Detection of chromatic and luminance distortions in natural scenes

Ben J. Jennings*†, Karen Wang*†, Samantha Menzies*†,

Frederick A.A. Kingdom*

* McGill Vision Research, Department of Ophthalmology, McGill University, Montreal, Canada.

† Optometry & Vision Science, University of Waterloo, Ontario, Canada.

Corresponding author: Ben Jennings (ben.jennings@mcgill.ca)

Abstract

A number of studies have measured visual thresholds for detecting spatial distortions applied to images of natural scenes. In one study, Bex (2010) measured sensitivity to sinusoidal spatial modulations of image scale. Here we measure sensitivity to sinusoidal scale distortions applied to the chromatic, luminance or both layers of natural-scene images. We first established that sensitivity does not depend on whether the undistorted comparison image was of the same or of a different scene. Next we found that when the luminance but not chromatic layer was distorted, performance was the same irrespective of whether the chromatic layer was present, absent or phase scrambled; in other words the chromatic layer, in whatever form, did not affect sensitivity to luminance-layer distortion. However when the chromatic layer was distorted, sensitivity was higher when the luminance layer was intact compared to when absent or phase-scrambled. These detection threshold results complement the appearance of periodic distortions of image scale: when the luminance layer is visibly distorted, the scene appears distorted, but when the chromatic layer is visibly distorted, there is little apparent scene distortion. We conclude that (a) observers have an in-built sense of how a normal image of a natural scene should appear, and (b) the detection of distortion in, as well as the apparent distortion of natural-scene images is mediated predominantly by the luminance not chromatic layer.

1 Introduction

When navigating the real world, be it a natural landscape or artificial environment, visual information is often spatially distorted. Distortions can arise in many ways (Fleming, Jäkel and Maloney, 2011; Schluter and Faul, 2014), for example they can be seen in the reflections from ripples on the surface of a lake (Fig 1a) or a curved metallic surface (Fig 1b). Distortions can also arise when light is diffracted during transmission through a geometrically curved transparent medium, e.g., glass (Fig. 1c). These spatial distortions are often modulated approximately sinusoidally; Fig. 1a illustrates a relatively high frequency example, Figs. 1b and c relatively low frequency examples.

Fig. 1a, b and c. Examples of spatially distorted images. Image (a) was produced when the surrounding natural environment was reflected from the surface of rippling water, (b) shows the distorted reflection of a cup produced by a curved metallic surface and (c) shows the distorted pattern of a tablecloth as viewed though a transparent wine glass.

A number of previous studies have measured observers' sensitivity to various types of natural-scene distortion, including uniform whole-image distortions (Kingdom, Field and Olmos, 2007), sinusoidally modulated distortions of scale (Bex, 2010) and distortion from blur (Wandell, 1995; Sharman, McGraw and Peirce, 2013). Kingdom et al. measured observers' sensitivities to a variety of transformations applied uniformly across images of natural scenes. The transformations included geometric (e.g., rotation, stretching), photometric (e.g., brightening, contrast-reducing) and noise-addition (e.g., Gaussian, fractal). Two of the geometric transformations they tested, shear and stretch, may be considered spatial distortions. Kingdom et al. found that for transformations equated in their Euclidean distance (the square root of the sum of squared differences between corresponding RGB pixels in the image pairs), observers were relatively insensitive to regularly experienced transformations, such as image translation, and relatively sensitive to uncommon transformations, such as the addition of noise. Kingdom et al. opined that the visual system tended to discard information about commonly experienced transformations as part of the process of achieving perceptual invariance during image transformation.

The transformations in Kingdom et al. were applied uniformly to the image. Bex (2010), on the other hand, applied sinusoidal modulations of image distortion, specifically of scale, resulting in images containing periodic expansions and contractions. He found that sensitivity to the distortions depended on a number of factors, including distortion spatial frequency, retinal eccentricity, the contrast of the scene and the particular scene structure - for example observers were most sensitive to distortions in regions containing edges. Bex demonstrated that the amplitude spectra of the distorted and undistorted stimuli were identical, thus ruling out this feature as a causal factor. He concluded that the detection of distortions is a high-level visual process that relies on observers' expectations of what is normal in natural scenes.

Fig. 2b shows the result of applying a spatial radial transformation to an undistorted image (Fig. 2a). In this case the distortion is a barrel distortion, as produced, for example by a fisheye lens (Hecht, 2001). Figs. 2c and 2d illustrate the same barrel distortion applied respectively to only the chromatic or luminance layer. When the distortion is applied to the whole image (Fig. 2b) or just the luminance layer (Fig. 2d), the distorted image appears like a reflection on a curved, shiny, possibly metallic, surface. On the other hand when only the chromatic layer is distorted (Fig. 2c) the resulting image is similar in appearance to the undistorted original, even though the disturbance to the colours is readily discernible. In this case a 'ghosting' can be observed whereby the colours appear to be smeared across the luminance boundaries, with the apparent structure of the scene being largely unaffected. This simple demonstration suggests that our sense of image scale distortion is largely mediated by luminance not chromatic information.

The two main aims of the study are as follows. First, we wanted to know whether sensitivity to image distortion reflects an in-built knowledge of what a "normal" scene looks like. Second we wanted know whether the appearance of the distortions in Fig. 2 is also reflected in measures of *sensitivity* to image distortion. The approach of separately manipulating the chromatic and luminance components of images of natural scenes in order to compare their relative contribution to a perceptual attribute has been applied in a number of previous studies (Wandell, 1995; Yoonessi & Kingdom, 2008; Kingdom, 2011; Sharman, McGraw and Peirce, 2013). As with Yoonessi & Kingdom's (2008) study of uniform color transformations and Bex's (2010) study of periodic scale distortions, we use phase-scrambled versions of the images in order to determine whether scene structure is a factor in transformation/distortion sensitivity. The potential importance of scene structure for detecting distortions applied to either or both of the luminance and chromatic layers becomes clear when one considers that in natural scenes most edges are both luminance- and color-defined (Fine, MacLeod & Boynton, 2003; Johnson, Kingdom & Baker, 2005; Hansen and Gegenfurtner, 2009).
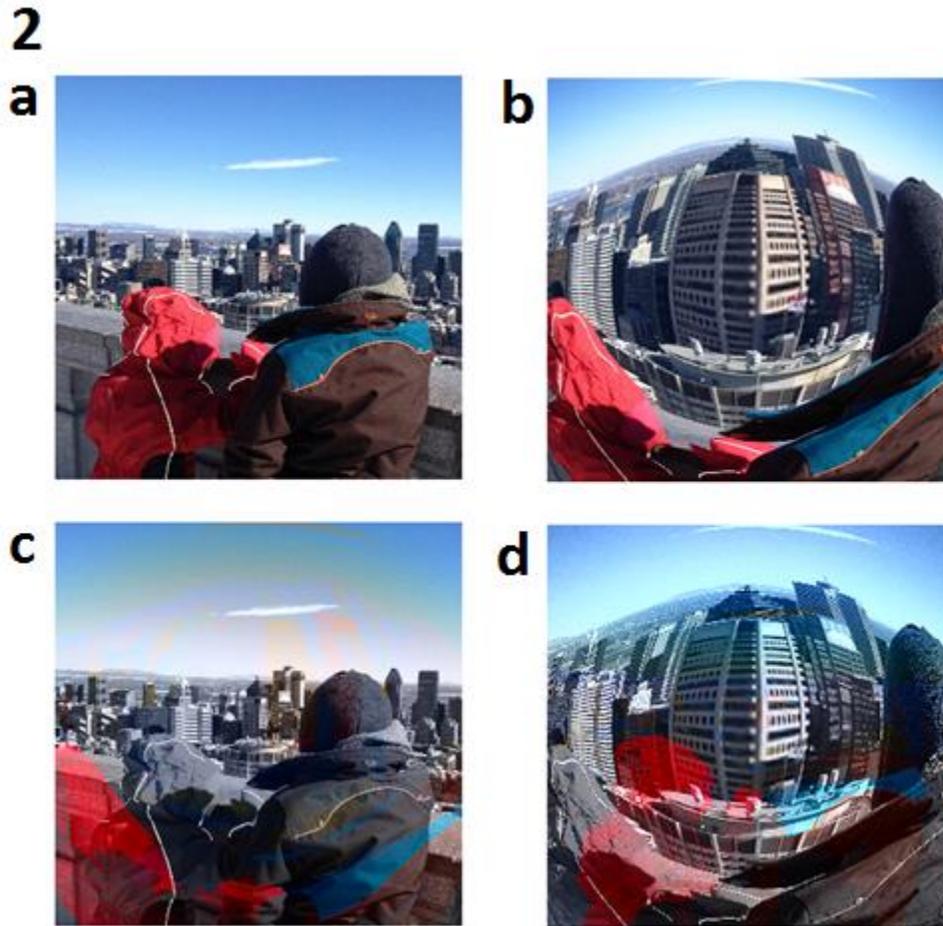
Fig. 2. **The same scene subject to chromatic/luminance distortion. Panel (a): original image. Panel (b) is the result of applying a barrel distortion to (a). Bottom panels are the result of applying the same barrel distortion to only the chromatic (c) or luminance layers (d) of the original.**

2. General methods

2.1 Observers

Seven observers participated in the experiments. Author KW participated in all experiments. Author SM participated in experiment 1 only. Author BJ participated in experiments 2, 3, and 4 only. The remaining four observers were naive to the purpose of the experiments. All observers had normal or corrected to normal (6/6) visual acuity. All observers additionally had normal colour vision, as tested by the Ishihara Colour Test (Isshinkai Foundation, published by Kanehara & Co., Ltd, 2001).

2.2 Equipment

The stimuli were presented on a CRT Sony Multiscan Trinitron G400 monitor, driven by a ViSaGe graphics display system (CRS ltd, UK) hosted by a DELL Precision T1650 computer. The display controlling software was programmed in C and utilised the ViSaGe Win32 application programming interface. The display was gamma corrected using a colorCAL (CRS ltd, UK) controlled via the vsgDesktop software. The spectral emission functions of the red (R), green (G) and blue (B) phosphors were measured using a SpectroCAL (CRS ltd, UK). The CIE xyY coordinates of the R, G and B phosphors at maximum luminance outputs were; red: xyY=(0.62, 0.34, 16.6 cd m^{-2}), green: xyY=(0.28, 0.61, 55.4 cd m^{-2}) and blue: xyY=(0.15, 0.07, 7.6 cd m^{-2}). The stimuli were presented on a mid-grey background located at xyY=(0.29, 0.31, 40.2 cd m^{-2}). The monitor was run with a refresh rate of 85 Hz and a resolution of 1280 x 1024 pixels, one pixel measured ~0.94 x 0.94 minute of arc at the viewing distance of 100 cm.

2.3 Stimuli

All stimuli were pre-generated using MatLab (Mathworks, Natick, MA) prior to the psychophysical testing. One hundred and fifty four digital photos were chosen from the McGill Calibrated Colour Image Database (Olmos and Kingdom, 2004). Fig. 3a provides examples of the images employed.

3a

3b

Distortion amplitude (vertical axis)

Distortion frequency (horizontal axis)

Fig. 3. (a) a pseudo-randomly selected subset of the 154 raw images employed in the study, illustrating the range of scene types, e.g. natural landscapes, foliage, flowers, urban scenes and human made objects is illustrated. (b) a series of distortion frequencies (increasing along the positive horizontal axis) and distortion amplitudes (increasing along the positive vertical axis) applied to a single image. These examples depict distortions well above threshold.

The images were distorted by applying a sine-wave distortion algorithm. The sinusoidal transformation was implemented in MatLab (Mathworks, Natick, MA) and applied to the image in both the horizontal and vertical directions simultaneously. For given distortion frequency and amplitude the horizontal and vertical distortion amplitudes were equal. The stimuli were pre-computed before each testing session as it took too much time to generate them on-line between trials. A series of six distortion spatial frequencies were selected, from ~0.08 to 2.7 cycles/°, and for each frequency a range of distortion amplitudes were employed. The distortion algorithm was applied to pseudo-randomly selected square subsections (512 x 512 pixels) cropped from the raw images. In experiment 1 the algorithm was applied to both the luminance and chromatic layers of the raw cropped image simultaneously. Examples of images distorted with various distortion amplitudes and frequency combinations are shown in Fig. 3b. Note the distortions illustrated in this figure are largely above threshold. In experiments 2, 3 and 4, the cropped images were first decomposed into their chromatic and luminance layers, then the distortion was applied just to one layer, chromatic or luminance, as required then the layers were recombined. Phase-scrambling, if required, was always applied prior to adding any distortion. Finally all stimuli had a circular Gaussian border applied to remove the sharp edge between itself and mid-grey background. Each image was displayed for 200 ms and ramped on and off according to a sine wave to avoid any artifacts being produced by a spontaneous stimulus onset.

2.4 Colour/luminance decomposition

The colour/luminance decomposition of each image was achieved by converting each gamma corrected pixel's RGB triplet into its corresponding YUV triplet. This was achieved by multiplying each RGB triplet by a 3x3 RGB to YUV transformation matrix. The Y layer of the YUV space contains the luminance information while the U and V layers contain the chromatic information. Hence a distortion and/or phase scrambling can be applied to the luminance layer by manipulating the Y layer, then recombining it back into RGB space with the unaltered U and V layers via the inverse RGB to YUV transformation matrix. On the other hand, a distortion and/or phase scrambling can be applied to both the chromatic layer (in RGB space) via manipulation of the U and V layers simultaneously before recombining them with the unaltered Y layer. To isolate the luminance information in an image the chromatic U and V layers pixels were set to zero. To isolate the chromatic information in the image the U and V layers were left unaltered, whist the Y layer was set to a constant value of 0.5, i.e., the background luminance, before the inverse transformation back to RGB space was applied. Note that the resulting isoluminant images are likely to have small luminance artifacts, as no attempt was made to make corrections based on individual variations in the luminosity efficiency function (Wyszecki and Stiles, 2000).

2.5 Phase scrambling

Phase scrambled versions of the images were generated according to the method outlined by Kingdom and Yoonsessi (2008). In this method, the absolute phases of the R, G and B layers were scrambled whilst preserving their relative phases. This ensured that as much colour information as possible was preserved in the scrambled images and that only the scene structure was destroyed. The algorithm employed a 2D fast Fourier transform to extract the amplitude A and phase P spectra, defined by Eqn. 1 and 2, respectively. The real and imaginary parts of each Fourier frequency component of the spectrum are F_r and F_i, respectively. The frequency variables are ω_x and ω_y.

$$A = \sqrt{F_r(\omega_x, \omega_y)^2 + F_i(\omega_x, \omega_y)^2} \qquad \text{Eqn. 1}$$

$$P = \arctan\left(\frac{F_i(\omega_x, \omega_y)}{F_r(\omega_x, \omega_y)}\right) \qquad \text{Eqn. 2}$$

The extracted phase spectrum subsequently had some randomly generated phase (in the range $\pm\pi$) added to it, before being recombined with the amplitude spectrum and the inverse Fourier transform performed. Any imaginary parts of the image produced as a side effect of rounding errors were discarded.

2.6 Experimental procedure

Distortion detection thresholds were obtained using a 2-IFC method and the method of constant stimuli. There were two main conditions, termed Same and Different (not to be confused with the "Same-Different" psychophysical task). In the Same condition the two images in each forced-choice pair were derived from the same image, i.e. the original plus a distorted version of the original. In the Different condition the two images images were of different scenes. Within each block the 6 distortion frequencies and 5 amplitudes were randomly interleaved. With 5 repeats of each combination of distortion frequency and amplitude this resulted in 150 trials per block (6 x 5 x 5 = 150 trials). There were four main conditions: Real-same, Real-different, Scrambled-same and Scrambled-different. Each of these four conditions was run 10 times in experiment 1, and 4 times in experiments 2, 3 and 4, resulting in 50 and 20 trials per frequency-amplitude combination, respectably. Psychometric functions of proportion correct versus distortion amplitude were fitted with a Weibull function using a maximum-likelihood criterion using the Palamedes toolbox (Prins and Kingdom, 2009). The Weibull estimates threshold at a proportion correct of ~0.82.

2.7 The different experiments

The following section outlines the various combinations of colour, luminance, distorted and phase scrambled conditions. Fig. 4 shows example stimuli.

Experiment 1: Chromatic and luminance layers were both distorted. Both real and phase scrambled images employed.

Experiment 2: One or other alone of the chromatic or luminance layer was distorted, with the other layer unaltered. Both real and phase scrambled images employed.

Experiment 3: One or other alone of the chromatic or luminance layer was distorted, while the other layer was phase scrambled.

Experiment 4: The chromatic or luminance layers were presented alone and distorted.

Raw image:

Experiment 1	Colour and luminance distorted		Colour and luminance phase scrambled and distorted	
Experiment 2	Luminance distorted		Colour and luminance phase scrambled, luminance distorted	
	Colour distorted		Colour and luminance phase scrambled, colour distorted	
Experiment 3	Colour phase scrambled, luminance distorted		Luminance phase scrambled, colour distorted	
Experiment 4	Luminance removed, colour distorted		Colour removed, luminance distorted	

Fig. 4. Top: unprocessed raw image. Below: examples of the different distortion conditions used in the experiments.

3. Results

The following sections present for each experiment distortion amplitude thresholds as a function of the spatial frequency of the distortion in log-log space, with all plots spanning equal ranges on the mantissas and abscissas, with error bars ±2SEM (standard error of the mean). All statistics are based on 2-tailed t-tests, with p-values Bonferroni corrected where appropriate.

3.1 Experiment 1 (*n* = 5)

In this experiment the distortion was applied equally to both chromatic and luminance layers. The four conditions were: Real-Same, Real-Different, Scrambled-Same, and Scrambled-Different (see experimental procedure). No significant differences were found between the Real-Same and Real-Different or between the Scrambled-Same and Scrambled-Different conditions (Fig. 5a) (*p*-values ranged $0.21 \leq p \leq 0.45$, i.e. all greater than 0.05). The Same and Different data were therefore combined and the resulting Real and Scrambled data is shown in Fig. 5b. Over the four highest spatial frequencies tested ('*' in Fig. 5b), all corrected *p*-values were in the range $0.01 \leq p \leq .028$, i.e. all less than 0.05. At the lowest two spatial frequencies tested ('†' in Fig. 5b) no differences exists (both *ps* ≥ .13). The slopes of the curves plotted in Fig. 5b are -0.97 and -1.1 for the real and scrambled stimuli, respectively, i.e. approximately -1. This implies that the product of amplitude threshold and distortion frequency should produce flat functions when plotted as a function of distortion frequency. Fig. 5c illustrates that this is the case, especially for the real image data. This shows that the rate of change of distortion across space is the main factor limiting performance in this experiment.
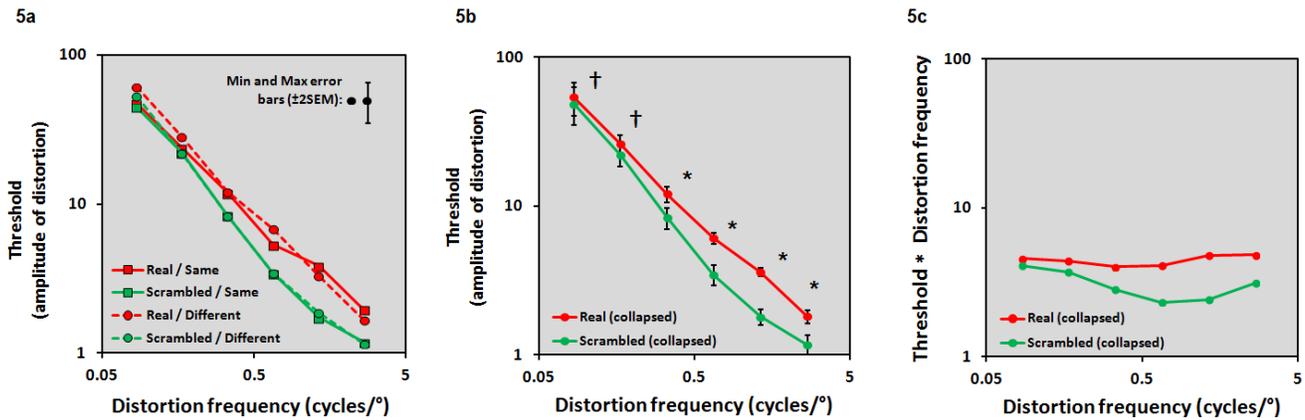
Fig. 5. Panels (a) and (b) show distortion amplitude thresholds as function of distortion spatial-frequency, when both chromatic and luminance layers are distorted. Panel (a) shows separately the Real-Same, Scrambled-Same, Real-Different and Scrambled-Different data. For clarity error bars are not displayed, however the top right corner of (a) shows minimum and maximum errors (±2SEM). Panel (b) shows data collapsed for the Real-Same and Real-Different, and for the Scrambled-Same and Scrambled-Different conditions. The symbol * indicates a significant difference between the Real and Scrambled conditions. Data is the average across 5 subjects and error bars are ±2SEM calculated across subjects' thresholds, i.e. are not errors on the fitted psychometric function thresholds. Panel (c) shows the same data as in panel (b) but plotted as the product of threshold amplitude and distortion frequency versus distortion frequency.

2.2 Experiment 2 (*n* = 3)

In this experiment only one layer, chromatic or luminance, was distorted. No significant differences were found between the Real-Same and Real-Different conditions, nor between the Scrambled-Same and Scrambled-Different conditions (corrected *p*-values ranged $0.16 \leq p \leq 0.56$, i.e. all greater than 0.05). Data was subsequently collapsed for both the chromatic (Fig. 6a) and luminance (Fig. 6b) conditions. A significant difference was found between the Real and Scrambled images when the chromatic layer only was distorted (*p*=.008); the detection of chromatic distortions in phase scrambled images being more difficult to detect. But there was no difference between the Real and Scrambled thresholds for the luminance defined distortions (*p*=.31). Almost significant (*p*=.06) were the lower thresholds for the luminance compared to chromatic distortions at the three highest frequencies. Finally, thresholds for phase-scrambled luminance distortions were significantly lower than for phase-scrambled chromatic distortions (*p*<.001).

6a

Chromatic distortions

6b
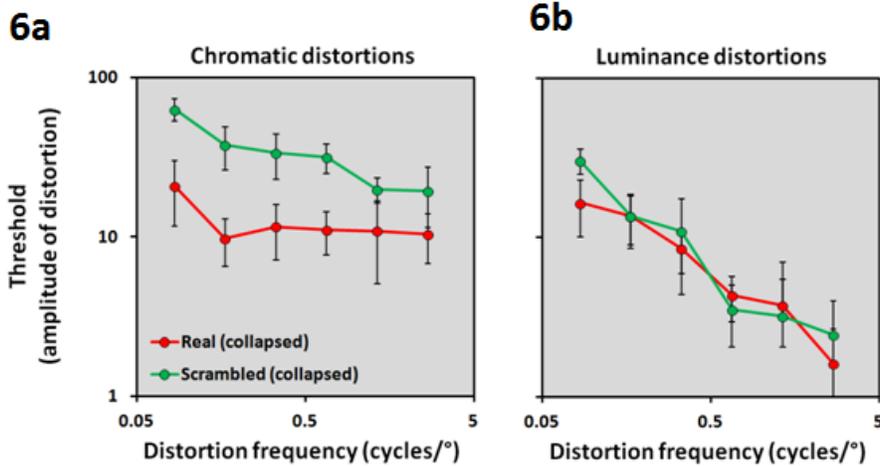
Luminance distortions

Fig. 6. Thresholds for detecting distortions when applied only to (a) the chromatic and (b) the luminance layer, as a function of distortion frequency. Thresholds for Real stimuli are plotted in red, Scrambled stimuli in green.

3.3 Experiment 3 (*n* = 2)

In this experiment one or other of the chromatic and luminance layers was distorted, while the other layer was phase-scrambled. Fig. 7a shows plots for the two subjects, BJ and KW tested, while Fig. 7b plots the mean values across the two subjects. Thresholds for detecting chromatic distortions with the luminance layer undistorted but phase-scrambled (purple curves) were significantly higher (p<.001) than thresholds for detecting luminance distortions with the chromatic layer undistorted but phase scrambled (dark grey curves). Moreover, the decline in thresholds with distortion spatial frequency on the log-log plots was steeper for the luminance compared to chromatic distortions.

7a

7b
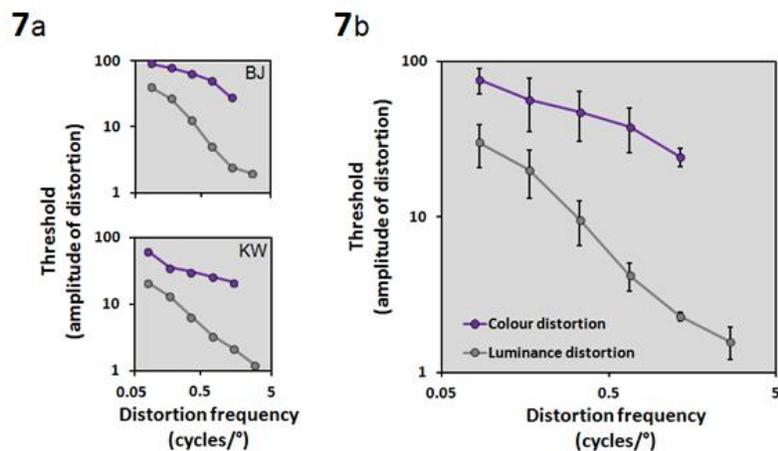
Fig. 7. Distortion thresholds for real images with either the chromatic or luminance layer distorted, with the other layer phase-scrambled. Panel (a) shows data for two subjects BJ and KW, while panel (b) shows mean thresholds across subjects. Chromatic layer thresholds are plotted in purple, luminance layer thresholds in dark grey.

3.4 Experiment 4 (*n* = 2)

Here, the chromatic and luminance layers were presented on their own, resulting in respectively isoluminant and isochromatic, in the latter case specifically achromatic stimuli. Fig. 8a shows individual data for BJ and KW, while Fig. 8b plots the mean values. Thresholds for detecting

chromatic layer distortions were significantly higher (p=.007) than thresholds for detecting luminance distortions. As with experiment 3 there is a steeper decline in thresholds with distortion spatial frequency for the luminance compared to chromatic distortions.
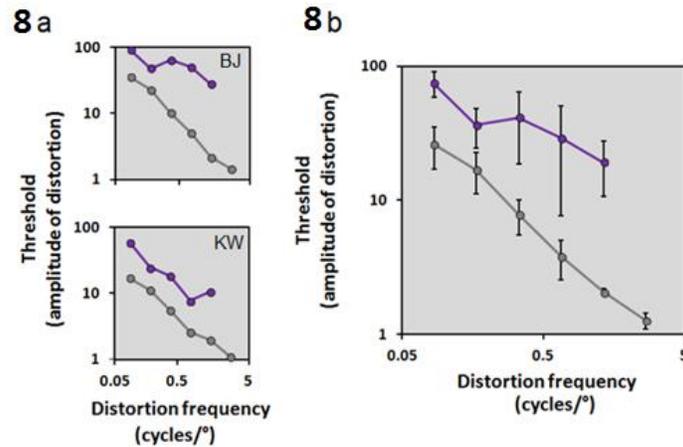
Fig. 8. Distortion detection thresholds for chromatic and luminance layers in isolation. Panel (a) shows plots of thresholds for BJ and KW, panel (b) shows mean thresholds. Chromatic thresholds are plotted in purple, achromatic thresholds in dark grey.

3.5 Comparison of results

To better grasp the significance of the results from experiments 2, 3 and 4, Fig. 9a and b compares their data. Consider first the data from experiments 3 and 4. No significant differences were found (p = .08) between chromatic distortion detection thresholds measured in the presence of a phase scrambled luminance layer and distortion detection thresholds for isoluminant stimuli (red and green curves in Fig. 9a, respectively). Furthermore, no significant differences were found between luminance distortion detection thresholds measured in the presence of a phase-scrambled chromatic layer and distortion detections for the pure luminance stimuli (red and green curves in Fig. 9b, respectively). This implies that observers were able to disregard the phase-scrambled "noise" in the irrelevant layer, irrespective of whether observers were detecting chromatic or luminance distortion.

On the other hand, if we compare experiments 2 with 3 and 2 with 4 one can see that observers were significantly more sensitive to chromatic distortions if the undistorted luminance structure was present (see blue curve in Fig. 9a; experiment 2 vs. 3: p = 0.01, experiment 2 vs. 4: p = .03). Whereas, detection thresholds for luminance distortions did not improve with the presence of undistorted chromatic structure (see blue curve in Fig. 9b). Surprisingly, as Fig. 9b illustrates, luminance thresholds are equal for a given distortion frequency, over the whole range (all p-values in range: $0.17 \leq p \leq 1$, (corrected upper p-value: 1.21)). Consequently, luminance distortion detection thresholds are independent of whether they are presented with an unaltered chromatic layer (experiment 2), a phase-scrambled chromatic layer (experiment 3), or in isolation (experiment 4).
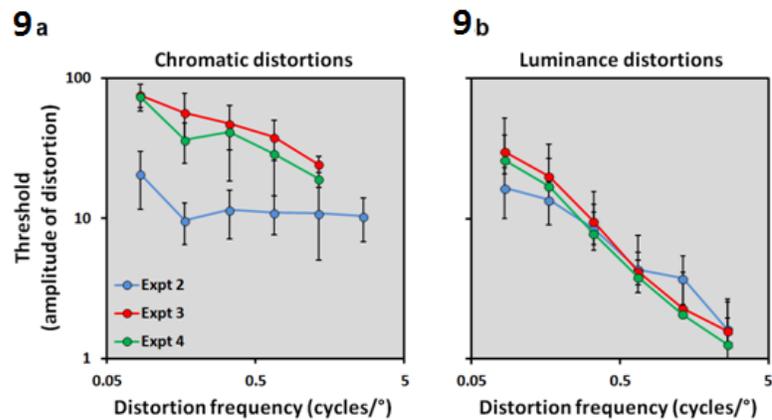
Fig. 9. Comparison of thresholds for experiments 2, 3 and 4. Panel (a) plots the chromatic distortion data while panel (b) plots the luminance distortion data. The blue curve in (a) shows chromatic distortion thresholds in the presence of an undistorted luminance layer and in (b) vice versa. The red curve 9n (a) shows chromatic distortion detection in the presence of a phase scrambled luminance layer and in (b) vice versa. The green curved in (a) shows isoluminant and in (b) isochromatic distortion detection thresholds.

4 Discussion

The following summarizes the main findings of the experimental part of the study.

1. Sensitivity for detecting sinusoidal distortions applied to the whole image (i.e. to both luminance and chromatic layers) is the same regardless of whether the undistorted comparison image is of the same or of a different scene.

2. Sensitivity for detecting whole-image distortions is higher for phase-scrambled compared to unscrambled scenes.

3. Distortion detection sensitivity increases (thresholds decline) as the distortion frequency increases, for both real and phase-scrambled scenes.

4. Sensitivity for detecting distortions is higher when the luminance layer is distorted compared to when the chromatic layer is distorted, especially at high distortion frequencies.

5. Sensitivity for detecting luminance layer distortions is independent of whether the chromatic layer is undistorted, phase scrambled or absent.

6. Sensitivity for detecting chromatic distortions is highest when the undistorted luminance structure is present.

Experiment 1 showed that for both real and phase-scrambled images, sensitivity for detecting sinusoidal distortions was independent of whether the comparison image was of the same or of a different scene. This implies that observers posses an internal representation of what is 'normal' in a natural scene, a conclusion also reached by Bex (2010). Although no difference was found between real and phase-scrambled thresholds for the two lowest distortions frequencies tested, at higher distortion frequencies observers were more sensitive to the distortions in the phase-scrambled images, also consistent with Bex's findings. The result is ostensibly consistent with Kingdom et al.'s (2007) conclusion that in general we are least sensitive to transformations that are normally experienced, i.e. the ones in real not phase-scrambled scenes. Unfortunately for the generality of this conclusion however, at least some types of image transformation show worse performance with phase-scrambled images, for example images subject to uniform colour transformations, such as rotations in colour space (Yoonessi and Kingdom, 2008).

An alternative explanation for the superiority in performance with the phase-scrambled images to that following Kingdom et al. (2007) is as follows. Applying distortion to a phase-scrambled image imposes structure on a stimulus that is relatively unstructured, whereas with the real image the structure is imposed on an already structured stimulus. Thus the observer's task can be loosely described as structure *detection* with the phase-

scrambled images as opposed to structure *discrimination* with the real-scene images. The analogy here is the difference between contrast detection and contrast discrimination: detecting a stimulus versus a blank requires less contrast than discriminating two high-contrast stimuli differing in contrast (e.g. Bradley and Ohzawa, 1986).

Experiment 2 on the other hand revealed a very different pattern of results. When only the chromatic layer was distorted (Fig. 6a), performance was *worse* with the phase-scrambled compared to the real scenes, and when only the luminance layer was distorted, performance was no better with the phase-scrambled compared to real scenes (Fig. 6b). How can we explain the discrepancy in results between the two experiments? The likely reason is that with the real but not phase-scrambled scenes, when the distortion is applied to just one layer, the undistorted layer provides a visible structure against which the presence or absence of distortions in the other layer can be compared. The real beneficiary of this process is colour: in the presence of in-tact luminance structure, distortions to the chromatic layer are easily detected, and at low frequencies as easily detected as luminance layer distortions (Fig. 6). With the structure of the image removed with phase-scrambling, distortion to the chromatic layer becomes much more difficult to detect.

Threshold results

A main experimental finding is that we are relatively insensitive to distortions to the chromatic layer. There are a number of possible reasons. First, if the amount of chromatic contrast relative to luminance contrast in natural scenes is low, we would expect sensitivity to chromatic distortions to also be relatively low. While doubtless a contributing factor, relatively low chromatic contrast is unlikely to be the sole cause. It does not, for example, explain why we are relatively good at detecting distortions to the chromatic layer when luminance structure is present, particularly at low distortion frequencies. A second possible contributory reason for the relatively poor chromatic distortion sensitivity is that in natural scenes colour is more sparse than luminance, meaning that chromatic information tends to be formed into relatively larger patches than luminance information (Yoonessi, Kingdom & Alqawlaq, 2008). This means that more of the image regions in the chromatic compared to luminance layer are devoid of edges and thence not subject to the effects of distortions. A third possible reason is the relative insensitivity of the chromatic system to high spatial frequencies (Mullen , 1985).

The appearance of image distortion

The experiments in this study measured thresholds, i.e. performance measures, using a conventional forced-choice task. How do the experimental results square with the appearance of image distortion as discussed earlier in relation to Fig. 2? We argue that there may be more to the appearance of Fig. 2 than mere lack of sensitivity to chromatic image distortions. In the natural world a large proportion of images reaching the eye are formed by light that has previously been reflected from the surfaces of other objects (Mandelstam, 1926; Kerker, 1969). Diffuse reflections occur when light is reflected equally in all directions, as for a Lambertian surface, with the result that the surface appears matte. Specular reflections on the other hand differ from diffuse reflections in that they are highly directional, with light reflected from the surface at an angle opposite to that of the incident light. An example of an ideal specular surface is a mirror, while a shiny metallic surface or glossy paint is a close approximation.

Do some surfaces predominately produce reflections that are composed largely of achromatic information, i.e., can colour information be lost during the process of reflecting an image? Different surfaces possess difference spectral reflectance curves, that is differences in the proportion of light reflected from a surface as a function of the light's wavelength. Spectral reflectance curves can reveal the selective absorption of a material; for example, copper and gold reflect almost 100% of long wavelength light, while only around 40% of short wavelength light (from violet to cyan). On the other hand, water reflects mostly high energy light from the bluish region of the spectrum. Hence, reflections originating from a number of materials would appear to have reduced colour spectra. Therefore the visual system might rely more on luminance than colour when dealing with patterns of specular reflection in order to encode the surface shape of curved, metallic objects, and this might in part explain the appearance of Fig. 2.

Conclusion

Our results suggest that human observers have an internal undistorted representation of the world, probably acquired via experience that can be relied upon when making judgments as to whether a scene, natural or artificial, has been spatially distorted. Luminance defined distortions are equally as salient regardless of whether the accompanying chromatic information is distorted, undistorted or phase scrambled. Chromatic distortions on the other hand are in general harder to detect, except in the presence of undistorted luminance structure. Our relative inability to use chromatic information to detect natural-scene spatial distortions is commensurate with the observation that the visual system appears to rely on luminance cues to encode the shapes of curved, metallic surface from the pattern of their specular reflections.

References

1. Bex, P. J. (2010). (In) sensitivity to spatial distortion in natural scenes. Journal of Vision, 10(2):23.1-15.
2. Bradley, A. Ohzawa,I. (1986). A comparison of contrast detection and discrimination. Vision Research, 26, 991-997.
3. Fine, I, MacLeod, D. A. & Boynton, G. M. (2003). Surface segmentation based on the luminance and color statistics of natural scenes, Journal of the Optical Society of America A, 20, 1283–1291.
4. Fleming, R. W., Jäkel, F., & Maloney, L. T. (2011). Visual perception of thick transparent materials. Psychological Science, 22(6), 812–820.
5. Hansen, T. & Gegenfurtner, K. R. (2009). Independence of color and luminance edges in natural scenes. Visual Neuroscience, 26, 35–49.
6. Hecht, E. (2001). Optics (4th Edition). Addison-Wesley. London, UK.
7. Johnson, A. P., Kingdom, F. A. A. & Baker, C. L. Jr. (2005) Spatiochromatic statistics of natural scenes: First- and second-order information and their correlational structure. Journal of the Optical Society of America A, 22, 2050-2059.
8. Kerker, K. (1969). The scattering of light and other electromagnetic radiation. Academic Press, New York: Academic.
9. Kingdom, F. A. A. (2011). Illusions of colour and shadow. Ch. in New Directions in Colour Studies, by C. P. Biggam, C. Hough, C. J. Kay & D. R. Simmons (eds.). John Benjamins Publishing Company.
10. Kingdom, F. A. A., Field, D. J., & Olmos, A. (2007). Does spatial invariance result from insensitivity to change? Journal of Vision, 7(14):11, 1-13.
11. Mandelstam, L.I. (1926). Light Scattering by Inhomogeneous Media. Zh. Russ. Fiz-Khim. Ova. 58: 381.
12. Mullen, K.T. (1985). The contrast sensitivity of human colour vision to red/green and blue/yellow chromatic gratings. J. Physiology, 359, 381-400.
13. Olmos, A., & Kingdom, F. A. A. (2004). A biologically inspired algorithm for the recovery of shading and reflectance images. Perception, 33, 1463-1473.
14. Prins, N., & Kingdom, F. A. A. (2009). Palamedes: Matlab routines for analyzing psychophysical data. http://www.palamedestoolbox.org.
15. Schluter, N., & Faul, F. (2014). Are optical distortions used as a cue for material properties of thick transparent objects? Journal of Vision, 14(14):2, 1–14.
16. RJ Sharman, R. J., McGraw, P. V., & Peirce, J. W. (2014). Luminance cues constrain chromatic blur discrimination in natural scene stimuli. Journal of vision, 13(4):14.
17. Yoonessi, A. & Kingdom, F. A. A. (2008) Comparison of sensitivity to color changes in natural and phase-scrambled scenes. Journal of the Optical Society of America A, 25, 676-684.
18. Yoonessi, A. & Kingdom, F. A. A. (2009). Dichoptic difference thresholds for uniform color changes applied to natural scenes. Journal of Vision, 9(2):3, 1-12.
19. Yoonessi, A., Kingdom, F. A. A. & Alqawlaq, S. (2008). Is color patchy? Journal of the Optical Society of America A, 25, 1330-1338.
20. Wandell, B. (1995) Foundations of Vision. Sunderland: Sinauer Associates, Inc.
21. Wuerger, S. M., Owens. H., & Westland,. S. (2001). Blur tolerance for luminance and chromatic stimuli. JOSA A, 18(6):1231-1239.
22. Wyszecki, G., & Stiles, W. S. (2000). Color science: Concepts and methods, quantitative data and formulae (2nd edition). New York: John Wiley & Sons.