The Color Parrot White Balance Tool, Version 1.2

Laboratory Tests

Douglas A. Kerr

Issue 1
April 3, 2008

ABSTRACT

In this article we report the results of laboratory tests performed on a sample Color Parrot white balance tool, version 1.2. The significance of the various measured parameters is discussed, the test procedures are described in detail, and the results are presented.

SUMMARY

The Color Parrot white balance tool is a *measurement diffuser* intended to be fitted to a camera to allow the camera to make measurements of the chromaticity of incident light for use in guiding white balance color correction of photographic images. We made qualitative tests of key technical properties of the latest design of the Color Parrot, identified as Version 1.2 (although that notation is not publicized in the promotional literature).

Directivity pattern

Directivity pattern describes how components of the incident light arriving at the diffuser from different directions are "weighted" in the composition of the light presented by the diffuser to the camera for measurement of its chromaticity.

The Color Parrot 1.2 has a directivity pattern that is substantially "narrower" than, for example, the classical "cosine" pattern often approximated by white balance diffusers. The manufacturer of the Color Parrot states that this is intentional and is a key to the "superior performance" of the Color Parrot.[1]

Chromatic neutrality

Chromatic neutrality refers to how closely a measurement diffuser preserves the chromaticity of the incident light that is collected when passing it to the camera for measurement.

[1] We can't discern the principle involved here.

The Color Parrot 1.2 exhibited very good chromatic neutrality. Its departure from perfect neutrality would likely be insignificant in normal photographic white balance processing.

Effective total transmission

Effective total transmission can be simplistically thought of as "what fraction of the total light incident upon the face of the diffuser is passed through to the camera." For the Color Parrot 1.2, this is approximately 3 times that of the ExpoDisc, another commercially-available white balance diffuser with which the Color Parrot is often compared. The manufacturer of the Color Parrot states that this is advantageous in allowing the Color Parrot to be used for effective measurements in "low-light" contexts with cameras that would otherwise balk at doing so.

INTRODUCTION

White balance color correction

If a scene is photographed by a digital camera while illuminated by light whose chromaticity does not correspond closely to the "reference white chromaticity" of the color space in which the delivered image is recorded, then to a viewer of the delivered image the various objects will not seem to have their "natural" chromaticities; a "color cast" will have been imparted to the image.

If we know, can guess, or can measure the chromaticity of the incident light at the time and place of the photograph, we can apply a "white balance color correction" to the image to restore the expected appearance of objects. This can be done during the processing of the captured image data (especially when working with the "raw" camera output), or it can be done in the camera (and will have its effect on the delivered image in JPEG form).

The camera as a colorimeter

If we utilize the camera itself for measurement of the incident light chromaticity, then the result can be automatically noted by the camera and used without further ado to guide in-camera white balance color correction of the subsequent "shots".

Use of a neutral target

One way to do this is to have the camera, in advance of the actual shot, regard a *neutral reflective target* ("gray card") placed at the subject location. The light reflected from such a target, and analyzed by the camera, will have essentially the same chromaticity as the "effective overall chromaticity" of the incident light.

The white balance diffuser

Another technique is to temporarily fit the front of the camera's lens with a *white balance measurement diffuser*, essentially a translucent disk (with certain carefully-controlled properties). We temporarily place the camera, fitted with the diffuser, at the subject location, facing generally toward the camera position to be used for the actual shots.

The face of the diffuser is thus exposed to the same incident light illumination as will be the surface of the subject that is to be photographed. It presents on its rear face a luminous disk whose chromaticity is (hopefully) the same as the net effective chromaticity of the incident light. The lens (operating in a gravely out of focus mode) "images" that luminous disk onto the camera sensor. Then the camera determines the chromaticity of that light (perhaps over only a central portion of the sensor) and prepares to use that chromaticity as the basis for a "custom white balance" color correction of the ensuing actual shots.

An alternate technique suggested by some photographers and some diffuser manufacturers is to leave the camera, with the diffuser in place, in the location it will have for the actual photographs and make the measurement in that context.

THE COLOR PARROT WHITE BALANCE TOOL, VERSION 1.2

The Color Parrot white balance tool[2] is a new product line recently introduced by Drew Strickland, proprietor of the Pro Photo Home online photography forums. We had earlier reported our observations on and the results of tests of the first version of this product (now identified as the Color Parrot white balance tool, version 1.0). The tool is essentially a white balance measurement diffuser.

The manufacturer has recently introduced an updated design, identified as the Color Parrot version 1.2 (although that notation is not publicized in the promotional literature). They were kind enough to provide us with an early production sample for evaluation.

SCOPE OF OUR REPORT

In this article, we will describe the Color Parrot 1.2 and report on various qualitative "laboratory" tests made of its technical properties. We will in some cases also report, for comparison, some of the same results on the ExpoDisc diffuser, a device of similar function which is often mentioned in the promotional literature for the Color Parrot.

---

[2] From this point on, we will refer to the Color Parrot white balance tool, for conciseness, as just the Color Parrot. Apologies to purists with regard to proprietary terminology ("don't call it a Frigidaire").

PHYSICAL MATTERS

Construction

The Color Parrot 1.2 has an overall configuration reminiscent of a photographic filter of the 82 mm size. The "rim" is evidently of aluminum. The construction is sturdy.

The active ingredient of the Color Parrot 1.1 is a "diffusive" disk which completely fills the area within the "rim". However, all of the disk except a central circular area about 1.05 inches (26.7 mm) in diameter is "blanked" at the rear by a thin opaque black disk. Thus, we in effect have a 1.05 inch diameter diffuser. Uncoated glass plates fore and aft complete the stack.

Handling

The Color Parrot 1.1 is not equipped with filter threads or any comparable facility for attaching it to the front of the lens. It is intended to be held, freehand, in front of the lens while in use. Note that ordinarily the taking of a single "white balance frame" or "white balance measurement" is done very quickly, so the "freehand mount" is not generally burdensome, and is in fact probably more expedient than, for example mounting via filter threads.

A consequence of this approach is that a single Color Parrot diffuser can be used equally well with lenses having a wide range of barrel or front thread diameters.

METHOD OF USE

The manufacturer of the Color Parrot indicates that the device may be used in either of two techniques:

<u>Technique A</u>

Here the Color Parrot is placed in front of the camera lens with the camera located at the subject (or subject location), aimed at the position where the camera will be for the actual shot(s) of the subject.

<u>Technique B</u>

Here the Color Parrot is placed in front of the camera lens with the camera located in the same position it will be for the actual shot(s) of the subject, "aimed at" the subject. [3]

The manufacturer's recommendations as to when each of these should be used are "evolving".

---

[3] We cannot discern the principle by which this is supposed to work.

LABORATORY TEST APPROACH

A number of laboratory tests were made to determine various technical properties that are of interest in connection with the behavior and performance of white balance diffusers.

These tests were all *in vivo*, rather than *in vitro*. That is, they were made with the Color Parrot (or other comparable device being tested) in place on the camera, and the basic “test data” was all “reported by the camera”. That camera is a Canon EOS 20D. Various lenses were used for different aspects of the tests.

It is important to note that tests of the properties of a diffuser *in vitro* (that is, using optical, photometric, and colorimetric instruments with the device “in a fixture”) cannot take into account the effects of the specific way in which the diffuser will interact with the camera. (And we don’t have any of those instruments anyway!)

DIRECTIVITY PATTERN

Background

Of interest as we contemplate the theoretical basis of white balance measurement is what I call the *directivity pattern* of the diffuser (or more correctly, the diffuser-equipped camera), a term I borrow from antenna practice.

This is, in effect, a plot of the relative “weight” given light landing on the face of the diffuser at different angles (that is, arriving from different directions). This is of interest in that the diffuser’s main job is to combine the light arriving from different angles so that the chromaticity of the “mix” is the same as the “effective chromaticity of the overall illumination on the subject.” To achieve that, (according to one widely held theoretical model) the diffuser, in making the mix, should accord the same weight to light components arriving from different angles as would be their relative influence on that effective chromaticity.

If we assume that the subject surface is an “ideal diffuse reflecting surface”[4], then the ideal directivity pattern for the diffuser would be what is called the “cosine” pattern. If the surface were not such, then the ideal pattern would be something else. But we would rarely be in a position to know the exact behavior of the subject surface and select a diffuser whose response matches that. Thus, in my opinion, the use of something like a cosine pattern is probably the best for overall use.

---

[4] Often said to be a “Lambertian” surface.

Test procedure

The procedure uses a "small" light source at a distance of about 20 feet from the camera. This is a cylindrical "bullet" fixture equipped with an incandescent lamp. A "snoot" is fitted to minimize light spill to the sides, which could have illuminated the room walls and polluted the results when the camera was not pointed directly at the light.

The camera's metering system (with the "partial" metering pattern in effect) is used to determine the relative illuminance on the central region of the focal plane for various angles of incidence. The camera reports its determination of that illuminance in terms of the indicated scene luminance (compensating for the lens maximum aperture, which is in effect during metering), which is reported in the "Measured Bv"[5] item in the makernote area of the Exif metadata.

The lens used is the Canon EF 24-105 mm f/4.0 IS, operated at a focal length of 24 mm with focus at infinity.

The camera is carefully aligned with the light source, and the azimuth scale on the tripod pan head set to zero. The diffuser is then put in place for the duration of the test (for the Color Parrot, blue painters' tape is very nice). A measurement exposure is taken at this 0° position. Then the camera is slewed in azimuth by 5° to the left and another measurement exposure taken. This is repeated at further 5° increments through an azimuth of 30°, and then at 15° increments through an azimuth of 90°. (We make no measurements with the camera slewed to the right, assuming that the directivity pattern is symmetrical.)

Then, the Measured Bv is read from each of the image files with ExifTool and entered into a data reduction Excel spreadsheet developed for the purpose. Relative sensitivity values are determined (with the one for 0° being defined as 1.00). A polar plot of the result is then generated in Excel (using Andy Pope's lovely polar plot add-in).

Test results

Figure 1 shows the measured directivity pattern of the Color Parrot 1.2 in the polar plot form just described. The classical cosine pattern is shown (in red) for comparison:

[5] This is actually designated "Measured Ev", but it is fact *bona fide* Bv. (It is not the other common misuse of Ev for luminance, in which "Ev" means "Bv + 5".)

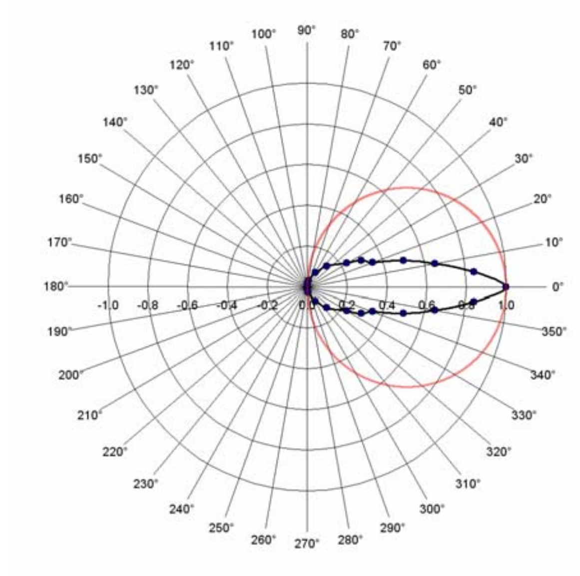

Figure 1. Color Parrot 1.2—directivity pattern

As mentioned above, data was only taken for one side of the pattern (0°-90°, but the pattern is plotted as if symmetrical.

Note that this should not be looked at as a cross-sectional sketch of the "acceptance beam". Rather, it is just a graph, of the polar type, in which the distance from the center of the chart of a point on a certain angle line represents the relative response for that angle.

The width of the pattern at the "50% response" level is about 30° in total (15° on each side). For the classical cosine pattern, the "50% response" width is 120° in total (60° on each side).

Figure 2 shows the pattern for an ExpoDisc diffuser (2007 design):

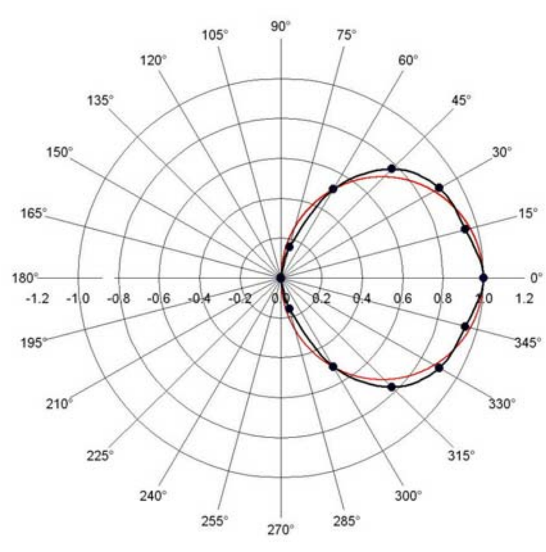

Figure 2. ExpoDisc (2007)—directivity pattern

Data points here were only taken at 15° intervals over the entire 0°-90° range, owing to the pattern being broader than for the Color Parrot. Again, the red circle shows, for reference, the classical cosine pattern.

The width of the pattern at the "50% response" level is about 30° in total (15° on each side), the same as for the cosine pattern.

Manufacturer's data

The manufacturer of the Color Parrot has not released any quantitative information on the directivity of any version of the device. He does state that the Color Parrot is "targeted", and "center-weighted", which we take to mean that it has a narrower directivity pattern than something else (supposedly more so for the v 1.1 than the v 1.0). This is indeed confirmed by our results.

Comment

The manufacturer claims that the relatively narrow directivity pattern of the Color Parrot 1.2 is advantageous, and allows it to provide better performance (especially in the "from the camera position" measurement technique—"technique B") than diffusers with a broader pattern. We have heard no credible rationale supporting this notion.

TRANSMISSIVE SPECTRAL UNIFORMITY

In the section after this, we will discuss determination of the *transmissive chromatic neutrality* of a diffuser, a property that indicates how the diffuser shifts the apparent chromaticity of the light passing through it on the way to the camera for measurement.

Transmissive chromatic neutrality is directly determined by a more fundamental (and comprehensive) physical property, the *transmissive spectral response* of the diffuser. Accordingly, some manufacturers emphasize the importance of that property of a device rather than its chromatic neutrality.

But there is a practical problem with "publishing" spectral response information. If indeed the transmissive spectral response were "absolutely uniform" (and of course, that can't be), then this would immediately tell us that the device was "absolutely chromatically neutral" (which nothing is).

But if we have a transmissive spectral response that is noticeably not uniform, can we tell by looking at the plot how far the device will be from chromatic neutrality? No, unless we can multiply functions of wavelength and integrate the products without taking off our shoes and socks. So in fact, while it would be very interesting to see the spectral response of these diffusers (which will always be "rather close to uniform" for any diffuser anybody would want to buy), it will tell us very little by which we can quantitatively assess the "accuracy" of the diffuser.

Manufacturer's data

The manufacturer of the Color Parrot has not released any *bona fide* transmissive spectral response information for any version.

TRANSMISSIVE CHROMATIC NEUTRALITY

Background

Of great interest in the matter of a white balance measurement diffuser is the matter of its "transmissive chromatic neutrality". Briefly, this means how far does the diffuser depart from the ideal situation in which the chromaticity of the light presented on its backside for observation by the camera is identical to the chromaticity of the overall illumination on its face.

Strictly, perfect neutrality is only attained in the general case (meaning, for light of any spectrum) with a perfectly uniform transmissive spectral response. If a diffuser had that (and none could), we could say "it is perfectly neutral" and be done with it.

But if its transmissive spectrum did depart from uniformity, we would like to be able to say how much that would cause the diffuser to shift the chromaticity of the light passing through it.

Unfortunately, as I mentioned in the previous section, that actually depends on a complex interaction between the the transmissive spectral response of the diffuser and the spectrum of the specific light of interest.

Fortunately, our interest is in a somewhat restricted situation. Firstly, we will generally be considering a device whose transmissive spectral response is not "too far from perfectly uniform". Otherwise, the amount of chromaticity shift given by the diffuser would make it wholly unusable for white balance measurements (and it would be "out of the running").

Secondly, we will generally be dealing with incident light whose chromaticity is not too far from "white" (which we can take here to mean the *white point* of the camera's color space), and which has a fairly "smooth" spectrum. If those stipulations are not so, it turns out that the process of white balance correction is probably doomed anyway, and precisely how a diffuser shifts the chromaticity of the collected incident light would be the least of our concerns.

Once we have adopted those stipulations, it turns out that a particular diffuser will actually shift the chromaticity of the light passing through it by a predictable amount and in a particular direction (as portrayed on some appropriate chromaticity diagram, using an appropriate system of "coordinates").

Thus, the amount of that shift is a good description of how the diffuser departs from perfect chromatic neutrality.

Interaction with the camera

Note that the "chromatic response" of the diffuser plus camera system (our real "instrument" in usage, and also in testing) may vary with exactly what part of the rear of the diffuser the camera regards and considers, and from what range of angles. Thus, the result could vary with other cameras (or even with lens settings).

Presenting the result

I mentioned just above that we need to express the chromaticity shift caused by a diffuser in some system of chromaticity coordinates.

There are several coordinate schemes by which we can express chromaticity. One is the CIE u'v' system. An advantage is that equal distances on its two-dimensional "plane", regardless of their starting point or direction, represent about equal perceptions of chromaticity difference (assuming a relatively-fixed luminance).

In fact, this chromaticity coordinate system is the basis of the chromaticity aspect of the CIE L*u*v* color space, an alternate version of the more-widely used CIE L*a*b* color space, both of them originally designed for description of reflective (or transmissive) color, but now both also used for the color of light. The coordinates "u" and "v" of the L*u*v* color space are derived from values of u' and v'.

We use the designation du' and dv' ("d" is evocative of "delta", for "difference") to label the difference between two chromaticities in the u' and v' directions of the u'v' coordinate system, and du'v' for the total difference.

Although the matter of "how much chromaticity difference can a human detect" is very complex, as a rough reference benchmark, a difference of 0.001 unit on the u'v' plane indicates a difference in chromaticity that can just be perceived in a side-by-side test under optimal viewing conditions.

Can we relate this scale to departure from neutrality as given by the a* and b* coordinates of the L*a*b* model, with which many workers in this field are more familiar? Well, sort of. A set of values of a* and b* does not directly indicate a chromaticity offset. It only does so when interpreted in connection with the accompanying value of L* (lightness, an indicator of the luminance of the color represented).

If we arbitrarily consider a color whose L* value is 70, then for a color whose chromaticity offset on the u'v' plane is:

+0.00071, +0.00071 (a du'v' of about 0.001)

its L*a*b* coordinates would be:

70, +0.24, +0.57

Test methodology

We unfortunately don't have any standardized source of light of a known chromaticity to use as the input to our measurement system, nor can we actually determine precisely on an absolute basis what chromaticities various camera responses indicate. So we have to make our determinations on a "differential" basis. Fortunately, our needed result is also "differential" (the departure of the transmissive chromatic response of the diffuser from "perfect neutrality").

Thus, we work by comparing (a) the effect of the diffuser on the chromaticity of light passing through it with (b) the effect of a reflective test target of "known" properties on the light reflected from it. Our test target, a WhiBal "gray card", has a "pedigree" in that the manufacturer has kindly supplied us with precision measurements of its reflective chromaticity.

The test environment and chromaticity indicator was, as before, a Canon EOS 20D camera, now equipped with a Sigma DC 18-200 mm f/3.5-6.3 OS lens.[6]

The chromaticity of the light reaching the camera sensor was read by examining the captured digital image. The average chromaticity over a small rectangular area in the center of the frame was used.

To minimize the impact of quantizing error, which we would have suffered had we read the chromaticities in terms of (8-bit) RGB values, we actually used the raw output files from the camera, converted them to 16-bit TIFF form (using Canon Digital Camera Professional), and read the RGB values from these files on a 16-bit basis (using Photoshop CS2).

To assure that the chromaticities of the delivered images were not too far from the reference white of the camera's color space (otherwise, some of our tidy assumptions about chromaticity shift would not have held up), we made all the actual measurements with the camera's

---

[6] This was for a pragmatic reason: the 200 mm focal length available on this lens allowed me to fill the critical area of the frame with the WhiBal gray card while standing where I wouldn't block any of the incident light!

custom white balance (CWB) mode in effect. The CWB balance reference frame was taken from our neutral test target, exposed to the same incident illumination as will be used for the actual test. The conversion of the raw data files to 16-bit TIFF form mentioned above was done with a white balance color correction applied in accordance with the camera's analysis of that reference frame (as would have been used for in-camera CWB color correction of a JPEG output file).

However, our method of ultimately determining the chromaticity shift of the diffuser was not dependent on "perfection" of this white balance process—it was just used to get the chromaticities "in the ballpark".

The conversion of the chromaticities from 16-bit RGB to u'v', and the comparisons that give the final result, were all done with an Excel spreadsheet developed for this purpose.

The result is reported as the departure of the chromatic transmission of the diffuser from perfect neutrality, as shown on a "du'v'" chromaticity chart.

The results

The results of two runs of the chromatic neutrality test are shown on figure 3. The results for the ExpoDisc diffuser (2007 version) are also shown for comparison. The origin (0,0) represents "perfect neutrality".

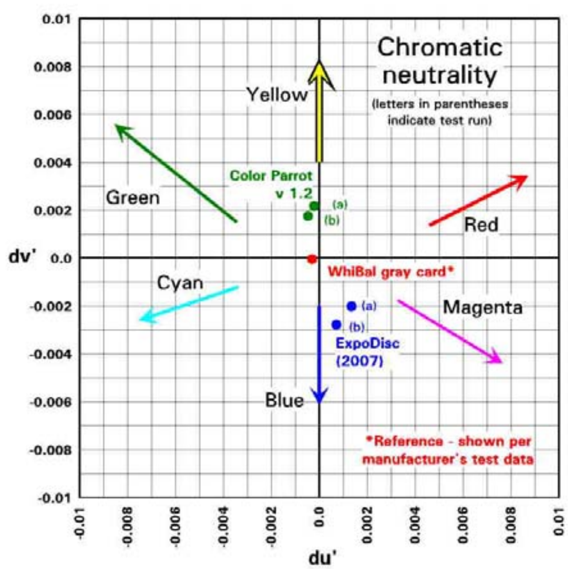

Figure 3. Transmissive chromatic neutrality

The two runs used the same incident light source and the same overall configuration, but the "setup" was remade for each run. The consistency between the two sets of data is within what we might consider the overall expected exposure to experimental error (less than 0.001 du'v' unit).

Also plotted on the chart is the assumed chromaticity of our reference surface, a "pedigreed" WhiBal gray card. That information is based on precise measurements made of our particular item by its manufacturer. The departure of the assumed response of the reference target from perfect neutrality not a source of error in the results; it is accounted for in the analysis of the results for the tested diffusers.

Having used it as our reference, we cannot state any "independently measured" neutrality for the WhiBal gray card. In any case, its evaluation is not an objective of this report.

The average measured neutrality error for the Color Parrot 1.2 is about 0.002 du'v' unit.

The average measured neutrality error for the ExpoDisc (1977) is about 0.0025 du'v' unit.

Manufacturer's data

The manufacturer of the Color Parrot has not released any *bona fide* transmissive chromatic neutrality information for any version of the product.

Comment

We would expect that the very modest chromatic neutrality errors of the Color Parrot 1.2 would not be of consequence in practical photographic white balance use. The same could be said for the ExpoDisc (2007).

EQUIVALENT TOTAL TRANSMISSION

Background

Of some interest is the matter of the equivalent total transmission of a diffuser. We can for the moment understand the concept (somewhat imprecisely) by thinking of this as indicating the fraction of the light incident on the face of the diffuser that exits from its backside.

But "fraction of the incident light . . ." isn't really quite accurate, since we don't really know the angular and spatial distribution of the emission from the backside of the diffuser, nor exactly how the camera deals with that. In any case the different dimensionality of the two quantities (one an illuminance and the other a luminance) wouldn't allow them to be directly compared anyway.

Our real interest is in the matter of, for any given lens aperture, and considering a certain illuminance on the diffuser face, how the illuminance on the sensor relates what it would be if the camera (without diffuser) regarded a 100% reflectance object illuminated by the same illuminance as we have on the face of the diffuser. The term "equivalent" is a cue as to this situation.

Why are we concerned with this? The manufacturer of the Color Parrot points out that some cameras, in certain of their "custom white balance" measurements modes, apparently balk at performing if the illuminance on the sensor is less that a certain level. He points out that the Color Parrot has a higher value of total transmission than certain competitive diffusers, and as a result, he says, it can be successfully employed (for those camera modes) in lower light than those other units.

Test procedure

We determined the effective total transmission of the Color Parrot by comparing the camera response to the light presented by the diffuser, illuminated by a certain arbitrary incident light, with the corresponding response of an ExpoDisc diffuser, whose total transmission had been measured by the manufacturer. It is recognized that this approach only yields an approximation (we believe a reasonable one) of the precise performance of the Color Parrot (which in any case would be quite variable depending on the nature of the interaction of the diffuser with the particular camera).

Our assessment of the relative luminance on the focal plane was made by way of the reading of the camera's internal exposure metering system, with the "partial" metering pattern in play. The reading was discerned from the Measured Bv value from the makernote area of the Exif metadata (read with the program ExifTool).

Over the various test situations used, the Color Parrot exhibited an equivalent total transmission of about 1.6-1.75 stops greater than the ExpoDisc (3.0-3.4 times as great). The ExpoDisc used has a rated effective total transmission of 0.182, and thus that of the Color Parrot would appear to be in the range of 0.55-0.60.

Manufacturer's data

The manufacturer of the Color Parrot has stated that version 1.2 should exhibit 2-3 times the total transmission of the ExpoDisc. Our results are consistent with this.

Comment

The fairly large equivalent total transmission of the Color Parrot will evidently allow the device to be used successfully for low light measurements with cameras and modes that are not willing to make measurements with a sensor luminance below a certain level. This seems to include some Nikon models when used in the "on the fly" white balance measurement mode, where a reference image is not taken for later examination). The property should not be of major concern otherwise, as for example with the custom white balance procedures of Canon EOS cameras.

CONCLUSION

The Color Parrot white balance tool is a sturdy and convenient-to-use white balance measurement diffuser. It exhibits a directivity pattern that is considerably narrower than that of other white balance measurement diffusers.[7] Its degree of chromatic neutrality is quite respectable, a desirable property. It has a higher effective total transmission than some other white balance diffusers, which can avert difficulty in making white balance measurements in low light with certain cameras.

#

[7] We cannot see just how this is advantageous.