]. (a)–(c) present such distributions using our Angle-Retaining Chromaticity diagram ARC. (d)–(f) show the distributions using rg chromaticity, which is affected by both hue and saturation distortions.

for each image. Figure 6(c) shows the color distribution for both cube faces, and highlights four to five clusters of illuminants, where the cluster closest to the neutral point represents the set of images from the original Cube dataset. The angle-retaining properties of ARC ensure that the dataset distribution can be observed and compared, with little-to-no representational bias involved.

B. Error Distribution Analysis of Color Constancy Algorithms

The comparison of existing methods for color constancy is commonly performed through quantitative analysis, i.e., reporting error statistics, or using more advanced tools such as the Wilcoxon signed-rank test. A visual comparison of the error distributions can also contribute to discovering the hidden characteristics of each analyzed solution. Through ARC, this can be obtained by computing the ratio between ground truth and estimated illuminant on a given benchmark dataset, and projecting the corresponding points in ARC. The resulting distributions provide a visual insight into the reproduction angular error, by considering the proximity of points with respect to the diagram center. Euclidean distances in ARC, in fact, directly map to RGB angular distances in degrees.

As a use case for data visualization and comparison of color constancy solutions, we purposely selected two methods that provide statistically equivalent performance on the Color Checker dataset in terms of the Wilcoxon test: the Bayesian Color Constancy (BCC) [35] and the General Grey-World (GGW) [23]. The simple statistics reported in Table 3, in fact, show comparable performance, having extremely similar values with the exception of the highly sensitive maximum error.

This is reinforced by the Wilcoxon signed-rank test, which reports statistical equivalence with $p-$ value >0.17. Following the proposed visualization procedure, then, the ground-truth-to-estimation ratio is converted into ARC. The resulting distributions are shown in Fig. 7(a), filtered through a bivariate Gaussian kernel density estimation (KDE) [38]. For reference, we also report the illuminants D40 to D150 from series D, as well as the blackbody radiation with a CCT from 3500 K to 15,000 K.

Coordinates α_X and α_Y are gathered into projection histograms to facilitate the comparison, and the same is done for polar coordinates α_A and α_R. The reported diagram effectively presents some hidden characteristics of the two distributions that cannot be captured through global statistics: GGW exhibits a more isotropic distribution, compared to the skewed results

Table 3. Reproduction Angular Error Statistics on the ColorChecker Dataset [35] for Two Statistically Equivalent Methods

Method	Minimum	Mean	Median	90th p.tile	95th p.tile	Maximum
BCC [35]	0.073°	5.627°	3.921°	12.516°	15.393°	29.298°
GGW [23]	0.068°	5.317°	3.981°	11.159°	14.021°	23.310°

of BCC. In particular, it can be observed how the reproduction error distribution of BCC roughly follows the reference curve of blackbody radiation. Furthermore, the two methods display a different hue-specific bias, with BCC being more spread toward the magenta region of the diagram, and GGW toward the opposite end.

For comparison, we report in Fig. 7(b) the same analysis conducted on the commonly adopted rg chromaticity diagram. In this case, the joint combination of diagram-specific distortions, and the intrinsic distribution of the estimations of BCC, hinders the visibility of the differences between the two distributions and distorts the appearance of hue-specific biases. Misinterpretation is, therefore, more likely to take place in rg-chromaticity, especially if no frame of reference is provided as a support for visualization, such as the concentric equal-angle neighborhoods of Fig. 2.

7. CONCLUSION

Computational color constancy is an active field of research, where methods are typically compared via numerical analysis. We suggest that visual inspection through chromaticity diagrams can also be insightful; however, commonly used diagrams introduce hue-specific and saturation-specific distortions of angular distances between RGB points, making them unreliable tools for visual assessment of the data distributions. We presented ARC: an angle-retaining chromaticity diagram that maintains angular distances with the neutral gray axis as Euclidean distances with the diagram origin. As a consequence, it can be used to faithfully represent the reproduction angular error, and to accurately approximate the recovery angular error. We defined and derived the transformation of any RGB triplet into ARC, as well as its numerically stable inversion back into the RGB color space. Finally, we showed two practical applications of the proposed diagram: to visualize dataset distributions, and to highlight the different behaviors of statistically equivalent methods for color constancy.

Concerning future developments, we plan to study the relationship between error distributions in ARC, and human-perceived quality. By construction, the ARC diagram is tightly related to the recovery angular error. As a consequence, it inherits the corresponding properties as to what can be considered "perceivable" and "acceptable" in terms of color reproduction error. These concepts have been analyzed in a literature survey by Gijsenij *et al.* [8]: a deviation of 1° in angular error with the

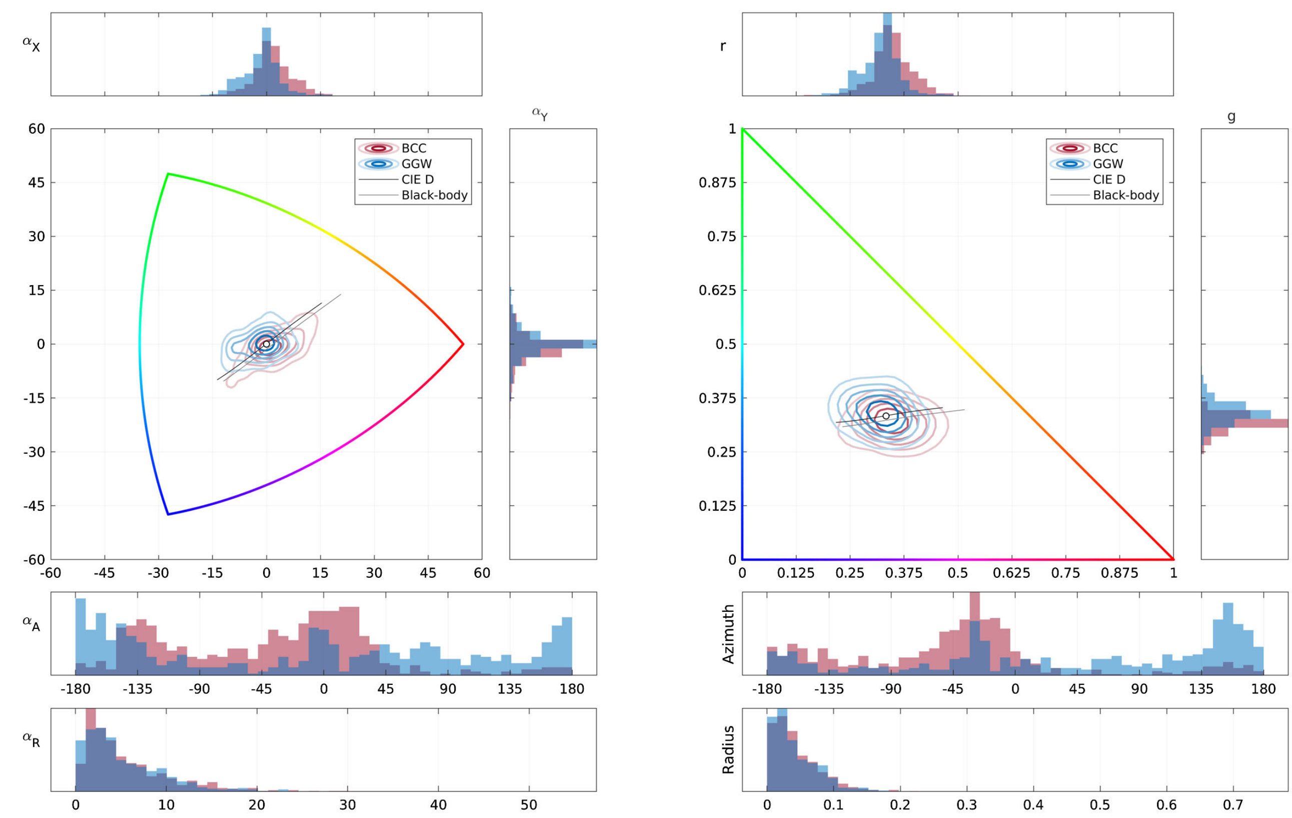

Fig. 7. Error distributions for color constancy methods BCC [35] and GGW [23]. The comparison is conducted through (a) the proposed Angle-Retaining Chromaticity, and (b) the commonly adopted rg chromaticity. Although the two algorithms are statistically equivalent, visualization with a chromaticity diagram such as ARC clearly highlights the difference in chromaticity patterns.

reference illuminant is considered below the threshold of what can be detected by a human being [39], while the range between 2° and 3° is considered perceivable but acceptable [40,41]. Other sources [42] identify a 2° recovery angular error as being acceptable for color constancy in complex scenes. Further analysis could be conducted by focusing on the chromaticity of the error itself, by relying on dedicated psychophysical experiments.

Finally, we will also consider expanding ARC from a pure visualization diagram to a working space representation for illuminant estimation and other forms of color-related processing.

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. D. H. Foster, "Does colour constancy exist?" Trends Cogn. Sci. **7**, 439–443 (2003).
2. J. J. McCann, "Do humans discount the illuminant?" Proc. SPIE **5666**, 9–16 (2005).
3. J. J. Granzier, E. Brenner, and J. B. Smeets, "Can illumination estimates provide the basis for color constancy?" J. Vision **9**, 18 (2009).
4. A. Hurlbert, "Colour constancy," Curr. Biol. **17**, R906–R907 (2007).
5. R. Ramanath, W. E. Snyder, Y. Yoo, and M. S. Drew, "Color image processing pipeline," IEEE Signal Process. Mag. **22**, 34–43 (2005).
6. J. von Kries, "Theoretische studien über die umstimmung des sehorgans," in *Festschrift der Albrecht-Ludwigs-Universität* (1902), pp. 145–158.
7. S. D. Hordley and G. D. Finlayson, "Reevaluation of color constancy algorithm performance," J. Opt. Soc. Am. A **23**, 1008–1020 (2006).
8. A. Gijsenij, T. Gevers, and M. P. Lucassen, "Perceptual analysis of distance measures for color constancy algorithms," J. Opt. Soc. Am. A **26**, 2243–2256 (2009).
9. G. D. Finlayson and R. Zakizadeh, "Reproduction angular error: an improved performance metric for illuminant estimation," Perception **310**, 1–26 (2014).
10. M. Buzzelli, J. van de Weijer, and R. Schettini, "Learning illuminant estimation from object recognition," in *25th IEEE International Conference on Image Processing (ICIP)* (IEEE, 2018), pp. 3234–3238.
11. F. Wilcoxon, "Individual comparisons by ranking methods," in *Breakthroughs in Statistics* (Springer, 1992), pp. 196–202.
12. S. D. Hordley and G. D. Finlayson, "Re-evaluating colour constancy algorithms," in *17th International Conference on Pattern Recognition, 2004 (ICPR 2004)* (IEEE, 2004), Vol. **1**, pp. 76–79.
13. C. Aytekin, J. Nikkanen, and M. Gabbouj, "A data set for camera-independent color constancy," IEEE Trans. Image Process. **27**, 530–544 (2018).
14. G. D. Finlayson, S. D. Hordley, and P. M. Hubel, "Color by correlation: a simple, unifying framework for color constancy," IEEE Trans. Pattern Anal. Mach. Intell. **23**, 1209–1221 (2001).
15. G. D. Finlayson, M. S. Drew, and C. Lu, "Intrinsic images by entropy minimization," in *European Conference on Computer Vision* (Springer, 2004), pp. 582–595.
16. J. T. Barron, "Convolutional color constancy," in *IEEE International Conference on Computer Vision* (2015), pp. 379–387.
17. M. Afifi and M. S. Brown, "Sensor-independent illumination estimation for DNN models," in *30th British Machine Vision Conference 2019* (BMVA Press, 2019), p. 282.
18. M. Afifi, B. Price, S. Cohen, and M. S. Brown, "When color constancy goes wrong: correcting improperly white-balanced images," in *IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 1535–1544.
19. W. Shi, C. C. Loy, and X. Tang, "Deep specialized network for illuminant estimation," in *European Conference on Computer Vision* (Springer, 2016), pp. 371–387.
20. V. C. Cardei, B. Funt, and K. Barnard, "Estimating the scene illumination chromaticity by using a neural network," J. Opt. Soc. Am. A **19**, 2374–2386 (2002).
21. D. Cheng, B. Price, S. Cohen, and M. S. Brown, "Effective learning-based illuminant estimation using simple features," in *IEEE Conference on Computer Vision and Pattern Recognition* (2015), pp. 1000–1008.
22. S.-B. Gao, M. Zhang, and Y.-J. Li, "Improving color constancy by selecting suitable set of training images," Opt. Express **27**, 25611–25633 (2019).
23. K. Barnard, V. Cardei, and B. Funt, "A comparison of computational color constancy algorithms. I: Methodology and experiments with synthesized data," IEEE Trans. Image Process. **11**, 972–984 (2002).
24. G. Hemrit, G. D. Finlayson, A. Gijsenij, P. V. Gehler, S. Bianco, M. Drew, B. Funt, and L. Shi, "Providing a single ground-truth for illuminant estimation for the colorchecker dataset," IEEE Trans. Pattern Anal. Mach. Intell. **42**, 1286–1287. (2019).
25. N. Banić and S. Lončarić, "Unsupervised learning for color constancy," in *Proceedings of the 13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications* (SciTePress, 2018), pp. 181–188.
26. J. P. Snyder and P. M. Voxland, *An Album of Map Projections, 1453* (US Government Printing Office, 1989).
27. C. Alsina and R. B. Nelsen, *Icons of Mathematics: An Exploration of Twenty Key Images, Volume 45 of Dolciani Mathematical Expositions* (MAA, 2011).
28. C. F. Gauss, *Disquisitiones Generales Circa Superficies Curvas* (Typis Dieterichianis, 1828), Vol. **1**.
29. A. Pressley, "'Gauss' theorema egregium," in *Elementary Differential Geometry* (Springer, 2010), pp. 247–268.
30. S. S. Stevens, "On the psychophysical law," Psychol. Rev. **64**, 153 (1957).
31. J. Bertin, *Sémiologie Graphique. Les diagrammes, les réseaux, les cartes* (Gauthier-Villars, 1967).
32. A. Bravais, *Analyse mathématique sur les probabilités des erreurs de situation d'un point* (Impr. Royale, 1844).
33. O. Rodrigues, "Des lois geometriques qui regissent les deplacements d'un systeme solide dans l'espace et de la variation des coordonnees provenant de deplacements consideres independamment des causes qui peuvent les produire," J. Mathematiques Pures Appliquees **5**, 380–440 (1840).
34. S.-B. Gao, M. Zhang, C.-Y. Li, and Y.-J. Li, "Improving color constancy by discounting the variation of camera spectral sensitivity," J. Opt. Soc. Am. A **34**, 1448–1462 (2017).
35. P. V. Gehler, C. Rother, A. Blake, T. Minka, and T. Sharp, "Bayesian color constancy revisited," in *IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2008), pp. 1–8.
36. D. Cheng, D. K. Prasad, and M. S. Brown, "Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution," J. Opt. Soc. Am. A **31**, 1049–1058 (2014).
37. D. B. Judd, D. L. MacAdam, G. Wyszecki, H. Budde, H. Condit, S. Henderson, and J. Simonds, "Spectral distribution of typical daylight as a function of correlated color temperature," J. Opt. Soc. Am. **54**, 1031–1040 (1964).
38. M. Rosenblatt, "Remarks on some nonparametric estimates of a density function," in *The Annals of Mathematical Statistics* (1956), pp. 832–837.
39. B. Funt, K. Barnard, and L. Martin, "Is machine colour constancy good enough?" in *European Conference on Computer Vision* (Springer, 1998), pp. 445–459.
40. G. D. Finlayson, S. D. Hordley, and P. Morovic, "Colour constancy using the chromagenic constraint," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)* (IEEE, 2005), Vol. **1**, pp. 1079–1086.
41. C. Fredembach and G. Finlayson, "Bright chromagenic algorithm for illuminant estimation," J. Imag. Sci. Technol. **52**, 40906 (2008).
42. S. D. Hordley, "Scene illuminant estimation: past, present, and future," Color Res. Appl. **31**, 303–314 (2006).