

Consistent shadow values for painters

Paul Centore

Email: paul@isletech.net

A simple visual aid is presented, that can be produced in black-and-white on a piece of paper. A painter can use this aid to insure the consistency of shadow values in a painting. Consistency is necessary when an object with multiple colours, such as a shirt with red and white stripes, is seen in both light and shadow. All the values (the values of white in light, red in light, white in shadow, and red in shadow) must be correctly related to each other to make the painting convincing. Consistent values can be found by drawing a few lines on the visual aid. The scientific derivation of the visual aid is described. Finally, an application is suggested, that could heighten Adelson's "checkerblock" illusions.

Received 13 February 2014; revised 19 June 2014; accepted 06 December 2014

Published online: 27 April 2015

Introduction

To portray two objects in the same illumination, or one multi-coloured object, a realistic painter would like his lights and darks to be consistent with the common illumination. For example, suppose an artist is painting the red and white striped smokestack pictured in Figure 1. The left side of the smokestack is strongly lit, while its right side is in shadow, receiving less light. The artist convincingly paints the left side of the white stripes with a high-value neutral colour, such as Munsell N8. He convincingly paints the right side of the white stripes with a darker neutral, N5, and the left side of the red stripes with 5R 6/8. Figure 2A shows the progress so far.

The question now is how dark to paint the shadowed portion of the red stripes. This decision introduces the issue of consistency. A red shadow colour that is believable when viewed in relation to the red colour in light, might not be believable when viewed alongside the lit and shadowed whites, which are in the same lighting condition. Figures 2B and 2C show two possible shadow colours for the red. The red can be seen in different degrees of shadow, and the shadows in Figure 2C are stronger (i.e. darker) than the shadows in Figure 2B. The darker shadow in Figure 2C would occur with a very directed light source, while Figure 2B would occur under more diffuse lighting conditions. The colours are shown, away from the white and its shadow colour, at the bottom of Figures 2B and 2C.

Seen without the white stripes for a comparison, the shadow colours at the bottom of Figures 2B and 2C are both plausible. Seen in the context of the smokestack, however, in the top of Figures 2B and 2C, either shadow colour is inconsistent. The error is most obvious in Figure 2C, where the shadowed white stripes seem to glow, relative to the shadowed red stripes. This effect occurs because the human visual system cannot determine a set of lighting conditions that is consistent for both the red and the white stripes.



Figure 1: Photograph of a real-world example of consistent shadow values.

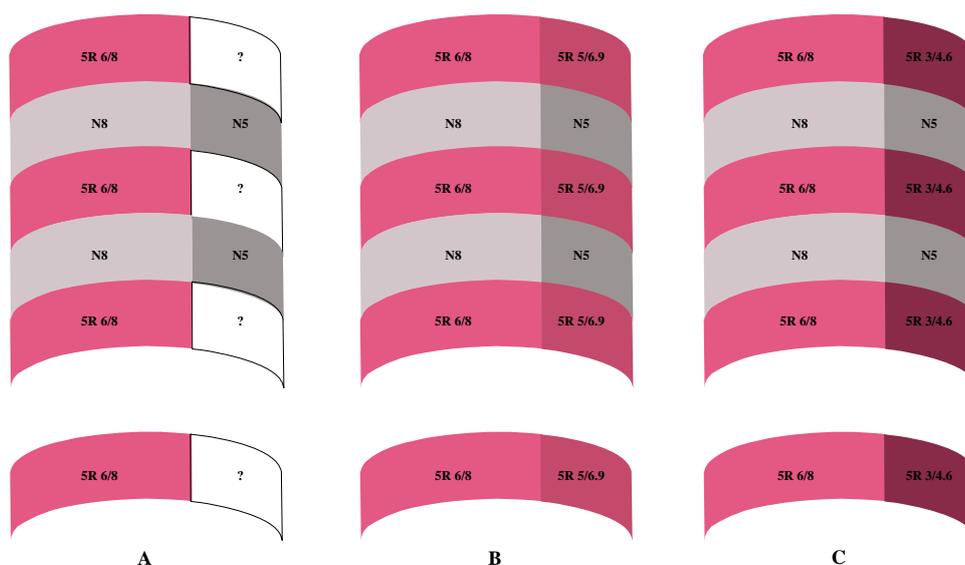


Figure 2: Choosing shadow values. A (left): just white in shadow. B (middle): lighter value for red shadow. C (right): darker value for red shadow.

The important variable in this example is Munsell value: if the value for a certain degree of shadow is decided, and the colour in light is known, then previous work showed how to find the hue and chroma for that degree of shadow [1-2]. The hue and chroma of the two red shadows in 2B and 2C were chosen so that the shadow colours would both be in the shadow series of 5R 6/8, and would have values 5 and 3. The key painting decision, then, is Munsell value.

This article presents a simple visual aid (see Figure 3) that an artist can use to determine consistent shadow values, so that different objects, or different parts of the same object, can be portrayed in the same lighting. This aid would be particularly useful when painting from imagination, or when combining studies, made in different lighting situations, into one painting. The physical and mathematical derivation of the visual aid is described. Also, a potential application is suggested, that could heighten Edward Adelson's "checkerblock" illusions.

3.0	4.0	4.5	5.0	5.5	6.0	6.5	7.0	7.5	8.0	8.5	9.0	9.5
2.8	3.8	4.2	4.7	5.2	5.7	6.1	6.6	7.1	7.6	8.1	8.5	9.0
2.6	3.5	4.0	4.4	4.9	5.3	5.8	6.2	6.7	7.1	7.6	8.1	8.5
2.4	3.3	3.7	4.1	4.6	5.0	5.4	5.8	6.3	6.7	7.1	7.6	8.0
2.2	3.0	3.5	3.9	4.3	4.7	5.1	5.5	5.9	6.3	6.7	7.1	7.5
2.0	2.8	3.2	3.6	4.0	4.3	4.7	5.1	5.5	5.8	6.2	6.6	7.0
1.8	2.6	2.9	3.3	3.7	4.0	4.3	4.7	5.1	5.4	5.8	6.1	6.5
1.6	2.3	2.7	3.0	3.3	3.7	4.0	4.3	4.7	5.0	5.3	5.7	6.0
1.4	2.1	2.4	2.7	3.0	3.3	3.7	4.0	4.3	4.6	4.9	5.2	5.5
1.2	1.8	2.1	2.4	2.7	3.0	3.3	3.6	3.9	4.1	4.4	4.7	5.0
0.9	1.5	1.8	2.1	2.4	2.7	2.9	3.2	3.5	3.7	4.0	4.2	4.5
0.7	1.3	1.5	1.8	2.1	2.3	2.6	2.8	3.0	3.3	3.5	3.8	4.0
0.4	0.7	0.9	1.2	1.4	1.6	1.8	2.0	2.2	2.4	2.6	2.8	3.0

Figure 3: Visual aid for consistent shadow values.

The Munsell system

Albert Munsell devised his colour specification system for painters and visual artists. It classifies surface colours by three perceptual attributes that are basic to painting: hue, value and chroma.

Hue is universally understood. It says whether a colour is red, yellow, purple, etc. Munsell designates 10 basic hues: R (red), YR (yellow-red, or orange), Y (yellow), GY (green-yellow), G (green), BG (blue-green), B (blue), PB (purple-blue), P (purple), and RP (red-purple).

Each basic hue is further subdivided into 4 steps, denoted with a prefix. For example, the four greens are denoted 2.5G, 5G, 7.5G, and 10G. 2.5G is a yellower green, that is closer to GY than it is to BG. 10G is a bluer green, that is closer to BG than it is to GY. In all, then, the Munsell system specifies 40 hues (4 steps for each of the 10 basic hues). These 40 hues are equally spaced perceptually. One

could interpolate any desired amount between two adjacent hues. For example, the hue 6GY is a yellowish green that is between 5GY and 7.5GY, but perceptually more similar to 5GY. White, black, and greys are not considered hues in the Munsell system. Rather, they are designated N, for “neutral.”

Many different colours can have the same hue. Figure 4, for example, shows the “hue leaf” for 6GY, a set of colours all of which have hue 6GY. The different colours within a hue leaf are specified further by value and chroma. The empty boxes indicate colours that are in the Munsell system, but that are beyond the gamut of the printing process used to produce the figure. The hue leaf shades smoothly into the neutral axis, consisting of greys, shown on the left.

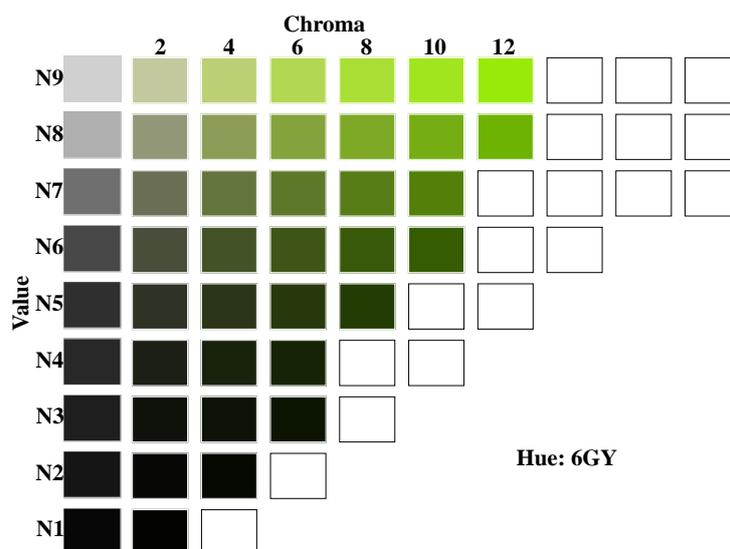


Figure 4: The hue leaf for 6GY in the Munsell System.

Munsell value designates how light or dark a colour is. The theoretically darkest black has a value of 0, and is denoted N0. The theoretically lightest white has a value of 10, and is denoted N10. Between N0 and N10 are 9 progressively lighter greys, denoted N1 through N9. The spacing between the greys is perceptually equal. All colours have a Munsell value, not just the neutrals. For example, there are light blues and dark blues. A blue with value 8.5 has the same lightness as N8.5.

Munsell chroma refers to how intense, or saturated, a colour is. For example, a lemon is an intense yellow, while masking tape is a dull yellow. A dull colour is closer to a neutral grey than an intense colour. The Munsell system denotes chroma numerically. Greys have chroma 0. A colour with a chroma of 10 or higher is generally perceived as saturated. Colours of low chroma, say 4 or less, are perceived as subdued, with a high grey content.

The Munsell notation for a colour takes the form H V/C, where H stands for hue, V stands for value, and C stands for chroma. For example, the colour 10R 9/6 would be a very light (V is 9), moderately intense (C is 6), orangish red (H is 10R). A colour with chroma 0 is a neutral grey, which is denoted NV, where V stands for value. For example, N5 is a grey that is midway between white and black.

A visual aid for consistent shadow value

Figure 3 shows a visual aid that helps a painter obtain consistent shadow values. The aid consists of rows of different sequences of greys. Each row shows how the first row would look in some degree of shadow, with darker shadows occurring nearer the bottom.

To use this aid, first determine the Munsell values of the local colours of interest, and mark those values on the top row of Figure 3. Then draw vertical lines down the figure, from the marked values. Figure 5 shows an example, in which the red and white stripes in a smokestack have local values 6 and 8, respectively, as illustrated in Figure 2. In a scene, different amounts of light would impinge on different regions of the smokestack. Suppose the lighting on the shadowed region caused the white stripes to appear to have value 5. Then the painter should identify the cell, or point within a cell, on the right-hand vertical line (which corresponds to the white stripes) that corresponds to value 5, and draw a horizontal line through it, as shown in the figure. The horizontal line in this example crosses the left-hand vertical line (which corresponds to the red stripes) at a value of 3.7, so consistency can be achieved by painting the red stripes with that value.

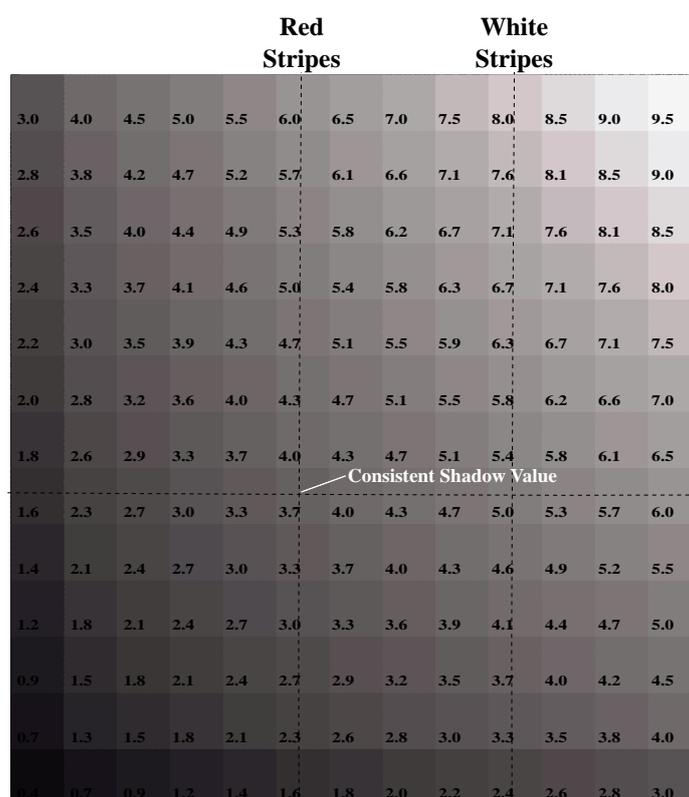


Figure 5: Example use of visual aid, for painting red and white stripes.

Once a shadow's value is known, its chroma can be determined [1]. A colour's shadow series falls approximately along a straight line drawn over the leaf for that colour's hue; the line runs from the colour itself to a point that would be two value steps below N1. Mathematically, the shadow series line gives chroma as a function of value. In the case at hand, a chroma of 5.4 corresponds to a value of 3.7. Figure 6 repeats Figure 2, with the uncertain shadow colours filled in as 5R 3.7/5.4. The shadows in 2A have consistent values: the red and white stripes appear to be in the same degree of shadow (on the right side), and at the same level of illumination (on the left side). Figure 6A should be compared with Figures 6B and 6C, where it is difficult to interpret the stripes in shadow and in illumination.

In theory, the visual aid applies to any situation in which two different local colours are lit the same way, which occurs very frequently. In practice, the aid would mainly be used when adjacent local colours are significantly different, or when a painter is experiencing difficulty. In this regard, the shadow value aid is much like perspective. Perspective can be applied to any realistic paintings, even

portraits, and a painter could spend hours making perspective constructions, even for simple scenes. In practice, however, perspective constructions are only made for complicated architectural scenes, or when a painter encounters difficulty in simpler scenes. Likewise, the shadow value aid would probably be used only occasionally, when the lighting in a painting does not seem convincing. For example, a portrait might depict a fair-skinned model with medium brown hair, where the skin and the hair are both partially in shadow. If the hair in shadow seems unrealistically light or dark, the visual aid could help adjust its value.

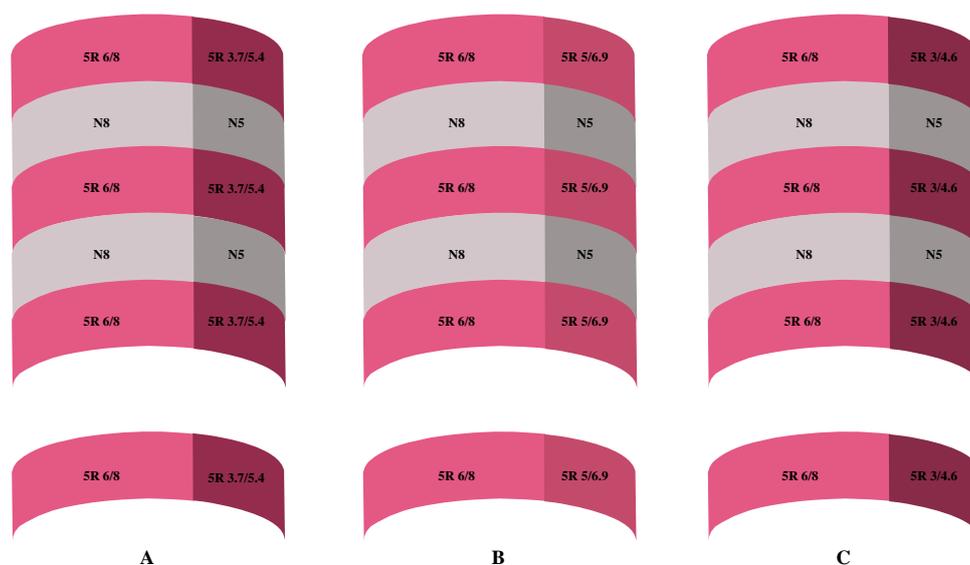


Figure 6: Shadows from visual aid. A (left): consistent shadow values. B (middle); red shadow values too light (same as Figure 2B); C (right): red shadow values too dark (same as Figure 2C).

Derivation of the shadow value aid

The aid was developed mathematically, from calculations involving Munsell value V ; and luminance factor Y : The conversion between them is given by a quintic polynomial that appeared originally in the Munsell renotation [3], and in a rescaled version in an ASTM standard [4].

Every point of a scene has a luminance factor, which measures the amount of light from that point (weighted with respect to human sensitivity to different wavelengths), which impinges on a painter's eye. Now suppose there is a painting of the scene, and the painting is evenly illuminated by Illuminant C. Every point of the painting has a Munsell value. The Munsell value, and indeed the entire Munsell specification, depends only on the reflectance characteristics of the paint at a particular point. In a realistic painting, each point on the canvas corresponds to a particular point in the scene. The relationship between the luminance factor and Munsell value at corresponding points is not completely understood.

A painter can "key" the values, within some limits, and still produce a convincing painting, much as a photographer has some latitude with exposure time. Nevertheless, there must be some consistency: darker parts of the scene should be painted darker than lighter parts of the scene, and by the same relative amount.

The luminance factor of a part of an object in a scene depends on two other factors: the reflectance properties of that part, and the illumination impinging on that part. A shadow occurs when there is

less illumination on some parts of the object than on others. In the smokestack, for example, there is less illumination on the right side than on the left side. Formally, the light side of the smokestack might be receiving light whose spectral power distribution (SPD) is given by $E(\lambda)$, while the shadowed side is receiving light whose SPD is $kE(\lambda)$, where k is positive but less than 1. The reflectance properties of the part in question are described by a reflectance function, $\rho(\lambda)$, where λ is a wavelength in the visible spectrum, which stretches approximately from 400 nm to 700 nm. ρ can take on values in the interval from 0 to 100%.

The luminance factor Y_L for the part in light is given by:

$$Y_L = \frac{\int_{400}^{700} \bar{y}(\lambda)\rho(\lambda)E(\lambda)d\lambda}{\int_{400}^{700} \bar{y}(\lambda)E(\lambda)d\lambda} \quad (1)$$

where \bar{y} is the photopic luminous efficiency function [5]. The luminance factor Y_S for the part in shadow is given by

$$Y_S = \frac{\int_{400}^{700} \bar{y}(\lambda)\rho(\lambda)kE(\lambda)d\lambda}{\int_{400}^{700} \bar{y}(\lambda)E(\lambda)d\lambda} \quad (2)$$

From Equations 1 and 2, it follows that:

$$Y_S = kY_L \quad (3)$$

A notable fact is that Equation 3 holds independently of the reflectance function $\rho(\lambda)$: in the smokestack example, the red and white stripes have very different reflectance functions. Despite this difference, Equation 3 applies to both of them: the ratio of the luminance factor in shadow, to the luminance factor in light, is the same k , for both the red and the white stripes. Physically, then, shadow values can be achieved by multiplying the luminance factor in light by the same constant, for all colours in the same light and shadow configuration.

Perceptually, however, multiplication by the same constant is not correct. The reason is that the perception of lightness, which is codified in the Munsell value, is not a linear function of the luminance factor. Instead, it is given by the quintic equation [4]:

$$Y = 0.00081939 V^5 - 0.020484 V^4 + 0.23352 V^3 - 0.22533 V^2 + 1.1914 V \quad (4)$$

where V is the Munsell value and Y is expressed as a percentage. Munsell values must be converted to luminance factors, after which multiplicative constants can be determined and applied. Then the resulting luminance factors can be converted back to Munsell values.

In the example in Figure 2, the white stripes in light and shadow have Munsell values 8 and 5, respectively. These correspond to luminance factors of 57.6% and 19.3%. The ratio, k , of these luminance factors is 19.3/57.6, which is about 0.34. The red stripe in light has a Munsell value of 6, which corresponds to a luminance factor of 29.3%. The luminance factor of the red stripe in shadow should therefore be 0.34 times 29.3%, which is about 10.0%. Inverting Equation 4 converts a luminance factor of 10.0% to a Munsell value of 3.7. The red shadow, then, should be painted with

Munsell value 3.7. Figure 5, of course, obtained this same result geometrically, freeing a painter from these computations.

The calculations behind the visual aid should now be clear. The first row of Figure 3 is a selection of Munsell values for colours in light. The second row is the Munsell values of the first row, when viewed in shadow defined by a constant k . The second row was scaled by requiring the cell on the far right to have Munsell value 9.0. Since the cell on the far right in the first row has Munsell value 9.5, the method of the previous paragraph gives a k of 0.87. The luminance factor of each cell in the second row is therefore 0.87 times the luminance factor of the cell above it in the first row. The luminance factors have been converted back to Munsell values, which are shown on the visual aid. Further rows were calculated by the same method, after choosing a set of values for the far right cells, that should provide a fine enough resolution for painting. Interpolation might be required when painters use the visual aid. Since a difference of less than 0.25 Munsell value steps is usually indistinguishable, however, and differences of less than 0.5 Munsell value steps are difficult for painters to mix, the aid should be sufficiently accurate.

The foregoing derivation makes some simplifying assumptions that might not hold perfectly in practice. First, the very use of the Munsell renotation implies that the ambient light has an SPD consistent with Illuminant C. Illuminant C was chosen as an average SPD for indirect daylight, and in fact the human assessments behind the renotation were made in indirect daylight, rather than under controlled illumination. While light in realistic settings can vary considerably from Illuminant C, it is also true that chromatic adaptation [6] makes local colours appear approximately constant, regardless of the lighting. For this reason, the Munsell system is accurate enough for painting. Second, it is assumed that the SPD of the light impinging on a shadowed area is a constant k times the SPD of the light impinging on a lit area. Possibly, however, these two lights have different relative SPDs. For example, the light on the shadowed areas might have been filtered naturally, say by a canopy of green leaves. While such changes would mainly affect hue and chroma, they might also have some effect on value, so the assumption of the same relative SPD is sometimes only an approximation. Thirdly, the derivation of the shadow value aid is theoretical, rather than empirical. While the theoretical approach is simple and reasonable, it is always possible that real-world factors have been overlooked, which would alter the results. While acknowledging the simplifications and assumptions behind the shadow value aid, it is nevertheless suggested as a helpful practical tool for painters, in some situations.

Applications to Adelson's checkerblocks

In 1993, Edward Adelson introduced his well-known "checkerblock" illusions [7], an example of which is shown in Figure 7. The illusion is that diamond A appears darker than diamonds C and E, when in fact they all have the same lightness. A plausible explanation is that the viewer interprets the figure as a piece of a 2-by-2 checkerboard, that has been extended into the third dimension. The diamonds are therefore not seen as flat shapes on a two-dimensional viewing surface, but as parts of a solid that exhibits shadowing (Figure 11.1(c) from a paper of Adelson and Pentland [8] shows a possible arrangement of shadows.). The viewer discounts the shadowing, in order to arrive at local colours. The simplest conclusion is that diamonds D, E, and G are facets of one cube, which has a uniform light colour, while A and B are facets of another cube, which has a uniform dark colour. Diamond E is therefore perceived as lighter than diamond A, even though their lightnesses are equal.

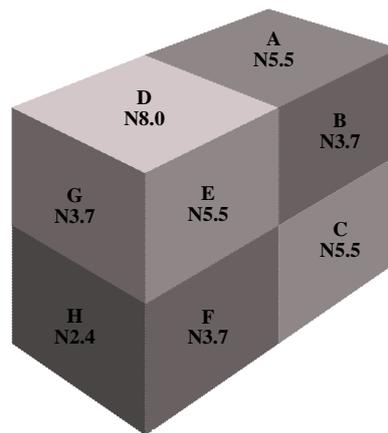


Figure 7: Checkerblock illusion.

A reasonable conjecture is that the checkerblock illusion would be most effective if the lightnesses of the diamonds were chosen to give consistent shadow values. This section finds the general form for such lightnesses, and could be used in testing the conjecture empirically.

The lightness relations depend on the shadow relations. For example, diamonds B and E are shadowed versions of A and D, respectively. In most realistic situations, the amount of light hitting one face will not vary appreciably, so B and E should have the same degree of shadow, relative to A and D. If the factor k_1 corresponds to this degree of shadow, then:

$$Y_B = k_1 Y_A \quad (5)$$

$$Y_E = k_1 Y_D \quad (6)$$

Similarly, diamonds G and H are shadowed versions of E and F, respectively, so:

$$Y_G = k_2 Y_E \quad (7)$$

$$Y_H = k_2 Y_F \quad (8)$$

The original illusion had $Y_A = Y_C$: Since C and E are presumably the same colour, we therefore have:

$$Y_A = Y_E = Y_C \quad (9)$$

To make the illusion simpler and even more compelling, we can also require G and F to have the same lightness, getting:

$$Y_G = Y_F = Y_B \quad (10)$$

Equations 5 through 10 can be combined to give a set of relationships that depend only on Y_A and Y_D :

$$Y_C = Y_E = Y_A \quad (11)$$

$$Y_B = Y_F = Y_G = \frac{Y_A^2}{Y_D} \quad (12)$$

$$Y_H = \frac{Y_A^3}{Y_D^2} \quad (13)$$

Given initial choices for Y_A and Y_D , Equations 11 through 13 can be used to determine all the lightnesses in a checkerblock. The lightnesses can be expressed as either luminance factors or Munsell values. Figure 7 shows the checkerblock that results when Munsell values of 5.5 and 8.0 are chosen for diamonds A and D , respectively. The colours are all chosen to be neutral greys.

It is unclear from the experimenter's account what values were used for experiments with the original checkerblock. The illusions were originally displayed on computer screens [7]. The luminances for some other illusions are given in milliLamberts. No mention, however, is made of the ambient viewing conditions, so luminance factors cannot be calculated. As Stevens and Stevens point out [9], perceived luminance differences vary when the eye is adapted to different viewing conditions. The quintic relationship given by Equation 4 is valid only when the lighting level is approximately that of indirect daylight. The perceived shadow relationships in the same illusion will therefore be different when viewed in a dim room and in a well-lit room. The instructions given here are intended to produce illusions that are effective when printed or painted, and viewed under everyday working conditions. The consistent shadow values should heighten the checkerblock illusion by making the three-dimensionality more convincing.

References

1. Centore P (2013), Shadow colours for painters, *Journal of the International Colour Association*, **11**, 18-27.
2. Centore P (2013), Shadow series in the Munsell system, *Color Research and Application*, **38** (1), 58-64.
3. Newhall S, Nickerson D and Judd DB (1943), Final report of the O.S.A. Subcommittee on the spacing of the Munsell colors, *Journal of the Optical Society of America*, **33** (7), 385-418.
4. ASTM (2008), *ASTM Standard D 1535-08: Standard Practice for Specifying Color by the Munsell System*, ASTM International.
5. CIE (2004), CIE Publication No. 15:2004: *Colorimetry* (Third Edition), CIE.
6. Fairchild MD (2005), Chromatic adaptation, in *Color Appearance Models* (Second Edition), John Wiley & Sons, 146-165.
7. Adelson EH (1993), Perceptual organization and the judgment of brightness, *Science*, **262**, 2042-2044.
8. Adelson EH and Pentland AP (1996), The perception of shading and reflectance, in *Perception as Bayesian Inference*, Knill D and Richards W (eds.), Cambridge University Press, 409-423.
9. Stevens JC and Stevens SS (1963), Brightness function: Effects of adaptation, *Journal of the Optical Society of America*, **53** (3), 375-385.