Shi, L., and Funt, B., "Monge's Unique Property of White Surfaces," in *AIC'2005 Proc. 11th Congress of the International Color Association*, Sydney, Sept. 2009.

MONGE'S UNIQUE PROPERTY OF WHITE SURFACES

Lilong Shi, Brian Funt
Simon Fraser University

ABSTRACT

As hypothesized by Gaspard Monge in the eighteenth century, white surfaces have a unique property that distinguishes them from other coloured specular surfaces. We test his hypothesis using modern digital imagery and find it to hold. We then incorporate it into an algorithm that identifies white surfaces when the illumination colour is unknown. Since white surfaces reflect the colour of the incident illumination, the method also has application to the problem of illumination estimation for colour constancy.

Keywords: Monge, specularities, illumination estimation, dichromatic

CONTACT

lshia@cs.sfu.ca

INTRODUCTION

John Mollon devoted his 2005 Verriest Lecture[1] to a discussion of the insights of Gaspard Monge. Monge had presented his ideas in a lecture to the Academy of Sciences in 1789. Many of these insights prefigure what Shafer[2] formalized as the dichromatic model of specular reflection in which specular reflection is analysed in terms of a body reflection component and an interface reflection component. In the context of specular reflection, one of Monge's insights that Mollon describes is that a white surface has the unique property that it necessarily has no variation in chromaticity across it. This property holds even when the illumination is not white, but coloured, and hence can be used to identify a white surface under illumination of unknown colour.

Mollon[1] (page 302) describes Monge's insight[3] as follows:

> He realized that a white surface has a unique property: provided there is only one source of illumination (and no secondary reflections from colored surfaces), there can be no variation in chromaticity across a white surface, because the chromaticity of the body color and the chromaticity of the illuminant are the same. We may add that this is also true for grey surfaces. In both cases, although there is no variation in chromaticity, there may well be variation in luminance across the surface. If this is how we identify a white surface, then—Monge argues—we can explain the whitening of a red object seen through a red filter. When the red surface is observed through the red glass, it has the property of a white object (i.e., there is no variation in chromaticity across its surface). The specular component reflected from any point on the surface has now the same color as the component that has undergone selective absorption. The same result will occur if, instead of viewing through a chromatic filter, we illuminate a scene with one predominant color: "For these homogeneous rays, being reflected to the eye from all the visible parts of the surface of colored objects, as is white light under ordinary conditions, we are led to take them for the white rays whose function they now perform, and thus to consider as white all those objects that reflect to the eye only rays of this type." (Monge, 1789, pp. 144–145.)

Intrigued by Monge's hypothesis that a white surface has such a unique property, we decided first to test his hypothesis, and then second to exploit it in an algorithm designed to distinguish white surfaces from other coloured surfaces in images of scenes containing both white and non-white specular surfaces. Although as Mollon points out, the property should hold for any achromatic surface, for simplicity we will simply refer to all of them as 'white'. While the problem of identifying white surfaces is of intrinsic interest, their reliable identification would also be very useful in illumination estimation and colour constancy. In addition to testing the Monge hypothesis directly by measuring the variation in the chromaticity of specular surfaces of various colours, and employing it an algorithm that distinguishes white surfaces from those of other colours, we also provide some initial results on using it for illumination estimation.

BACKGROUND

Shafer's dichromatic reflection model for inhomogeneous dielectric objects states that the colour signal is a linear combination of two components, one being associated with the interface reflection and the other describing the body reflection part[2]. This is expressed as

$$C(\theta,\lambda) = m_i(\theta)C_i(\lambda) + m_b(\theta)C_b(\lambda) \tag{1}$$

where $C_i(\lambda)$ and $C_b(\lambda)$ are the spectral power distributions of the interface and the body reflection respectively, and m_i and m_b are the corresponding weighting factors which depend on the geometry θ, which includes the incident angle of the light, the viewing angle, and the phase angle.

Suppose R, G, and B are the red, green, and blue pixel value outputs of a digital camera, then each color vector $[R,G,B]^T$ is determined by a linear combination of a surface reflection component $[R_i,, G_i,, B_i]^T$ and a body reflection $[R_b, G_b, B_b]^T$ component. Equation 2 shows that the resulting $[R,G,B]^T$ colour can be expressed as the weighted sum of these two reflectance components. Thus the colours from a single specular surface must lie in a plane.

$$[R,G,B]^T = w_i[R_i, G_i, B_i]^T + w_b[R_b, G_b, B_b]^T \tag{2}$$

Given two specular surfaces under the same illumination then there must be two such RGB planes. Both planes, however, contain the same illuminant RGB. This implies that their intersection must represent the colour of the illuminant itself. Based on this observation, Lee[4] introduced a method for computing the chromaticity of the scene illuminant from the distribution of chromaticities observed from two or more coloured specular surfaces. Lee's method is based on the fact that dicrhromatic planes in 3D colour space project to lines in 2D chromaticity space. Just as two dichromatic planes intersect along a line representing the colour of illumination, in chromaticity space, the lines from two dichromatic planes intersect at a point representing the chromaticity of the illumination. Variations on Lee's basic method have led to a significant performance improvement[5,6,7].

TESTING THE MONGE HYPOTHESIS

To test the Monge's hypothesis about the unique property of white, we use images of a scene containing both white and non-white surfaces taken under several different illuminants. Fig 1(a) shows two juice containers, one orange and the other white. Fig 1(b) and fig 1(c) show the containers under 2 different illuminants with the white and orange regions within each image segmented out by hand. The *rg*-chromaticities ($r=R/(R+G+B)$, $g=G/(R+G+B)$) of the pixels were then histogrammed in a histogram of 64x64 bins. Fig 1(e) and fig 1(f) show the histograms obtained under the 2 illuminants. In both cases, the chromaticities from the white region form a narrow peak created by the white surface.

As a quantitative test of Monge's hypothesis, table 1 lists the sum of the variances of the *r* and *g* chromaticities from the white versus orange regions under 8 different illuminants. Table 1 also includes the variances of the grey versus red surfaces in Grey Bag Scene of fig 2, and the white versus green surfaces of the Flower Scene shown in fig. 1(d).

(a) Containers Scene

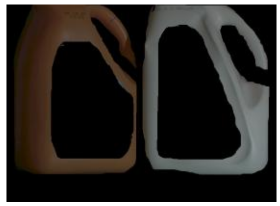

(b)

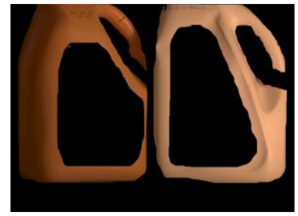

(c)

(d) Flower Scene

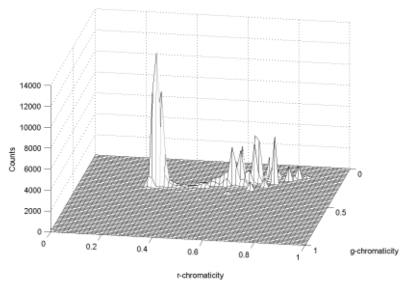

(e)

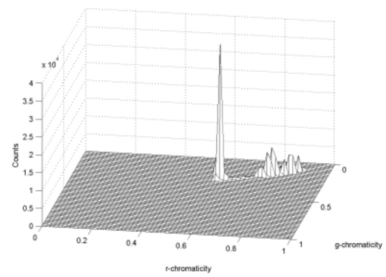

(f)

Fig 1. Top row: (a) Original unsegmented containers scene under one of the illuminants; (b) containers under bluish illuminant; (c) containers under reddish illuminant. In (b) and (c) black regions are excluded by hand from further processing so that only pixels from the orange and white surfaces are histogrammed. Bottom row: (e) histogram of counts of rg-chromaticities (r-chromaticity left to right, g-chromaticity back to front) corresponding to (b); (f) histogram corresponding to (c). As predicted by Monge's hypothesis, independent of the illumination colour, the white surface generates the tall narrow peaks. The orange surface generates a broader distribution of chromaticity values. Results for (d) "Flower Scene" (from Corel Gallery Copyright © 1999 Corel Corporation) are in Table 1.

Day

CWF

U30/TL84

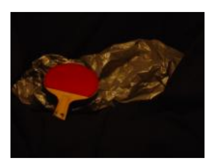

A

Fig 2. Grey Bag Scene under 4 different illuminants. This scene is used to test the Monge hypothesis with an achromatic (grey not white) surface. The variances of chromaticities from the grey and red surfaces from hand-segmented regions are listed in Table 1.

IDENTIFYING WHITE SURFACES USING MONGE'S HYPOTHESIS

In the above tests of the Monge hypothesis, the regions were hand-segmented. If the regions could be segmented automatically then it would be a simple matter to measure the variance of each region and classify them as white or non-white; however, as is well known, region segmentation is not an easy problem. Some previous work on the analysis of specularities has avoided the need for pre-segmentation either by the use of Hough transforms[6], by looking for a global fit for intersecting planes[7,9], or by using dichromatic planes generated from image subwindows[5]. In this paper, we avoid pre-segmentation by using a Hough transform[10], in fact, two Hough transforms in sequence.

The key to avoiding the need for region segmentation is to determine the dichromatic planes based on using all image pixels and their distribution in colour space. Each RGB colour defines a set of dichromatic planes on which it could lie. If we assume that there is only a single source of illumination, then these planes are constrained to pass through both the RGB and the origin (i.e, black), and so must contain the line defined by the RGB and the origin. Using a Hough transform of RGBs to candidate planes, the likelihood of each dichromatic plane occuring in the image data is represented by its count in Hough space. For each candidate plane, the variances of the *r* and *g* chromaticities of the pixels that contributed to each plane are computed and summed. A dichromatic plane with small combined *r* and *g* variance should correspond to a white surface; one with a large surface to a non-white specular surface. Monge's argument was that a white surface implies a small amount of colour variation. Although Monge's white-surface property is a sufficient, not necessary condition, we are now trying it in the reverse direction; namely, that a small amount of colour variation may imply a white surface.

Table 1. Total variance of chromaticities obtained from white (achromatic) versus non-white (chromatic) specular surfaces under a variety of illuminants. In agreement with the Monge hypothesis, the variance in the chromaticity of the pixels from the white surface remains low in comparison to that of the non-white surface even though the mean chromaticity of the white surface varies significantly with the choice of illuminant. The 11 images of the containers scene are from the SFU on-line colour constancy dataset[8].

Scene	**Illuminant**	**Achromatic surface**		**Chromatic surface**	
		Mean chromaticity	**Variance x10^3**	**Mean chromaticity**	**Variance x10^3**
Containers scene under 11 illuminants from the SFU dataset	1	(0.49,0.38,0.13)	0.56	(0.74,0.23,0.03)	4.99
	2	(0.48,0.36,0.16)	1.70	(0.74,0.22,0.04)	7.36
	3	(0.35,0.37,0.29)	0.89	(0.67,0.25,0.08)	8.93
	4	(0.43,0.37,0.20)	1.16	(0.71,0.23,0.05)	7.76
	5	(0.32,0.35,0.33)	1.22	(0.65,0.26,0.09)	11.74
	6	(0.40,0.37,0.22)	1.10	(0.70,0.24,0.06)	7.92
	7	(0.31,0.34,0.35)	3.33	(0.63,0.26,0.10)	22.88
	8	(0.52,0.34,0.14)	1.11	(0.77,0.20,0.03)	5.70
	9	(0.39,0.36,0.25)	0.99	(0.69,0.24,0.07)	7.18
	10	(0.46,0.34,0.20)	0.41	(0.72,0.23,0.05)	4.46
	11	(0.60,0.28,0.13)	0.43	(0.81,0.17,0.02)	3.80
Grey Bag	Day	(0.28,0.34,0.37)	2.10	(0.76,0.09,0.14)	8.75
	CWF	(0.38,0.37,0.25)	1.57	(0.66,0.20,0.14)	11.09
	U30/TL84	(0.39,0.35,0.26)	1.81	(0.62,0.21,0.16)	10.37
	A	(0.56,0.32,0.12)	2.94	(0.90,0.07,0.03)	9.33
Flower	N/A	(0.33,0.34,0.33)	0.02	(0.21,0.58,0.21)	25.87

The information about the variance of the colours from each possible dichromatic plane occurring in the input image is spread throughout the Hough space. To accumulate the evidence about the variance, a second Hough transform is used. The cells of this second Hough space represent chromaticities. In the second transform, each plane (i.e., Hough cell from the first transform) votes for all the chromaticies that contributed to it. This vote is the inverse of the variance scaled by the square of the number of chromaticites it represents. In other words, Hough cells with small counts contribute very little to the final result. The scaling is effectively a thresholding step, but with a soft boundary. In the completed second transform, the Hough cell with the highest score represents the chromaticity of the white surface. Note that because the illumination is not necessarily white, the white surface chromaticity similarly will not necessarily be white (i.e., may not have $r = g$) either.

IMPLEMENTATION DETAILS

Under the assumption that there is only a single illuminant spectrum lighting the scene, all dichromatic planes must pass through the origin. In this case, the planes can be parameterized by the two angles ϕ and θ as:

$$D(\phi,\theta,r,g,b) \Leftrightarrow R\cdot\cos(\phi)\cos(\theta)+G\cdot\sin(\phi)\cos(\theta)+B\cdot\sin(\theta)=0 \quad (3)$$

All pixels from the same surface belong to a single plane defined by angle ϕ relative to the z-axis, and angle θ relative to the y-axis. The distance D indicates the closeness of a given RGB to the plane specified by (ϕ,θ). In the discrete case, the parameter space (ϕ,θ) is quantized into bins, so the Hough Transform[11] is represented as a two-dimensional "histogram" $\mathbf{H_1}(\phi,\theta)$ of potential dichromatic planes. The "counts" for the bins $\mathbf{H_1}(\phi,\theta)$ are computed as follows. First, find the set of pixels **P** with chromaticity lying on or near the dichromatic plane specified by (ϕ,θ).

$$\mathbf{P}=\{\forall <r,g,b>\mid D(\phi,\theta,r,g,b)<\xi\} \quad (4)$$

The "count" for bin $\mathbf{H_1}(\phi,\theta)$ is set as

$$\mathbf{H}_1(\phi,\theta)=(var(\mathbf{P}_r)+var(\mathbf{P}_g))^{-1}\cdot|\mathbf{P}|^2 \quad (5)$$

where *var()* is the variance function applied to the chromaticities in **P**, and $|\mathbf{P}|$ is the cardinality of **P**. The effect of the scaling by $|\mathbf{P}|^2$ is to render insignificant the contributions of planes created by only a small number of pixels.

A high value of $\mathbf{H_1}(\phi,\theta)$ indicates a dichromatic plane with low variance, and hence, following Monge, a high likelihood of emanating from a white surface. The next step is to accumulate the evidence that is spread across multiple dichromatic planes in order to establish what the chromaticity of the white surface as it appears in the image is. Of course, a white surface will match the chromaticity of the illuminant. All dichromatic planes generated from a white surface must intersect along the direction of the illuminant colour; hence, that colour vector must be perpendicular to the normal of each of those dichromatic planes. The vector perpendicular to the normals of the hightest number of such dichromatic planes should, therefore, indicate the chromaticity of the illumination, and hence white surface as it appears in the image.

To determine this illumination vector, we use a second Hough Transform to create an *illumination histogram* $\mathbf{H}_2(\alpha,\beta)$. Let $\boldsymbol{n}=(u,v,w)$ be the normal of dichromatic plane (ϕ,θ) and the illumination vector be represented by angles α and β. For small ξ, they are (approximately) perpendicular when

$$u\cdot\cos(\beta)\cos(\alpha)+v\cdot\sin(\beta)\cos(\alpha)+w\cdot\sin(\alpha)<\xi \quad (6)$$

The second histogram is formed as

$$\mathbf{H}_2(\alpha,\beta)=\sum\mathbf{Q} \quad (7)$$

where $\mathbf{Q}=\{\forall\mathbf{H_1}(\phi,\theta)\mid u\cdot\cos(\beta)\cos(\alpha)+v\cdot\sin(\beta)\cos(\alpha)+w\cdot\sin(\alpha)<\xi\}$.

Then, when the normal of dichromatic plane (ϕ,θ) is perpendicular to illumination axis (α,β), the score from the corresponding bin of $\mathbf{H_1}$ is added to that of the corresponding bin of $\mathbf{H_2}$. The value of $\mathbf{H_2}(\alpha,\beta)$ represents the number of dichromatic planes intersecting at (α,β) weighted by their respective variance measures from Eqn 5. The value of (α,β) at which $\mathbf{H_2}(\alpha,\beta)$ is the largest represents the chromaticity of the white surface (or equivalently, the colour of the illumination).

RESULTS IDENTIFYING WHITE SURFACES AND THE ILLUMINATION CHROMATICITY

The procedure described above identifies the chromaticity of white surfaces illuminated by light of unknown colour. Once the chromaticity of the white surfaces has been identified, it can be used to label them in the image. Fig 3 presents this labelling as a pseudo-coloured image encoding the distance between the chromaticity of the predicted white and each pixel's chromaticity. The "whiteness" map is pseudo-colored by representing the chromatic distance from red to blue, where red means short distance yet blue means long distance.

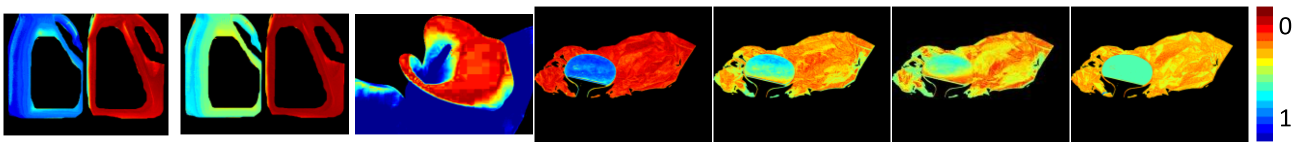

Fig 3. Results for the identification of white surfaces for the images from fig. 1(b)-(d) and fig. 2(a)-(d) represented by a pseudo-colouring of the distance between each pixel's chromaticity and that of the estimated white. The colouring is from red (short distance) to blue (long distance) as shown in the color bar on the right.

Since the chromaticity of white surfaces is the same as that of the illuminant chromaticity, white-surface identification can also serve as an illumination-estimation method for colour constancy. Table 2 compares the Monge-based estimate of white to the illumination estimates of some standard illumination-estimation methods. Although the results are quite good in this case, the Monge-based method cannot be expected to succeed as an illumination-estimation method in all circumstances since it requires that the scene contain specular surfaces and not any large matte surfaces. The Monge hypothesis is most likely to best used in conjunction with other illumination-estimation methods.

Table 2. Angular error between the estimated (r,g,b) chromaticity of the illuminant and its chromaticity as measured from a calibration target for the Gray World12, Shades of Gray12, Gray Edge13, and the proposed method based on Monge's hypothesis for 850 images from the SFU data set. Of the 900 images in the dataset, 50 were excluded because their resolution was insufficient (less than 107,000 total pixels).

Method	Angular Error		
	Median	Mean	Max
Gray World	4.70	5.89	35.18
Shades of Gray	3.40	4.28	22.13
Gray Edge	4.06	6.26	35.80
Monge White	2.78	4.79	35.87

CONCLUSION

As Mollon[1] pointed out, in 1789 Monge made some very interesting observations about the colour of specular reflections. Our experiments show that his hypothesis as to the unique properties of white surfaces does in fact hold and, furthermore, that it can be used both to identify white surfaces as well as the chromaticity of the scene illumination.

REFERENCES

1. J. Mollon, “Monge: the Verriest lecture, Lyon, July 2005.” *Visual Neurosci.* **23**(3-4), pp. 297-309, 2006
2. S.A. Shafer, “Using Color to Separate Reflection Components.” *COLOR research and application*, **Vol. 10**, No. 4, pp. 210-218, 1985
3. G. Monge, *Mémoire sur quelques phénomènes de la visión*, Annl. 3, de Chimie. 131–147, 1789.
4. H. Lee, “Method for computing the scene-illuminant chromaticity from specular highlights.” *J. Opt. Soc. Am.* A3, pp. 1694-1699, 1986
5. G. Schaefer, “Robust dichromatic colour constancy.” *Springer Lecture Notes in Computer Science. LNCS* **Vol. 3212**, pp. 257-264, 2004
6. L. Shi, and B. Funt, “Dichromatic Illumination Estimation via Hough Transforms in 3D.” In *Proc. CGIV2008 IS&T Fourth European Conf. on Colour in Graphics, Imaging and Vision,* Barcelona, 2008
7. J. Toro, and B. Funt, “A Multilinear Constraint on Dichromatic Planes for Illumination Estimation.” *IEEE Transactions on Image Processing*, **Vol. 16**, No. 1, pp. 92-97, 2007
8. K. Barnard, L. Martin, A. Coath and B. Funt, “A Comparison of Computational Colour Constancy Algorithms. Part Two: Experiments on Image Data.” *IEEE Transactions on Image Processing*, **No. 11**, pp. 985-996, 2002
9. J. Toro, D. Ziou, and B. Funt, “Illuminant Estimation from Dichromatic Planes.” In *Proceedings of the Thirteenth Color Imaging Conference: Color Science, Systems & Applications*, Scottsdale, AZ, pp.95-98, 2005
10. R. O. Duda, and P. E. Hart, “Use of the Hough Transformation to Detect Lines and Curves in Pictures.” *Comm. ACM*, **Vol. 15**, pp. 11–15, 1972
11. K. Okada, S. Kagami, M. Inaba, H. Inoue, “Plane segment finder: algorithm, implementation and applications.” In *Proc. of Int. Conf. on Robotics and Automation*, pp. 2120–2125, 2001
12. G.D. Finlayson, E. and Trezzi, “Shades of Gray and Colour Constancy.” In *Proceedings of 12th Color Imaging Conference*, pp. 37-41, 2004
13. J. Weijer, T. Gevers, A. Gijsenij, “Edge-Based Color Constancy.” *IEEE Transactions on Image Processing* 16(9): pp. 2207-2214, 2007