

A Robust Illumination Estimate for Chromatic Adaptation in Rendered Images

Alexander Wilkie¹ and Andrea Weidlich²

¹Charles University in Prague, Czech Republic

²Vienna University of Technology, Austria

Abstract

We propose a method that improves automatic colour correction operations for rendered images. In particular, we propose a robust technique for estimating the visible and pertinent illumination in a given scene. We do this at very low computational cost by mostly re-using information that is already being computed during the image synthesis process. Conventional illuminant estimations either operate only on 2D image data, or, if they do go beyond pure image analysis, only use information on the luminaires found in the scene. The latter is usually done with little or no regard for how the light sources actually affect the part of the scene that is being viewed. Our technique goes beyond that, and also takes object reflectance into account, as well as the incident light that is actually responsible for the colour of the objects that one sees. It is therefore able to cope with difficult cases, such as scenes with mixed illuminants, complex scenes with many light sources of varying colour, or strongly coloured indirect illumination.

Categories and Subject Descriptors (according to ACM CCS):

Computer Graphics [I.3.7]: Three-Dimensional Graphics and Realism—Colour, shading, shadowing, and texture

Image Processing and Computer Vision [I.4.8]: Scene Analysis—Colour

1. Introduction

In digital photography, the chromatic adaptation of the viewer is routinely – though not always entirely successfully – compensated by the white balance post-processing performed by the camera software. This issue does not only affect photography, though, but is potentially relevant for all colour imagery, even renderings.

For most of current 3D graphics work, the luminaire colours are directly chosen so that the appearance of the resulting image meets the expectations of the designer, so no actual white balance step is needed. However, if the luminaire colours for a scene are fixed, e.g. for the evaluation of an interior design that involves a particular type of light, a white balance step has to be performed. For such applications that attempt to predict scene appearance, image- or scene-based illumination colour estimation methods are normally used as part of the tone reproduction process. These have shortcomings, though, and can fail for certain types of scene; figure 1 shows a such a problematic setting.

In this context, the contribution of our research is to

demonstrate that one can perform highly robust white balance computations for rendered images by appropriately re-using some of the information that has to be computed during image synthesis anyway. Our technique is very simple, and can completely replace conventional image-based illumination colour estimates for rendered images.

2. The Phenomenon of Chromatic Adaptation

The processes that govern the chromatic adaptation state of a human observer are complex, subtle, and not entirely understood in all their details, although we can claim to understand their most important aspects [Boy79, LM71, Fai98, Ebn07]. From an end-user perspective, it is usually sufficient to consider this to be a process that can identify the colour cast caused by the illumination in a scene; the actual object colours in a scene only have a very small influence [Bae99].

One of the numerous problems with the derivation of accurate models of this phenomenon is that it is not something that happens on a sensor level within the eye alone; cognitive functions do play a significant role as well. Any objects

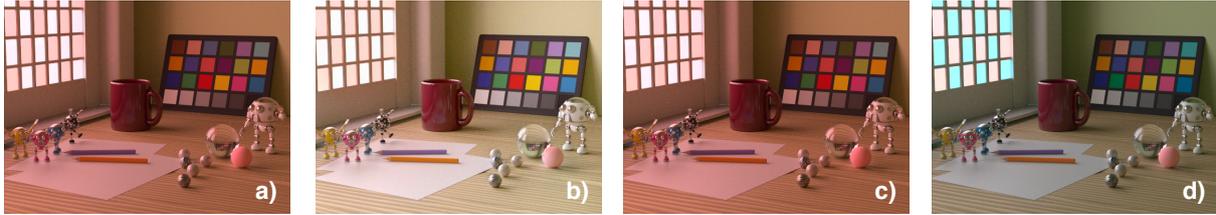


Figure 1: A rendering where colour correction of the result is necessary, but where conventional techniques have problems in determining what to compensate for. The room is lit by daylight, but the window consists of alternating pink and purple stained glass panes. **a)** raw image, **b)** colour correction performed with the Retinex method – a slightly problematic residual colour cast is evident, **c)** luminance analysis [WEV02] yields no compensation at all, and **d)** a fully automatic compensation performed by our method. The image-based technique **b)** only performs passably well because this particular scene contains white objects; if these were removed, it could start to have problems. By comparison, our proposed technique simply works, regardless of scene content, and luminaire arrangements. In image **d)**, the cyan colour of the purple glass panes is the result of applying the same Bradford transform to all pixels in an HDR image, even the very bright ones. This artefact does not affect the point we wish to make here, though, namely that our method is capable of automatically determining the correct colour to compensate for.

of neutral colour that are present in a scene (usually referred to as the *reference white objects*) apparently play a major role in the identification of the illumination colour cast; the perceived colour of these objects closely correlates with the illumination. The mechanism that identifies these objects in a scene is a higher brain function, and not easy to analyse. This situation makes it hard to replicate this process in software, especially if only 2D image data is available as input.

2.1. Models of Chromatic Adaptation

Most of the work on computational models of this phenomenon has been in the field of image processing and machine vision, with sidelines in digital camera development, and actual computer graphics. The main goal is to estimate the unknown illuminant(s) found in an image, and to use this information to correct any colour cast that is present. Figure 2 gives an overview of the workflow.

The actual colour correction that is performed in the second stage is usually done by applying a von Kries or a Bradford [Lam85] transform, or the more recent CIECAT02 transform [MFH*02]. The correction step is not particularly complex, so the main technical difficulty of image based white balance is the identification of the illuminant colour.

2.1.1. Image Based Colour Constancy Methods

All the techniques discussed in this section are only based on the analysis of a single 2D input image, and therefore limited in one way or another. No one-size-fits-all image based technique has been found so far, although existing methods offer a reasonable level of reliability for common usage scenarios. Due to the large amount of theoretically relevant work in this very active [HMP*08] and diverse research area [Bar99,Ebn07], we can only give a broad overview over the main existing techniques. Moreover, since the focus of

this research is usually the derivation of a colour constant image descriptor that can be used for shadow removal or similar processing tasks, some of the methods are only of indirect use if one wants to simulate human perception.

Gray world techniques [Buc80] are based on the assumption that, if all surface colours in a scene are averaged, the scene is gray (neutral); any difference that is found can be assumed to be a colour cast, and will be compensated. *White Patch algorithms* assume that at least one object in the scene is white. Any colour cast found there will be compensated; this is also the idea behind the Retinex method [LM71]. A more sophisticated approach is that of the *Gamut Constraint Methods*, which consider the possible gamut of object colours under a known illuminant, compare this to the gamut of colours found in the image, and derive the necessary shift from the difference between the two [FFB95,Fin96]. *Local Shift* techniques [Ebn07] go further, attempt to analyse any local variations in illumination across the image.

2.1.2. Scene Based Colour Constancy for Renderings

The same techniques that were developed for image processing can also be used on renderings. Although a rare functionality in rendering engines, this does get discussed and implemented from time to time in a tone reproduction context [FPSG96,PFFG98,DD00], or more recently in conjunction with models of human colour perception [KWK09]. One crucial difference to photography is that for renderings one actually has information about the depicted scene at one's disposal. [WEV02] were the first to make use of this, and proposed a direct correction for the illuminants defined in a scene. This approach works very well if only a single illuminant is present.

However, knowing the reflectances in a scene, or the colour of the illuminants, does *not* necessarily help with the automatic white balance computations, at least not for all

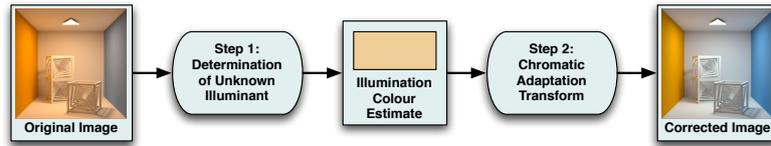


Figure 2: The two stages of a white balance workflow: a raw image is analysed to determine the illuminant, and then an appropriate correction is applied. Our contribution is an improvement to step 1, which so far was the weak link in this chain.

possible scene types. This information is in *object space*, but the white balance has to be computed for the resulting *image*. There is, for example, no guarantee that any of the white objects in a scene will actually be visible in a rendered image. The same holds for the illuminants: although their colour is known, not all of them will necessarily contribute to the image in question.

Of course it would be possible to derive scene analysis heuristics that deal with such cases. But as we will see in the next section, it is much simpler and more efficient to store and analyse related information that can be obtained as a by-product during rendering, instead of directly analysing the 3D model of the scene. Also, no such analysis can yield information on indirect illumination colour, which can be very significant in some scenes. Our proposed technique is also capable of this without any additional modifications.

3. Our Approach: Scene-Based Chromatic Adaptation

The basic idea of our technique is to extract some easily computable additional information during rendering, and to use this later during the colour correction. More specifically, we re-use data that is already being computed during image synthesis to reduce the colour constancy problem to a setting where the assumptions that are typical for gray world techniques reliably hold. This yields a robust illumination colour estimate, and allows us to closely mimic the colour adaptation process behaviour of the human visual system. We do this by biasing the process so that the illumination that falls on white objects becomes a preferred source for the estimate.

3.1. Storing Scene Information in Auxiliary Images

The key concept to realise these goals is that we generate two additional, auxiliary images during rendering:

- an image that shows what a given scene would look if direct and indirect illumination were neutral, the **CWNI** image (**C**oloured **W**orld **N**eutralised **I**lluminants),
- and how this scene would look like if all BRDFs were replaced with surfaces of neutral colour, the **NWCI** image (**N**eutralised **W**orld **C**oloured **I**lluminants).

These two auxiliary images are then later used to guide the actual colour constancy computations. It is fairly straightforward to modify an existing global illumination renderer

so that these two additional output images are created as by-products of the rendering calculations, along with the normal image that is being computed. Both images are cheap to obtain, since no additional ray-casting operations or BRDF samples have to be done. Figure 3 gives an overview of the basic principle. Although we demonstrate our technique for a ray-based renderer, it would also work for a scanline/GPU renderer, or even with radiosity methods, where the needed incident illumination information can be taken from the radiosity data during the shading computations.

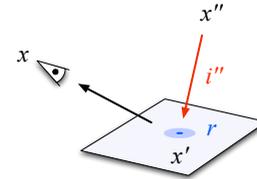


Figure 3: The content of the auxiliary images interpreted in accordance with the rendering equation [Kaj86]. The **CWNI** (**C**oloured **W**orld **N**eutralised **I**lluminants) image is computed by replacing the incident illumination i'' (red) with neutralised light of similar brightness, and the **NWCI** image (**N**eutralised **W**orld **C**oloured **I**lluminants) by replacing the reflectance r at point x' (blue) with a neutralised reflectance of similar albedo.

3.1.1. The CWNI Image

The first of the two auxiliary images, the **CWNI** image, shows the scene under a neutral illumination – regardless of whether this incident light is direct or indirect illumination, and regardless of its original colour. This allows us to reliably identify those objects in the scene that are genuinely white, irrespective of their appearance in the main image. We later use this information to bias the chromatic adaptation simulation towards giving greater emphasis to the illumination that falls onto white objects, since the human brain apparently works along similar lines in this regard. The **CWNI** image is computed by doing the following at the first intersection of the viewing ray with the scene (and *only* there):

- the energy spectrum of the incident illumination at this point is averaged, to obtain a neutralised light of similar brightness. It is important to use a *neutralised* version

of the incident light, and not a pre-defined neutral illuminant here. Neutralising the actual incident light retains the power relationships between the illuminants in a scene if there are several of them, and also retains the indirect illumination.

- this neutralised incident light is multiplied with the current BRDF sample; the result is written to the CWNI output image.

The resulting image looks more or less like a normally rendered version of the scene under white light, even though it is *not* an unbiased GI solution of this type of scene; see figure 4 for an example. The main "inaccuracy" compared to a real solution is that while the magnitude of any indirect illumination is correct, its colour is not, so any colour bleeding effects are removed.

However, this is exactly what we want. Due to the neutralisation of all incident light, the identification of white objects based on the CWNI image delivers correct results even in the presence of arbitrarily strong, coloured indirect illumination that is distinct from any light source found in the scene.

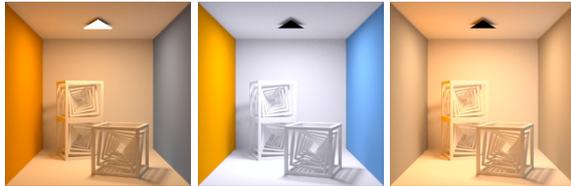


Figure 4: Raw image (left), and the corresponding CWNI (middle) and NWCI images (right). Directly viewed luminaires are omitted so they do not interfere with the identification of white objects.

3.1.2. The NWCI Image

The second auxiliary image, the NWCI image, shows the scene under its original illumination, but with all the surfaces replaced with a neutralised version of themselves. This is simply done by

- replacing the BRDF sample value with a neutralised reflectance spectrum of similar overall albedo
- multiplying the original illumination with the neutralised BRDF, and writing this result to the NWCI output image.

The purpose of the NWCI image is to yield a reliable estimate of incident illumination colour for all pixels, regardless of the surface colours found in the scene; figure 4 shows an example. This image is more obviously "wrong" than the CWNI image – note that all colour bleeding effects are still clearly visible, even though there would be nothing in a genuinely neutral scene that could cause them.

In theory, the NWCI image could already be used to get an illumination estimate, at least for mostly Lambertian environments. Scenes with specular objects are more problem-

atic; as figure 5 demonstrates, objects seen in specular reflections retain their colours, so direct use of NWCI images is inadvisable. But as we discuss in the next section, one can both make the estimate robust for specular scenes, and also improve the quality of the chromatic adaptation so that reference white objects are taken into account, by combining the NWCI information with the CWNI image.

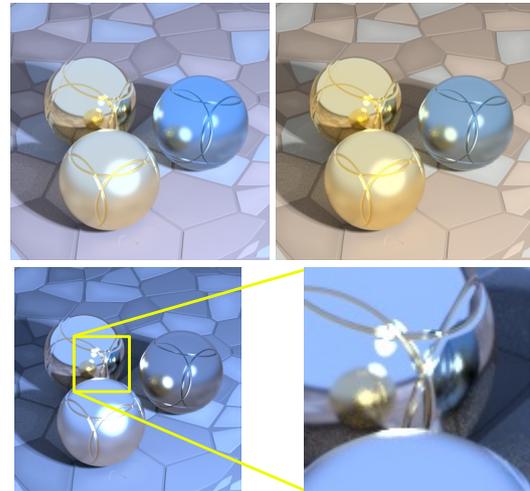


Figure 5: Issues with NWCI images: a scene with specular metallic spheres on beige tiles under a strongly blue illumination before (top left) and after colour compensation (top right). The NWCI enlargement (bottom) shows that only the colour of the mirror is neutralised, but the objects that are being reflected retain their colours.

3.2. Combining the CWNI and NWCI Images

The CWNI and NWCI images are created as by-products of the rendering process, and are only used later, during the tone reproduction step. Making the most of the information contained in them requires the creation of two additional auxiliary images. Both additional images are only created during the tone reproduction computations, so that the effect on rendering performance is negligible. The entire proposed white balance workflow is outlined in figure 6.

3.2.1. The "1 - Chroma" Image

The first additional image is created by extracting a grayscale image, which we refer to as the "1 - Chroma" image, from the CWNI image. Its purpose is to increase the contribution of those parts of the image which contain white objects. When it performs its own colour correction, the human visual system places an emphasis on the illumination that falls on such objects, and the 1-Chroma image is a tool for emulating this: a monochrome version of the rendered image which is dark in areas of high surface chroma, and

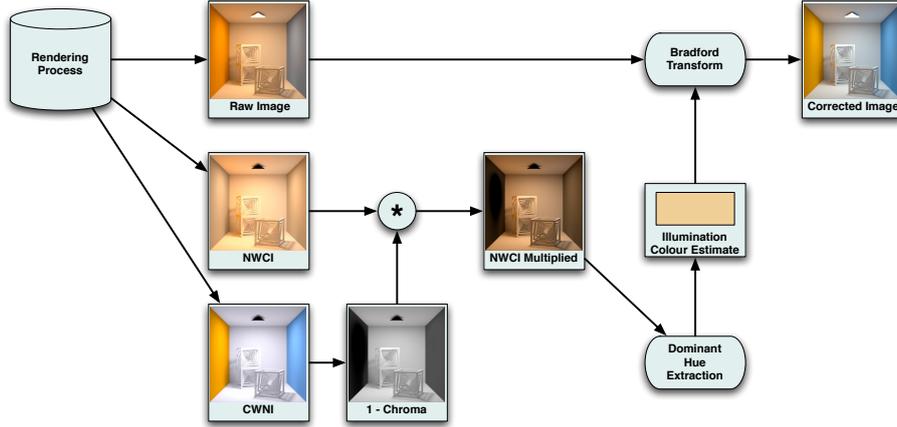


Figure 6: Flowchart of our entire proposed algorithm. The detailed meaning and creation of the individual auxiliary images (CWNI, NWCI, 1-Chroma, NWCI Multiplied) are discussed in sections 3.1.1, 3.1.2, 3.2.1 and 3.2.2, respectively.

bright in areas of low surface chroma. Figure 7 shows an example of this.

To do this, we convert the image to CIE LCH [WS82] space. Then, given the minimum and maximum chroma and luminance values C_{min} , C_{max} and L_{max} of all pixels in the CWNI image, and we store the value

$$g = \left(\left(1 - \frac{C_{pixel} - C_{min}}{C_{max} - C_{min}} \right) * \frac{L_{pixel}}{L_{max}} \right)^w \quad (1)$$

for each pixel. The parameter w controls the influence of this "white object information"; if it is left at 1, a wider range of pixels is used, while higher values create a successively stronger bias towards only using the most neutral pixels. The largest practical value of w appears to be at about 4; larger values do not seem to provide any additional useful bias, and a value of 2 seems to be optimal. Unfortunately, no quantitative model of how reference white objects influence human perception exists, so we had to resort to trial and error to find these values. Figure 12 demonstrates the effect of varying w .

3.2.2. The "NWCI Multiplied" Image

The content of the "1 - Chroma" image is then used to scale the NWCI image on a pixel-by-pixel basis. This yields the "NWCI Multiplied" image (figure 7), from which the illumination colour is later extracted. Areas in the NWCI image where illumination falls onto reasonably neutral objects are emphasised by this multiplication. A reliable and robust identification of these areas would be hard by using only 2D image information. As shown by the images in figure 8, it also prevents any negative impact of the peculiar feature of plain NWCI images – shown in figure 5 – that specular reflections retain their colour.

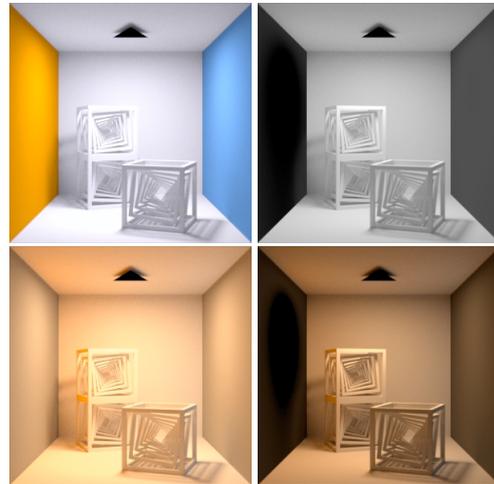


Figure 7: *Top:* comparison of a CWNI image (left) and the corresponding "1 - Chroma" image (right), which is bright in exactly those areas that are well lit and that contain the pixels with the least chroma. It is also dark in those with the highest chroma values, and in those areas that are not well lit, and therefore not perceptually important. *Bottom:* comparison of an NWCI image (left) and the corresponding "NWCI Multiplied" image (right). The latter emphasises illumination that falls onto neutral objects.

3.3. Performing the Colour Correction

As the actual colour correction operation we use a Bradford [Lam85, Lin] transform, that is applied after the image has been converted to colour space and tone mapped, but before it is reduced from floating point colour values to a 24 bit integer image. The transform operates on colours given

in CIE XYZ coordinates, and moves colours from an uncorrected source space to the corrected destination space via multiplication with a matrix M :

$$[X_d, Y_d, Z_d] = [X_s, Y_s, Z_s][M]$$

The formulas for the matrix M only require the CIE XYZ chromaticities of source and destination white as input:

- The *source white* that we use for our corrections is derived from the "NWCI Multiplied" image through averaging of the pixel chromaticities. This is equivalent to a simple gray world approach; the result is directly assumed to be an estimate of the unknown illumination. Since the underlying assumption of gray world algorithms is guaranteed to be met (the "NWCI Multiplied" image is tailored to satisfy it), this result can be directly used.
- The *destination white* is the default white point of the output colour space one chooses for the tone mapped image. Usually, this is CIE D50, the default white point for the CIE $L^*a^*b^*$ colour space, or the white point of the display targeted by the tone mapping step.

We use a Bradford transform because it is sufficient for our purposes; however, other, more sophisticated techniques could easily be substituted in its place [MFH*02, CGZ07].

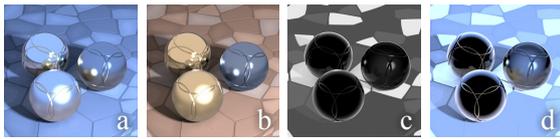


Figure 8: A demonstration of how the interaction between the NWCI (a), CWNI (b) and "I-Chroma" (c) images removes the potentially unwanted coloured reflections in the NWCI image of the scene shown in figure 5. The "NWCI Multiplied" image (d) contains only information from diffuse, neutral surfaces, and neutral specular reflections.

3.4. Performance Impact

For an unbiased global illumination renderer, we found that the computation of the two additional output images causes an almost negligible overhead (see table 1) even if they are created at the same resolution as the original image. If their creation turned out to be a performance problem, e.g. if our technique were to be used in a real-time rendering application, it is important to note that the CWNI and NWCI images can – and probably even *should* – be rendered at a much smaller resolution than the original image. Our experiences suggest that *very* small resolutions (in the order of 50×50 pixels and less) are completely sufficient. The only reason we did not implement some form of sub-sampling for the auxiliary images was lack of incentive in the case of the GI renderer we used; any performance gains would have been hard to measure, given the already low impact of writing the two additional images at full resolution.

	without - plain path tracer	with our method
Figure 5	238.7 seconds	239.8 seconds
Figure 9	92.76 seconds	94.86 seconds

Table 1: Overall rendering times for two typical images.

4. Evaluation of Our Approach

Our goal is to reproduce a scene as it would be perceived by an observer who is fully adapted to the surroundings. Since there are no formal models of this phenomenon, we have the same problem as all other colour correction algorithms – there is no "correct" solution, and we have to be content with delivering a reasonable approximation of what the visual system does. So we have to demonstrate that our technique

1. works as advertised for normal, realistic scenes with complex illumination and object reflectances, and
2. can also handle a number of problematic scenarios that are difficult or impossible for image-based techniques.

The test images can be seen in table 4. For all "problematic" test cases there is at least one conventional algorithm that also works reasonably well, but none of them works for *all* test cases. Apart from a couple of pathological configurations where no colour correction is possible anyway (such as scenes illuminated by monochrome illuminants – in this case the theory behind the Bradford transform fails; our method still gathers the correct illumination information), we were not able to find any scenarios where our approach fails. Together with its relative simplicity and low computational cost, this suggests that our method is a viable technology for colour constancy computations on rendered images. In the following paragraphs, the circled number corresponds to the test case number in table 4.

① ② **Standard Scene:** the test scene is illuminated by CIE Illuminant A and a green light, with a small contribution from a dark blue late evening sky outside the window. The point of these test cases is not to show that one can perform an automatic white balance computation here – after all, most conventional, image-based algorithms would work for this scene. What we demonstrate is that our technique works reliably for standard environments.

③ **Transmission Colour:** the scene from figure 1.

④ **Indirect Illumination:** a problematic scenario for methods that analyse the luminaires are scenes with strongly coloured indirect illumination. The luminaire in the scene is white, but is oriented so that it shines towards the blue wall. Our technique would also work properly if the coloured wall were not visible, if the luminaire was also strongly coloured, and if no white objects were present.

⑤ ⑥ **Dominant Overall Hue:** images with a strong colour bias towards a specific hue are a problem for image-based

methods, because they can violate some of the common assumptions made about the scene. To demonstrate that our technique also works for such scenes, we created four types of "3D Mondrians" that contain a several randomly coloured diffuse blocks under a diffuse illumination. They are arranged by increasing difficulty for image-based techniques:

1. The first contains an even selection of random colours, as well as one white block. Illumination is a reddish orange. This is potentially easy for image-based methods.
2. Similar to the first type, but without the white block. This is harder for image-based techniques, but still solvable, especially if the distribution of the colours is even.
3. A type that contains a strongly skewed sample of colours, as well as one white block. The illumination is cyan. This is still solvable for all those image-based techniques that are capable of identifying the white block.
4. Similar to type 3, but without the white block. This is a very hard problem for all image-based algorithms.

Figure 9 shows Mondrians of all four types. Note that our method performs reliably and *consistently* for all of them. One usage scenario where this kind of consistency is vital are animations – the colour correction should not change just because the only white object in the scene moves outside the field of vision during a camera zoom.

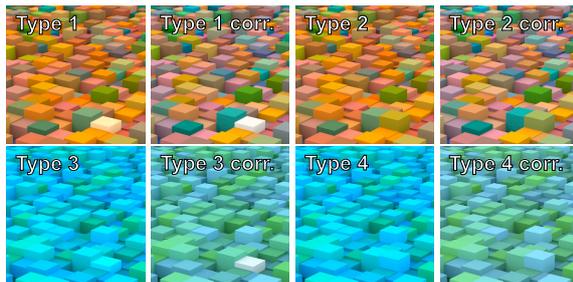


Figure 9: Results for all four types of 3D Mondrian scenes introduced as "Dominant Overall Hue" test cases in section 4. Our technique consistently applies the same correction, regardless of whether white objects are present or not.

⑦ ⑧ **World of One Reflectance:** a study by [RB07] shows that human observers can distinguish between neutral scenes that are illuminated by a coloured light, and coloured scenes under neutral illumination. This confirms earlier results by [GJ84], and implies that a chromatic adaptation algorithm should also behave accordingly. The awareness of observing a neutral scene under a coloured illumination should lead to a colour compensation, while coloured scenes should remain much the way they are. White balance techniques that only use 2D image data as input have a hard time here, while our method is able to distinguish between these two cases.

Multiple Coloured Illuminants and Reference White Objects: one scenario where the CWNI image is essential to delivering good results are scenes with just a few white objects,

and multiple coloured illuminants. Figure 10 shows such a set-up: the scene is mostly made up of brightly coloured objects, and only one of the crate-like objects in the foreground is actually white. This object is lit by a light source that has a different colour than all the others in the scene.



Figure 10: A box with coloured walls and a green luminaire on the ceiling, and where the only prominent white object in the scene is being illuminated by a strong red light source (*left* the uncorrected image, *middle* only based on the contents of the NWCI image and *right* the corrected version).

Human observers base their chromatic adaptation on reference white objects in a scene. The degree by which our method favours information from white surfaces over others is a user-controllable parameter: figure 12 shows an image sequence where the selectivity parameter w of the "1-Chroma" image is varied between 1 and 4 for the example scene from figure 10. For the comparatively high value of $w = 4$ used to generate figure 10, the compensation is done almost exclusively against the red illuminant; figure 11 shows the CWNI and NWCI images for this scene.



Figure 11: The CWNI and NWCI images for the scene shown in figure 10. The CWNI image shows that the object in the foreground is the only white object, and the NWCI image shows that the incident illumination is green and red.

Figure 10 shows what the results would be, if only the contents of the NWCI image were used without any scaling by the "1-Chroma" image. The resulting compensation is ambiguous, and without the additional information found in the CWNI image the method lacks the capability to selectively focus on the illumination that falls on the white object.

As already mentioned, the *extent* to which the presence of reference white objects influences the chromatic adaptation of the human visual system has not been quantified in a rigorous fashion yet. We are not aware of a mathematical model that would allow a prediction of how much influence a given white object in a scene context will have on the adaptation

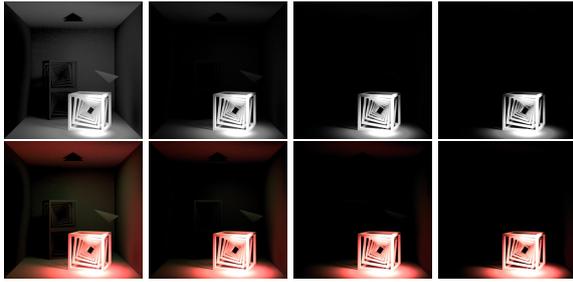


Figure 12: The effect of the "1-Chroma" selectivity parameter w (as defined in equation 1) for the scene shown in figure 10. The top row shows the "1-Chroma" images, while the bottom row contains the corresponding "NWCi Multiplied" images. The values of w are (from left to right): 1, 2, 3, 4.

state of an observer. The almost complete compensation of the red illuminant applied in figure 10 is almost certainly too strong, and was mainly done to demonstrate the capabilities of our technique. But even if we do not yet know how much influence a given white object should have on the overall correction of a scene, the development of the *technical capability* to easily perform such a selective correction (which is what our technique is about) is still a useful improvement.

5. Conclusion and Outlook

We demonstrated a simple and robust way to automatically determine the necessary colour shift for chromatic adaptation computations of rendered images. Our technique behaves in a fashion that is consistent with what one would expect in a human observer. Possible future improvements would involve investigation of more sophisticated analysis algorithms on the data contained in the "NWCi Multiplied" image. One might also want to investigate the derivation of efficient gamut constraint techniques that use CWNI images.

Acknowledgements

The authors would like to thank Greg Ward for his help with improving the presentation of the paper. We would also like to thank Thomas Gamper for his constructive comments on early versions of the paper, and the anonymous reviewers for their helpful input.

References

[Bae99] BAEUML K. H.: Color constancy: the role of image surfaces in illuminant adjustment. *Journal of the Optical Society of America* 16, 7 (July 1999), 1521–1530.

[Bar99] BARNARD K.: *Practical colour constancy*. PhD thesis, Burnaby, BC, Canada, Canada, 1999. Adviser-Brian Funt.

[Boy79] BOYNTON R. M.: *Human Color Vision*. Holt Rinehart and Winston, 1979.

[Buc80] BUCHSBAUM G.: A spatial processor model for object color perception. *Franklin Inst.* 310 (1980), 1–26.

[CGZ07] CHONG H. Y., GORTLER S. J., ZICKLER T. E.: The von kries hypothesis and a basis for color constancy. In *International Conference on Computer Vision* (2007), pp. 1–8.

[DD00] DURAND F., DORSEY J.: Interactive tone mapping. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000* (London, UK, 2000), Springer-Verlag, pp. 219–230.

[Ebn07] EBNER M.: *Color Constancy*, 1 ed. Imaging Science and Technology. John Wiley & Sons, Apr. 2007, pp. 198–204.

[Fai98] FAIRCHILD M. D.: *Color Appearance Models*. Addison-Wesley, 1998.

[FFB95] FINLAYSON G. D., FUNT B. V., BARNARD K.: Color constancy under varying illumination. In *International Conference on Computer Vision* (1995), pp. 720–725.

[Fin96] FINLAYSON G. D.: Color in perspective. *IEEE Trans. Pattern Anal. Mach. Intell* 18, 10 (1996), 1034–1038.

[FSPG96] FERWERDA J. A., PATTANAIK S. N., SHIRLEY P., GREENBERG D. P.: A model of visual adaptation for realistic image synthesis. In *SIGGRAPH 96 Conference Proceedings* (Aug. 1996), ACM SIGGRAPH, Addison Wesley, pp. 249–258.

[GJ84] GILCHRIST A., JACOBSEN A.: Perception of lightness and illumination in a world of one reflectance. *Perception* 13 (1984), 5–19.

[HMP*08] HSU E., MERTENS T., PARIS S., AVIDAN S., DURAND F.: Light mixture estimation for spatially varying white balance. *ACM Transactions on Graphics* (Aug. 2008).

[Kaj86] KAJIYA J. T.: The rendering equation. *Computer Graphics (SIGGRAPH '86 Proceedings)* 20, 4 (Aug. 1986), 143–150.

[KWK09] KIM M. H., WEYRICH T., KAUTZ J.: Modeling human color perception under extended luminance levels. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2009)* (2009).

[Lam85] LAM K.: *Metamerism and Colour Constancy*. PhD thesis, University of Bradford, Bradford, Great Britain, 1985.

[Lin] LINDBLOOM B.: Colour science website. <http://www.bruce.lindbloom.com/>.

[LM71] LAND E. H., MCCANN J. J.: Lightness and retinex theory. *Journal of the Optical Society of America* 61, 1 (1971), 1–11.

[MFH*02] MORONEY N., FAIRCHILD M. D., HUNT R. W. G., LI C., LUO M. R., NEWMAN T.: The CIECAM02 color appearance model. In *Color Imaging Conference* (2002), IS&T - The Society for Imaging Science and Technology, pp. 23–27.

[PFFG98] PATTANAIK S. N., FERWERDA J. A., FAIRCHILD M. D., GREENBERG D. P.: A multiscale model of adaptation and spatial vision for realistic image display. In *Computer Graphics Proceedings* (1998), ACM SIGGRAPH, pp. 287–298.

[RB07] RUPPERTSBERG A. I., BLOJ M.: Reflecting on a room of one reflectance. *J. Vis.* 7, 13 (10 2007), 1–13.

[WEV02] WARD G., EYDELBERG-VILESHIN E.: Picture perfect RGB rendering using spectral prefiltering and sharp color primaries. In *13th Eurographics Workshop on Rendering* (Pisa, Italy, 2002), Eurographics Association.

[WS82] WYSZECKI G., STILES W. S.: *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley, 1982.

Test Case	Gray World	Retinex	Local Shifts	Ward	Our Method
① Standard Scene					
② Arbitrary Illuminant					
③ Transmission Colour					
④ Indirect Illumination					
⑤ Type 3 Mondrian					
⑥ Type 4 Mondrian					
⑦ White World - Orange Light					
⑧ Orange World - White Light					

Table 2: Comparison of different white balance algorithms for several representative scenarios. The leftmost column shows the uncorrected images, and arrows mark instances where an algorithm fails. Arrows in brackets indicate questionable results.