Prediction of transparency perception based on cone-excitation ratios

Caterina Ripamonti

Department of Psychology, University of Pennsylvania, Philadelphia, Pennsylvania 19104

Stephen Westland

Colour & Imaging Institute, Derby University, Derby DE22 3HL, UK

Received September 10, 2002; revised manuscript received April 23, 2003; accepted April 25, 2003

Perceptual transparency was measured in two experiments by using simulations of illuminated surfaces presented on a CRT monitor. In a two-alternative forced-choice paradigm, observers viewed two simulated Mondrians in temporal sequence. In one sequence the Mondrian was simulated to be partially covered by a transparent filter; in the other sequence the filter color over each Mondrian patch was modified. Observers were instructed to select the sequence containing a transparent filter. Observers' selections corresponded to sequences in which the cone-excitation ratios for each adjacent pair of Mondrian patches were approximately the same as the cone-excitation ratios for the pair of patches covered by a filter. The results suggest that cone-excitation ratios may be a cue for perceptual transparency. © 2003 Optical Society of America

OCIS codes: 330.7310, 330.1690, 330.1720, 330.4060, 330.5510.

1. INTRODUCTION

In this study we conduct two psychophysical experiments to investigate whether the invariance of spatial cone-excitation ratios predicts the perception of transparency. In one experiment we test whether the invariance of cone-excitation ratios could be used as a cue for the perception of transparency by presenting images for which the invariance of cone-excitation ratios is deliberately violated for all three cone classes. In a second experiment, cone-excitation invariance is manipulated for individual cone classes, for pairs, or for all cone classes to test the effect of ratio-invariance violations for individual cone classes. In the following introduction we describe the nature of transparency perception and then describe the cone-excitation, invariant-ratios hypothesis.

Perceptual transparency is the phenomenon of seeing one surface behind another. For example, in Fig. 1, four opaque areas give rise to the perception of two opaque surfaces (large rectangles) seen behind a transparent filter (small rectangles). The question of whether it is possible to perceive both the color of the filter and the color of the opaque surface in the overlapping region has been a matter of controversy since the nineteenth century, when Helmholtz described the perception of transparency as "seeing through." Also, Koffka stated that the filtered region is the area in which we can simultaneously perceive both the filter and the opaque surface behind the filter.[1,2] Moreover, according to Morinaga *et al.*, we see two surface colors, one behind the other in the filtered region.[3] Hering had denied the possibility of perceiving one color behind another and stated that the light reflected by two different colors and projected onto the same retinal region gives rise to one color, called fusion color. However, recovering the color of two surfaces from one set of strictly local cone excitations would seem to be intractable. Furthermore, if an opaque surface is partially covered by a filter whose color is complementary (for example, a red surface and a green filter), the filtered patch will appear very dark and the red and green surfaces will not seem to be simultaneously observed.[4,5]

Our work is not concerned directly with the nature of transparency perception but rather with the stimulus constraints that are required for transparency perception to occur. It has been conjectured that the presence of X junctions formed by the junctions of borders of opaque surfaces and the transparent medium at the overlapping area is a necessary element in the image.[6–8] An X-junction requires a minimum of four areas, but other studies have shown that transparency perception can occur even with only three areas.[9] Moreover, Khang and Zaidi have shown that in spite of spatially rotating the filtered area such that the figural unity between the opaque surfaces and their corresponding filtered surfaces would be destroyed, observers are still able to identify filters across different illuminants almost as well as they can discriminate them under the same illuminant.[10]

Luminance relationships may be necessary for the perception of transparency in achromatic images, but they are not sufficient when the objects in the image do not have the same chromaticities. In fact, it has been shown that the perception of transparency holds even when opaque and filtered surfaces have identical luminance and differ only in their chromaticities.[11,12] D'Zmura *et al.* considered the overall effect of luminance and chromatic constraints and suggested a constraint on some tristimulus representation of the colors at the X junctions. They showed evidence that observers are able to adjust the color of a filtered surface to make the central region appear transparent and that the colors of the filtered regions then converge toward a point in color space.[11] For

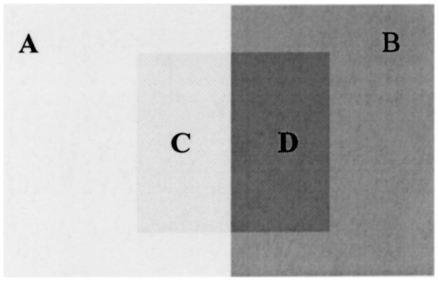

Fig. 1. Opaque areas A and B are partially covered by a single filter, giving rise to the filtered areas C and D.

example, if the tristimulus values of two opaque surfaces A and B are given by the three-dimensional vectors $\mathbf{x}_A$ and $\mathbf{x}_B$, respectively, the adjusted colors of the filtered regions would lie on lines passing through $\mathbf{x}_A$ and $\mathbf{g}$ and through $\mathbf{x}_B$ and $\mathbf{g}$, respectively, where $\mathbf{g}$ defines the tristimulus values of the convergence point. The convergence model can be expressed by the following equations:

$$\mathbf{x}_C = (1 - \alpha)\mathbf{x}_A + \alpha\mathbf{g},$$
$$\mathbf{x}_D = (1 - \alpha)\mathbf{x}_B + \alpha\mathbf{g}, \tag{1}$$

where α defines the amount by which the surface colors $\mathbf{x}_A$ and $\mathbf{x}_B$ are shifted towards the convergence point $\mathbf{g}$. In terms of linear algebra, the term α is a 3×3 diagonal matrix (where only the three diagonal elements have nonzero values). The model represented by Eqs. (1) is known as the generalized convergence model. A special case of the generalized convergence model is known simply as the convergence model in which the three diagonal elements of α are identical. Note that in the generalized convergence model, the term for scaling the color of the opaque area may be different for each of the three dimensions of the space in which the color is defined, whereas for the convergence model, the scaling is the same for all three dimensions or color channels.

The performance of the convergence model has been compared with other models—including models based only on cone scaling, such as von Kries's—and has been quantitatively demonstrated to fit the color shifts that correspond to transparency better than the other models.[13] More recently it has also been shown that the convergence model can account for the color changes that take place when surfaces are viewed through a fog.[14]

The generalized convergence model and the convergence model define chromatic conditions under which transparency perception can occur. The approach that we adopt in this paper, which we call the invariant-ratios model, similarly defines chromatic conditions for transparency perception and, like the convergence model, can be shown to be a special case of the generalized convergence model. Our approach was inspired by a simple computational model of color constancy based on the invariance of cone-excitation ratios.[15]

Perception of transparency poses the general question of how the visual system can correctly recognize the color of a surface when its color has been altered in some way, for example, by covering the surface with a transparent filter. An analogous problem has been investigated in color constancy: When the color signal of a surface is altered by illuminating it with a different light source, its color appearance remains approximately constant (although the constraints on the color signals introduced are different from those in the case of transparency). In the case of a change in the illumination, it has been found that, within each cone class, cone-excitation ratios between surfaces seen under one illuminant and cone-excitation ratios for the same surfaces seen under a different illuminant are almost invariant, and this may be a cue for color constancy.[16] In the case of perception of transparency, we make the same predictions, since certain changes to the illuminant are approximately equivalent to passing the illuminant through a transparent filter.

The principle of invariance of cone-excitation ratios states that the ratio of the cone excitations between two opaque surfaces and the ratio between the same surfaces covered by a filter are almost statistically invariant within each cone class. This can be expressed by the equation

$$e_{i,1}/e_{i,2} = e'_{i,1}/e'_{i,2}, \tag{2}$$

where the cone excitation is given by $e_{i,j}$ for cone class i (where $i \in \{L, M, S\}$, denoting long-, medium-, and short-wavelength-sensitive cone classes) and a surface j seen directly, and the prime superscript denotes the excitations for the surface viewed through a filter. However, if we represent the constraint implied by Eq. (2) in the form of the generalized convergence model, we arrive at Eqs. (3):

$$\mathbf{e}_C = (1 - \alpha)\mathbf{e}_A,$$
$$\mathbf{e}_D = (1 - \alpha)\mathbf{e}_B, \tag{3}$$

where $\mathbf{e}_A$ and $\mathbf{e}_B$ are the cone excitations of the opaque surfaces and $\mathbf{e}_C$ and $\mathbf{e}_D$ are the cone excitations of the corresponding filtered areas. It is clear, therefore, that both the convergence model and the invariant-ratios model are special cases of D'Zmura's generalized convergence model. For the convergence model the special case is that the scaling is the same for all three color channels, whereas for the invariant-ratios model the special case is that there is no translation term, only a scaling.

For many conditions the predictions made by the invariant-ratios model and the convergence model are very similar.[16] This is not surprising, since the two models are mathematically identical if the special cases are relaxed.

In the present study we report the results of two psychophysical experiments to investigate whether the invariant-ratios model can predict the perception of transparency. In experiment 1 we tested whether the invariance of cone-excitation ratios can be used as a cue for the perception of transparency by presenting images for which the invariance of cone-excitation ratios was deliberately violated for all three cone classes. Presentations whose cone-excitation ratios were invariant were compared with presentations whose cone-excitation ratios were not invariant. In experiment 2, cone-excitation invariance was manipulated for individual cone classes, or for pairs, or for all cone classes to test the effect of ratio-invariance violations for individual cone classes.

2. EXPERIMENT 1

In the first experiment simulations of transparency yielding different ratios of cone excitations were compared. In particular, a simulation of a physically plausible filter was compared with the following cases: (a) a simulation in which cone-excitation ratios were manipulated to be perfectly invariant, (b) a control simulation (the two trials had the same cone-excitation ratios), (c) a simulation in which cone-excitation ratios were manipulated by adding 25% of noise, and (d) a simulation in which cone-excitation ratios were manipulated by adding 50% noise. Observers were presented with two successive Mondrian patterns. In the simulation, each Mondrian was partially covered by a transparent filter that could be either a physically plausible filter or one of the four comparison filters (a)–(d). Observers were required to indicate by a button press which presentation showed a Mondrian covered by a homogeneous transparent filter.

The invariant-ratios hypothesis predicts that in each trial, the presentation containing the most-invariant cone-excitation ratios would be preferred. We note that in most cases both stimuli appeared transparent and therefore observers might have faced the question of which presentation to choose. However, we later show that the variances within and between observers were very small. On the basis of this result, we argue that observers may have used the same strategy, which could be: "choose the stimulus condition that evokes the strongest impression of transparency." If observers were using different strategies, then significant differences might be expected between observers. Nevertheless, we cannot discard the possibility that observers might have adopted different strategies that led to similar results, although we consider it to be quite unlikely. Our experimental paradigm tested which stimulus condition evoked the strongest impression of transparency.

A. Observers

Four naïve observers participated in the experiment, all of whom had normal or corrected-to-normal visual acuity and had been assessed as color normal on the Farnsworth–Munsell 100-hue test. None of them was aware of the nature or purpose of the experiment.

B. Apparatus

A Sony Trinitron GMD500 color monitor driven by a VSG2/3 video card of a personal computer was used for presenting the stimulus patterns. The resolution was 1152×864 pixels and the frame rate was 120 Hz. The monitor had been characterized[17] by using a Minolta spectroradiometer and was gamma corrected.[18]

C. Stimuli

Stimuli consisted of a Mondrian pattern (5.14×5.33 deg of visual angle) partially overlaid by a simulated transparent filter (1.14×4.76 deg). Spectral reflectances of the opaque surfaces were selected from 1269 samples[19] of the Munsell Book of Color (1976). A subset of 100 Munsell samples was selected whose chromaticities all lay in the gamut of the monitor even when the reflectances were filtered and noise was applied to them. Effective spectral reflectances $R'(\lambda)$ for the filtered surfaces $R(\lambda)$ were computed according to Wysecki and Stiles.[20] The formula is illustrated by Eq. (4),

$$R'(\lambda) = R(\lambda)[T(\lambda)(1 - r)^2]^2, \tag{4}$$

where r is the internal reflectance of the filter and filter transmittance $T(\lambda)$ is defined by the modified Gaussian distribution shown in Eq. (5),

$$T(\lambda) = 0.4 + 0.6 \exp[-(\lambda - \lambda_m)^2/2\sigma^2], \tag{5}$$

where λ_m was randomly (with a flat probability distribution) selected in the range $400 \leqslant \lambda_m \leqslant 700\,\text{nm}$ and σ could be 5, 25, or 50 nm.

The assumption that the spectral transmittance of most natural filters is Gaussian (or inverted Gaussian) is based on the fact that most natural organic and inorganic molecules in solution form Gaussian distributions with respect to wavelength.[21] The offset of 0.4 in Eq. (5) was used to ensure that a strong percept of transparency would be generated by the stimuli.[22] Filter transmittance for all the filters was normalized so that the total transmittance was invariant to changes in σ. The internal reflectance r was also fixed at 0.1 throughout the experiment. Note that our choice of method for stimulus generation ensures that to a first approximation our stimuli are physically reasonable, although it is not important to our experiment that our model of physical transparency be accurate.

The first step in generating the four comparison simulation cases (a)–(d) was the simulation of a physically plausible filter. The second step was the computation of the cone-excitation ratios for CIE illuminant D65 for each opaque–transparent pair. The third step was the manipulation of those ratios to give rise to cone-excitation ratios that (a) are perfectly invariant, (b) result from the physically reasonable model, (c) are perturbed by adding 25% noise to the cone excitations, and (d) are perturbed by adding 50% noise to the cone excitations. Note that case (b) involved no extra manipulation of cone excitations and was included as a control.

Given the ratio between two opaque surfaces, case (a) was achieved by adjusting the cone excitation of one of the two filtered surfaces such that the resulting cone-excitation ratio was exactly the same as the ratio for the opaque surfaces. For example, consider two opaque surfaces with cone excitations $e_{i,1}$ and $e_{i,2}$ and their corresponding cone excitations $e'_{i,1}$ and $e'_{i,2}$ when filtered. To have $e_{i,1}/e_{i,2} = e'_{i,1}/e'_{i,2}$, the cone excitations of one of the filtered surfaces (for example, that corresponding to $e'_{i,2}$) was manipulated such that Eq. (2) was satisfied. Cases (c) and (d) were generated by adding noise to the cone-excitation ratios of the filtered surfaces by multiplying the cone excitations by a term $(1 + N)$, where N was randomly selected from a uniform variable in the range $[-0.25, +0.25]$ for (c) and $[-0.50, +0.50]$ for (d). It is important to note that noise was added on a patch-by-patch basis rather than on a pixel-by-pixel basis, and therefore the noisy areas remained spatially uniform.

Simulated filters in the two intervals were displayed with different orientations that could be vertical or horizontal. A schematic representation of the stimuli is illustrated in Fig. 2. Each pair of presentations consisted of a

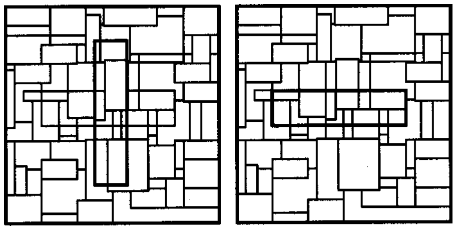

Fig. 2. Each pair of sequential presentations showed a simulation of a Mondrian partially covered by either a vertically or a horizontally oriented filter.

vertically and horizontally oriented filter, but the order of presentation was randomized. Vertical and horizontal filters lay over different opaque surfaces; thus, even for filters simulated in case (b), the two stimuli showed filters with different orientation. The underlying surfaces were randomly selected for each stimulus.

Each trial consisted of a pair of stimuli; in one stimulus the filter was physically reasonable [case (b)], and in the other the filter was a comparison filter that was selected from case (a)–case (d). The two stimuli will be referred to as "real" and "comparison" hereafter.

D. Procedure

In a two-alternative forced-choice paradigm, observers sequentially viewed simulations of two Mondrian patterns partially overlaid by transparent filters. Each presentation lasted 1 s on screen. The interstimulus interval also lasted 1 s. The next trial was presented 2 s after the observer indicated his or her response with a button press. Observers indicated which interval contained a homogeneous transparent filter that covered the Mondrian pattern by pressing one of two buttons on a pushbutton switch box. Each trial was repeated three times and the session of 72 trials was run three times. A training run of 20 trials was given before each session and the data subsequently discarded. No feedback was provided during the experiment.

E. Results

For each trial we calculated the degree of deviation from invariance in spatial cone-excitation ratios for all the possible pairs of surfaces seen directly and under the filter displayed in each image. The degree of deviation was equal to

$$\text{deviation} = 1 - r_i \quad \text{if } r_i \leqslant 1,$$
$$\text{deviation} = 1 - 1/r_i \quad \text{if } r_i > 1, \tag{6}$$

where r_i is the ratio of cone-excitation ratios, defined as

$$r_i = (e_{i,1}/e_{i,2})/(e'_{i,1}/e'_{i,2}). \tag{7}$$

For any pair of surfaces seen directly and under the filter, three ratios of cone-excitation ratios [as indicated in Eq. (7)] can be defined: the ratio between their S cones (for brevity r_s), the ratio between their M cones (r_m), and the ratio between their L cones (r_l). Mean deviations from an invariant ratio (i.e., r_i equal to 1) were then calculated for each single presentation separately for each of the three cone classes. Results are shown in Tables 1–3.

Table 1. Mean Deviations for S-Cone Class[a]

	S-Cone Deviations		
σ	Real	Comparison	
5 nm	0.0035	Perfect	0.0000
	0.0012	Real	0.0038
	0.0012	Noise 25%	0.1685
	0.0044	Noise 50%	0.3527
25 nm	0.0025	Perfect	0.0000
	0.0043	Real	0.0074
	0.0077	Noise 25%	0.2021
	0.0032	Noise 50%	0.3773
50 nm	0.0048	Perfect	0.0000
	0.0029	Real	0.0023
	0.0059	Noise 25%	0.1387
	0.0027	Noise 50%	0.3803

[a] Mean deviations have been computed for the two presentations (real versus comparison) in each trial.

Table 2. Mean Deviations for M-Cone Class[a]

	M-Cone Deviations		
σ	Real	Comparison	
5 nm	0.0787	Perfect	0.0000
	0.0195	Real	0.1540
	0.0077	Noise 25%	0.2317
	0.0041	Noise 50%	0.6222
25 nm	0.0140	Perfect	0.0000
	0.0155	Real	0.0222
	0.0132	Noise 25%	0.1299
	0.0167	Noise 50%	0.6078
50 nm	0.0053	Perfect	0.0000
	0.0075	Real	0.0041
	0.0035	Noise 25%	0.1783
	0.0025	Noise 50%	0.5235

[a] Mean deviations have been computed for the two presentations (real versus comparison) in each trial.

Table 3. Mean Deviations for L-Cone Class[a]

	L-Cone Deviations		
σ	Real	Comparison	
5 nm	0.0777	Perfect	0.0000
	0.0777	Real	0.0794
	0.0428	Noise 25%	0.2252
	0.0053	Noise 50%	0.2231
25 nm	0.0108	Perfect	0.0000
	0.0200	Real	0.0300
	0.0113	Noise 25%	0.1321
	0.0363	Noise 50%	0.7427
50 nm	0.004	Perfect	0.0000
	0.0126	Real	0.0036
	0.0041	Noise 25%	0.1016
	0.0034	Noise 50%	0.4864

[a] Mean deviations have been computed for the two presentations (real versus comparison) in each trial.

When the comparison filter is a perfect filter, r_i is always equal to 1, and therefore its deviations are equal to zero.

Both the filter with 25% noise and the filter with 50% noise have r_i values that are far from unity and, as illustrated in Tables 1–3, exhibit very high deviations.

The observers' ability to discriminate between a physically plausible filter (real) and any one of the four comparisons (a)–(d) was tested by measuring the discriminability index d prime (d') of signal detection theory. For the purposes of our analysis, a correct response was deemed to be the real filter and an incorrect response was deemed to be the comparison filter. Values of d' equal to zero indicate chance performance; d' greater than zero indicates preference for the real filter, d' less than zero indicates preference for the comparison filter.

In Fig. 3, mean values of d' for all the four conditions have been plotted against σ levels. For each condition, a one-way analysis of variance was performed to test for a significant effect of the variable σ. There was no significant effect ($p > 0.10$) of σ for any of the four conditions (a)–(d). Since there is no significant effect of σ on performance, d' values were pooled for all values of σ and the pooled values were plotted for each condition (Fig. 4). There is a significant effect ($F_{3,105} = 42.26$, $p < 0.001$), depending on the condition. There was no significant effect ($t_{35} = 0.71$, $p = 0.48$) in the control condition (real versus real), nor when the physically plausible filter is compared with a filter with approximately the same value of r_i (perfect versus real) ($t_{35} = 1.89$, $p = 0.07$). A simple analysis of the proportions correct gave qualitatively results very similar to those obtained by the d' analysis.

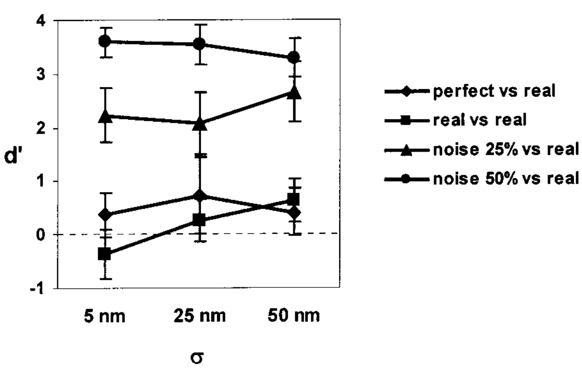

Fig. 3. Mean d' values for three different values of σ. Observers prefer the real filter when $d' > 0$. The control condition and the perfect-versus-real condition are close to chance performance ($d' = 0$).

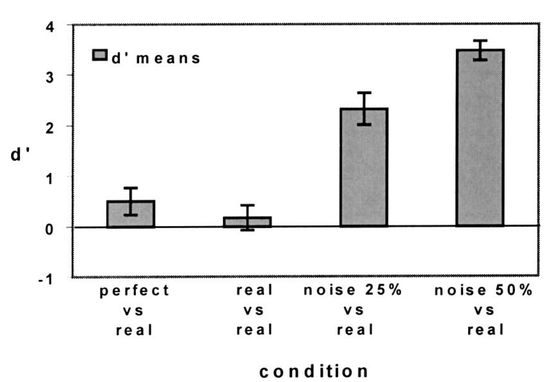

Fig. 4. Mean d' values for each of the four experimental conditions.

When the comparison filter was either case (c) or case (d), observers' performance was significantly greater than chance ($p < 0.01$). In each condition, observers reliably chose the real filter rather than the noisy filter where the invariance had been violated (recall that the noisy filters were no less spatially uniform than the real and perfect filters). We believe that this demonstrates that the ratios of cone-excitation ratios may provide a cue for perceptual transparency.

F. Summary

In summary, experiment 1 shows that when observers were asked which of two patterns shows a simulation of a transparent filter over a Mondrian and the two patterns differed markedly in their deviations from invariance, they reliably selected the pattern whose cone-excitations ratios were closer to invariance. We note, however, that observer performance was the same for ratios that were perfectly invariant and for ratios that were generated by a simple physical model of transparency. This is not surprising since it has previously been shown[22] that many physically transparent systems have ratios that are close to being invariant, and many models of physical transparency generate ratios that are close to being invariant. The changes that were made to convert the output of our model of physical transparency to a stimulus with invariant ratios were small and presumably below the threshold of perceptual discrimination.

3. EXPERIMENT 2

In experiment 1, presentations containing cone-excitation ratios close to invariance ($r_i \approx 1$) were compared with presentations in which the cone-excitation ratios were systematically violated. The dependence of observers' performance on the invariance of cone-excitation ratios did not reveal whether ratios of all cone classes must be invariant or whether the invariance of only one or two cone classes is sufficient to perceive transparency. This question is addressed by experiment 2.

In experiment 2, a simulation of a physically plausible transparent filter partially covering a Mondrian was compared with (a) a simulation in which cone-excitation ratios for all three cone classes were perturbed, (b) simulations in which cone-excitation ratios for single cone classes were perturbed, and (c) simulations in which cone-excitation ratios for pairs of cone classes were perturbed. In total, there were seven different combinations.

If transparency perception is mediated mainly by the L- and M-cone classes, then we expect a relatively small change in performance when noise is added to the S-cone class alone. We also expect the task to be easier when we add noise to both L- and M-cone classes than when we add noise to only one of these classes simply because the amount of noise added is greater. A similar experiment to measure the effect of cone-excitation ratios on relational color constancy found that sensitivity did depend on cone class.[23]

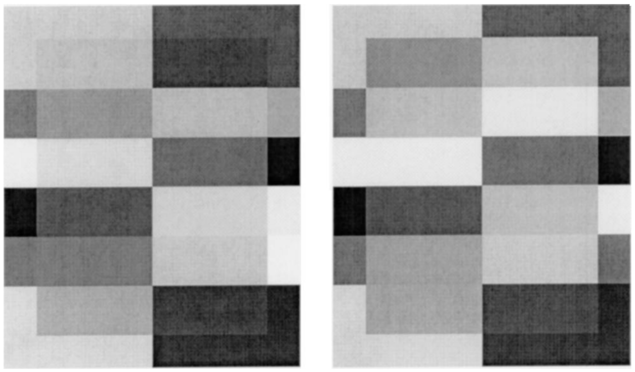

Fig. 5. Each trial in experiment 2 consisted of sequential displays of two Mondrians partially covered by a transparent filter. In one display (shown left), the spatial cone-excitation ratios were indirectly determined by the properties of the filter [Eqs. (4) and (5)]. In the other display (shown right), the cone excitations were systematically perturbed.

A. Observers

Three naïve observers participated in the experiment. All of them had normal or corrected-to-normal visual acuity and had been assessed as color normal on the Farnsworth–Munsell 100-hue test. None of them was aware of the nature and purpose of the experiment nor had participated in experiment 1.

B. Apparatus

The apparatus used was identical to that used in experiment 1.

C. Stimuli

Stimuli consisted of a Mondrian pattern (4.52×3.58 deg) partially overlaid by a simulated transparent filter (3.34×0.95 deg). The number of surfaces contained in the pattern was 12 and they were displayed in a 2×6 arrangement (Fig. 5). A transparent filter was simulated to be partially covering each of the 12 opaque surfaces. Spectral reflectances for the opaque surfaces were selected from the same set as in experiment 1. In order to use only those Munsell colors lying in the color gamut of the monitor, the preselection procedure applied in experiment 1 was also applied in experiment 2.

Spectral reflectances R' for the filtered surfaces were simulated according to Eq. (4). Filter transmittance T was defined by Eq. (5), and the value of λ_m was randomly selected. The standard deviation σ was fixed at 15 nm. As in experiment 1, filter transmittance was normalized, the total transmittance was independent of σ, and the internal reflectance r was set equal to 0.1 throughout the experiment.

We added noise to the cone-excitation ratios (derived from cone excitations computed for CIE illuminant D65) of the filtered surfaces by multiplying the cone excitations by a term $(1 + N)$, where N represented randomly selected, uniform variables in the ranges $[-0.04, +0.04]$, $[-0.08, +0.08]$, and $[-0.10, +0.10]$ corresponding to 4%, 8%, and 10% noise, respectively.

D. Procedure

In a two-alternative forced-choice paradigm, observers viewed two successive Mondrian patterns. Each pattern was partially covered by a transparent filter that could be either a physically plausible filter or one of the seven comparison filters. Each presentation lasted 1 s on screen. The interstimulus interval lasted 2 s, as did the intertrial interval. The next trial was not presented until the observer's answer was given.

Observers indicated whether a homogeneous transparent filter covering the Mondrian pattern was displayed in the first or in the second presentation by pressing one of two buttons on a pushbutton switch box. In half the trials the physically transparent filter was displayed in the first presentation; in the remaining trials the perturbed filter was displayed in the first presentation. The order of presentation was randomized. Each trial was repeated 4 times for a total of 168 trials [4 (repetitions) $\times$ 3 (amount of noise) $\times$ 7 (combinations of cone classes perturbed) $\times$ 2 (randomized presentation order)]. The complete session of 168 trials was repeated twice. A training run of 20 trials was given before each session and the data subsequently discarded. No feedback was provided during the experiment.

E. Results

Figure 6 shows d' values plotted against levels of noise and combinations of perturbed cone classes. Values of d' equal to zero indicate chance performance; d' greater than zero indicates preference for the physically plausible filter; and d' less than zero indicates preference for the noisy condition.

Excluding the case in which noise is added to the S-cone class alone, observers' performance is always above chance. Tables 4–6 show results of independent Student t-test analyses performed on means of d' tested against zero.

Figure 7 summarizes the observers' performance for different levels of noise. The independent variable noise does not significantly influence observers' performance

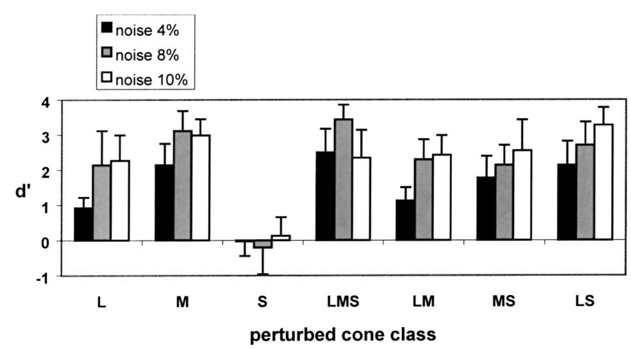

Fig. 6. Mean discrimination performance when noise is added to individual cone classes (L, M, and S), pairs of cone classes (LM, MS, and LS), and all three cone classes (LMS). Data for three levels of noise are shown.

Table 4. Means of d' Tested Against Zero[a]

	L	M	S	LMS	LM	MS	LS
$t_{1,5}$	3.098	3.523	−0.079	3.716	2.923	2.857	3.126
Signal	0.027	0.017	0.940	0.014	0.033	0.036	0.026

[a] In each trial the comparison filter was perturbed by 4% noise.

Table 5. Means of d' Tested Against Zero[a]

	L	M	S	LMS	LM	MS	LS
$t_{1,5}$	2.211	5.528	−0.269	8.114	4.016	3.727	4.093
Signal	0.078	0.003	0.799	0.000	0.010	0.014	0.009

[a] In each trial the comparison filter was perturbed by 8% noise.

Table 6. Means of d' Tested Against Zero[a]

	L	M	S	LMS	LM	MS	LS
$t_{1,5}$	3.098	6.523	0.239	2.915	4.221	2.907	6.505
Signal	0.27	0.001	0.821	0.033	0.008	0.034	0.001

[a] In each trial the comparison filter was perturbed by 10% noise.

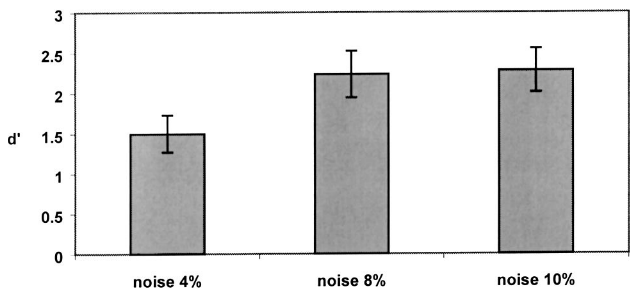

Fig. 7. Summary of observer performance at three different noise levels (averaged over all conditions).

($F_{2,4} = 1.173$, $p = 0.3973$). However, a comparison between 4% noise versus 8% and 10% noise is significant ($F = 11.585$, $p = 0.0023$).

F. Summary

In summary, experiment 2 supports the results found in experiment 1 in that observers prefer patterns for which the cone-excitation ratios are closer to invariance. More interesting, however, the data show that the invariance of the S-cone excitations is less useful for the prediction of perceptual transparency than that of the L- and M-cone classes.

4. DISCUSSION

We have demonstrated that the human visual system is able to discriminate simulations of filtered surfaces in which ratios of cone-excitations remain invariant from similar simulations in which the cone-excitation ratios are not invariant. In a previous study,[22] we used a simple model of physical transparency to show that the cone-excitation ratios for pairs of surfaces viewed directly and for the same surfaces viewed through a physically transparent filter were close to being equal in certain conditions. We also used direct measurements on physically transparent systems to show that the ratios of some physically transparent systems are close to being invariant. Although it is interesting that some physically transparent systems can be shown to have invariant cone-excitation ratios (at least to a first approximation), it is clear that such ratios will not be invariant for all physically transparent systems. It is important to note that the hypothesis that perceptual transparency can be predicted by the invariance of cone-excitation ratios does not rely on such ratios being invariant for all physically transparent systems. Indeed, it relies on such ratios not being invariant for all physically transparent systems, since not all physically transparent systems are perceptually transparent. In the previous study,[22] we did find that although the cone-excitation ratios were close to invariance for some physical conditions according to our model, the invariance was poor for filters with narrow-band spectral transmission properties.

This paper addresses the question of whether the invariance of cone-excitation ratios for surfaces covered by a transparent filter is a cue for perceptual transparency. Two experiments were carried out, and in each experiment, observers were asked to discriminate between two simulations of Mondrians partially covered by a transparent filter. Observers were asked to select which of the two presentations simulated a homogeneous transparent filter.

We found that observers were able to discriminate between simulations with ratios close to invariance and similar simulations in which noise had been added to violate the invariance (experiment 1). We also tested whether observers could discriminate between simulations of physically plausible filters and similar simulations in which the ratios had been manipulated to make them perfectly invariant. Observers were not able to discriminate between these two conditions, and we suggest that this is because the ratios of the physically transparent systems that we simulated were close to being invariant. Our model of physical transparency was only a first approximation of a real physical system. We suggest that if we had used a more complex and realistic model, we might have found greater deviations from invariance, and it is possible that under these conditions we would have found a significant difference between our physically plausible filters and our perfect filters.

We also found that noise added to the S-cone class alone did not affect discrimination performance (experiment 2). Simulations of transparent systems were compared with similar simulations in which noise had been added selectively to the individual cone classes or to specific combinations of cone classes. Observers' performance was always above chance except when noise was added to the S-cone class alone. There was some indication that observers' performance increased with the level of noise. Experiment 2 might lead to the suggestion that the S cones are less important to the task of perceptual transparency than M and L cones, since discrimination performance in a transparency-perception task is enhanced to a far greater extent when noise is added to the M- and L-cone classes than when noise is added to the S-cone class. Although it seems likely that the S-cone class is less important for transparency perception, we do not suggest that the luminance signal is the primary signal used for transparency perception. Indeed, we note that perceptual transparency has been measured for images that are isoluminant.[11,12,16]

In a related study we investigated discrimination performance in experiments similar to those in experiment 2, but in which the number of surfaces in the simulated Mondrian was varied.[24] We found that in discrimination experiments between simulations of filters giving small deviations from invariance and filters giving large deviations, performance improved with the number of surfaces in the display. Our hypothesis was that with an increased number of surfaces, there would be a corresponding increase in the number of pairs of surfaces from which invariant cone-excitation ratios could be recovered. For computational models of perceptual transparency that make use of X junctions and T junctions, the number of X junctions also increases with the number of surfaces in the image.[6,22] It is likely that the number of surfaces is one of several factors that could affect the strength of psychophysical cues for transparency; we might reasonably expect other factors to include the variance of the surfaces and their spatial relationships. For real scenes, we would expect many additional factors (such as surface specularity and the degree of spatial uniformity of surfaces) to be involved.

Given that both the convergence model and the invariant-ratios model are special cases of the generalized convergence model, it is reasonable to expect that there may be a wide range of conditions for which the two models perform similarly but certain conditions for which the predictions of the two models would be different. The performance of the convergence model has been quantitatively demonstrated to fit the color shifts that correspond to transparency better than a simple cone scaling model.[13] However, in recent studies we have demonstrated that the invariant-ratios model makes better predictions than the convergence model for certain stimuli.[16] Thus it seems that there are stimuli that are better predicted by the convergence model, while at the same time there are other stimuli that are better predicted by the invariant-ratios model. The convergence model contains an additive component and might be expected to make good predictions for stimuli with substantial specular reflectance. Stimuli with scattering components such as turbid media may also be predicted well by the convergence model (however, such systems would be more accurately described as translucent rather than transparent). Similarly, we might predict that the situations in which the invariance model would outperform the convergence model would be those in which the ratios of the cone excitations are close to invariance but are very different for each of the cone classes. At the very least, our recent data would suggest that the generalized convergence model (with six parameters) is necessary to predict transparency perception for all stimuli.[16]

In conclusion, our psychophysical data support the hypothesis that invariant cone-excitation ratios may provide a cue for transparency perception. However, such invariance seems to be a necessary but not sufficient condition for transparency perception. Furthermore, the invariance constraint may not be uniquely represented as invariant cone-excitation ratios; indeed, other models (such as the convergence model[11]) may provide alternative representations of approximately the same constraint.

ACKNOWLEDGMENTS

We thank D. H. Foster, M. G. A. Thomson, and S. M. C. Nascimento for a critical reading of this manuscript and for their helpful suggestions.

Corresponding author C. Ripamonti may be reached by e-mail to ripamont@psych.upenn.edu.

REFERENCES

1. K. Koffka, *Principles of Gestalt Psychology* (Harcourt, Brace, New York, 1935).
2. S. C. Masin, "The luminance conditions of fuchs's transparency in two-dimensional patterns," Perception **27**, 851–859 (1998).
3. S. Morinaga, K. Noguchi, and A. Ohishi, "Dominance of main direction in the apparent transparency," Jpn. Psychol. Res. **4**, 113–118 (1962).
4. O. Da Pos, *Trasparenze* (Icone, Padua, Italy, 1989).
5. B. Khang and Q. Zaidi, "Accuracy of color scission for spectral transparencies," J. Vision **2**, 451–466 (2000a).
6. G. Kanizsa, *Organization in Vision* (Praeger, New York, 1979).
7. P. Cavanagh, "Reconstructing the third dimension: Interactions between color, texture, motion, binocular disparity and shape," Comput. Vision Graph. Image Process. **37**, 171–195 (1987).
8. E. H. Adelson and P. Anandan, "Ordinal characteristics of transparency," in *Workshop on Qualitative Vision* (American Association for Artificial Intelligence Boston, Mass., 1990), pp. 77–81.
9. S. C. Masin, "An experimental comparison of three- versus four-surface phenomenal transparency," Percept. Psychophys. **35**, 325–332 (1984).
10. B. Khang and Q. Zaidi, "Cues and strategies for color constancy: perceptual scission, image junctions and transformational color matching," Vision Res. **42**, 221–226 (2002).
11. M. D'Zmura, P. Colantoni, K. Knoblauch, and B. Laget, "Color transparency," Perception **26**, 471–492 (1997).
12. J. V. Chen and M. D'Zmura, "Test of a convergence model for color transparency perception," Perception **27**, 595–608 (1998).
13. M. D'Zmura, O. Rinner, and K. R. Gegenfurtner, "The colors seen behind transparent filters," Perception **29**, 911–926 (2000).
14. J. Hagedorn and M. D'Zmura, "Color appearance of surfaces viewed through fog," Perception **29**, 1169–1184 (2000).
15. D. H. Foster and S. M. C. Nascimento, "Relational colour constancy from invariant cone-excitation ratios," Proc. R. Soc. London, Ser. B **257**, 115–121 (1994).
16. C. Ripamonti and S. Westland, "Perceptual transparency for chromatic and isoluminant stimuli," Invest. Ophthalmol. Visual Sci. **42**, S50 (2001).
17. R. S. Berns, "Methods for characterizing CRT displays," Displays **16**, 173–182 (1996).
18. D. L. Post and C. S. Calhoun, "An evaluation of methods for producing desired colors on CRT monitors," Color Res. Appl. **14**, 172–189 (1989).
19. J. P. S. Parkkinen, J. Hallikainen, and T. Jaaskelainen, "Characteristic spectra of Munsell colors," J. Opt. Soc. Am. A **6**, 318–322 (1989).
20. G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae* (Wiley, New York, 1982).
21. L. T. Maloney, "Evaluation of linear models of surface spectral reflectance with small numbers of parameters," J. Opt. Soc. Am. A **3**, 1673–1683 (1986).
22. S. Westland and C. Ripamonti, "Invariant cone-excitation ratios may predict transparency," J. Opt. Soc. Am. A **17**, 255–264 (2000).
23. S. M. C. Nascimento and D. H. Foster, "Detecting natural changes of cone-excitation ratios in simple and complex coloured images," Proc. R. Soc. London Ser. B **264**, 1395–1402 (1997).
24. C. Ripamonti and S. Westland, "Invariant cone-excitation ratios predict the perception of transparency," Invest. Ophthalmol. Visual Sci. **41**(4), S238 (2000).