Visual perception of thick transparent materials

Roland W. Fleming[1, 2] Frank Jäkel[1,3] Laurence T. Maloney[4]

1 Max Planck Institute for Biological Cybernetics, Tübingen, Germany

2 Department of Psychology, University of Gießen, Germany

3 Institute of Cognitive Science, University of Osnabrück, Germany

4 New York University, New York, USA

In press, *Psychological Science*, January 27, 2011

Corresponding author:

Roland W. Fleming
Justus-Liebig-Universität Gießen,
FB06 – Psychologie,
Otto-Behaghel-Str. 10/F2 -338
53594 Gießen
Germany

Running head: Perception of transparent materials

Word counts:

Text:	3855 (excluding references and figure captions)
Abstract:	149
Figures:	3

ABSTRACT

Under typical viewing conditions, human observers readily distinguish between materials such as silk, marmalade or granite, an achievement of the visual system that is poorly understood [Adelson, 2001]. Previous work on the perception of transparency [Koffka, 1935; Metelli, 1974a,b; Singh & Anderson, 2002] has focused on the perception of transparent objects that are flat, infinitely thin filters. Here, by contrast, we consider thick transparent objects, such as ice cubes, which are irregular in shape and can vary in refractive index. An important part of the visual evidence signalling the presence of such objects is distortions in the perceived shape of other objects in the scene. We propose a new class of visual cues based on the *distortion field* induced by thick transparent objects and provide experimental evidence that cues based on the distortion field predict both the successes and failures of human perception in judging refractive index.

[147 words]

Keywords: surface perception, material perception, lightness, grouping, material properties, maximum likelihood difference scaling (MLDS)

INTRODUCTION

Transparent objects, such as gemstones and glass can be fascinating and beautiful to look at. However, they pose the visual system with a formidable problem. Under typical viewing conditions, light passing through such objects refracts and reflects several times before reaching the eye, yielding image intensities that are a complex mixture of background and foreground surfaces. In order to recover the intrinsic properties of a transparent object—and any surfaces visible through it—the brain must somehow separate the various contributions to the observed image intensities.

For over a hundred years, perceptual psychology has tried to understand how the visual system identifies transparent surfaces and estimates their intrinsic properties [von Helmholtz, 1867; Metelli, 1970, 1974a,b; Gerbino, 1994; Singh and Anderson, 2002; Robilotto et al. 2002, 2004]. Almost all research to date models transparent objects as ideal thin, neutral density filters perpendicular to the line of sight (Figure 1A). The primary effect of such a transparent surface on the retinal image is to modify the intensities or chromaticities of patterns visible through it. Specific photometric and geometric pre-conditions for the perception of transparency have been derived from this assumption. For example, it is commonly argued that transparency leads to *contour junctions* in the image (contour junctions are locations at which a number of contours meet in the image), as highlighted in Figure 1a [Beck et al. 1984, 1988; Adelson and Anandan, 1990; Anderson, 1997]. These “X-Junctions” occur when an edge feature on the background layer passes behind a transparent surface, undergoing a reduction in contrast.

However, many transparent objects that we encounter on a daily basis, such as ice cubes or chunks of glass, are thick, refractive bodies (Figure 1CD). The physics of light transport through such real-world objects is markedly more complex than for neutral density filters, and this has profound consequences for the image cues that the visual system can exploit to infer the physical properties of a transparent object. In particular, thick transparent materials not only transform the photometric properties of patterns visible through them, but also their spatial properties.

When a light ray strikes the surface of a transparent object, some proportion of the light is reflected, and the remaining light refracts as it enters the body of the object (Figure 1B). The angle of refraction depends on (i) the local geometry of the surface and (ii) an intrinsic property of the material: its *refractive index* (RI). Light exiting the objects is similarly transformed.

Insert Figure 1 Here

Using a computer graphics simulation of refraction, it is possible to systematically modify RI, while holding constant other scene variables, such as lighting and object shape. In Figure 2, RI varies from 1.1 to 2.3, and we experience a concomitant change in the apparent material properties of the object. With increasing refractive index, the object appears more lustrous and 'substantial'. Refraction can evidently play a key role in the appearance of many transparent objects.

As evident in Figure 1C, refraction can substantially distort the patterns visible through the object and can prevent X-junctions from occurring at the boundaries of the transparent object. Despite this, we readily experience a vivid impression of a transparent object in the scene. This discrepancy demonstrates that there is a large class of real transparent objects that current theories of perceptual transparency do not consider. How does the visual system estimate the properties of such objects?

The fact that transmitted patterns are systematically distorted by refraction suggests that the visual system could use patterns of distortions to estimate RI. We can quantify the distortion effect by measuring how the positions of features in the background are displaced when viewed through the transparent object. Specifically, if the image coordinate of a feature in plain view is $\mathbf{p}_i = (x, y)$ and its position when refracted through the transparent object is $\mathbf{p}_r = (x', y')$, then we can define a vector field in image space $\mathbf{D}(x, y) = \mathbf{p}_r - \mathbf{p}_i$ which measures the displacement of all features in the background when seen through the transparent object. We call this the *displacement field*.

Given that the visual system does not have access to the initial, non-refracted positions of features in the background, $\mathbf{p}_i$ there is no way for it to measure the displacement field directly. However, it is plausible that the visual system may be able to estimate the *relative* compression and expansion of a texture pattern compared to its local context. The relative magnitude of compression can be captured mathematically as the divergence of the displacement field, $\mathbf{d}(x, y) = \nabla \circ \mathbf{D}(x, y)$, which we call the *distortion field*. Where d is positive, the displacement field is diverging, leading to a local magnification of the refracted pattern; where d is negative, the pattern is compressed. For an object of arbitrary shape, the distortions vary continuously across the image of the object. Previous work has suggested that local patterns of compression and magnification can be used to infer three-dimensional shape from textured and specular surfaces [Fleming et al., 2004]. Here, we suggest that a related strategy could provide the brain with information about refractive index.

Insert Figure 2 Here

Central to this line of reasoning is the observation that distortion fields tend to vary systematically with refractive index. As refractive index increases, the pattern of compressions and rarefactions remains similar, but the magnitude of the distortions increases. This observation suggests that the visual system could use some summary statistic of the magnitude of distortions—pooled across the image of an object—to estimate its refractive index.

If the visual system makes use of cues based on the distortion field, then we should expect judgments of the material appearance of transparent object to correlate with changes in the distortion field. We put this hypothesis to the test in three experiments.

EXPERIMENT 1

In Experiment 1, we used a psychophysical scaling method to measure how the perceived material properties of a transparent object change as a function of changes in the physical RI, while holding all other properties of the scene constant.

Maximum likelihood difference scaling (MLDS) is a method that can be used for estimating the function relating a physical parameter (in this case refractive index) to its perceptual correlate [Maloney and Yang, 2003; Knoblauch & Maloney, 2008]. The subject is presented with two *pairs* of images (a total of four images), and must report which pair appears to be more different (*i.e.*, for which pair the difference between the two stimuli within the pair is greater). From the pattern of responses, it is possible to estimate the maximum likelihood perceptual scale that accounts for the data. We applied this method to the perception of materials that differed only in their RI, using stimuli spanning the same range as shown in Figure 2. We asked subjects to base their judgments on the apparent material that objects were made from. No reference was made to RI in the instructions.

METHODS

Stimuli. Stimuli consisted of computer generated images of a smooth, irregularly shaped 'pebble' of homogeneous refractive material inside a 20x20x20 cm textured box, as shown in Figure 2. The frame was included to enhance the perception of 3D distance and to provide a plausible means of support for the object.

The object shape was created in the application 3DSMax® by generating a unit 'Geosphere' primitive with approximately 10^5 triangular faces and applying various modifiers to create a flattened but curvaceous pebble. The scene was illuminated by 2 light sources. Refractive index was varied linearly from 1.1 to 2.3 in 10 steps.

Rendering was performed using the global illumination software DALI, by Henrik Wann Jensen. Trace depth (*i.e.*, the number of times a traced ray can spawn additional rays due to surface interactions) was set to 24. The rendering of caustics (light focused by refraction) was enabled, although no visible caustics were generated with this scene configuration. The renderings were tone-mapped for display with a gamma of 2, using the DALI program *vism*, and the final images were JPEGs of 320x320 pixels. Subjects viewed the images in a dark room on a 24" Sony Trinitron CRT monitor (1280 x 1024 resolution), at a distance of 55cm (set using a chin rest). The monitor gamma was calibrated to 1.8, leading to natural looking contrasts. The images were monocular (i.e. no stereoscopic depth information), and subjects viewed the screen with both eyes.

Subjects. 6 naïve observers from the Max Planck Institute subject database were paid to participate in the experiment. All participants had normal colour vision (Ishihara Test) and normal or corrected-to-normal acuity. Subject age ranged from 20-40 years.

Procedure. On each trial, subjects were presented with two pairs of stimuli simultaneously (*i.e.*, four images, or a 'quadruple'), and were asked to indicate in which pair the material composition of the objects appeared to be more different. The RIs of the four objects on each trial were drawn without replacement from ten possible values, such that all four objects had different RIs, but the physical intervals between pairs varied from trial to trial. Subjects viewed all ${}^{n}C_{k}$ combinations of refractive index (n=10, k=4, yielding ${}^{n}C_{k}$ = 210 trials) in random order. Subjects were given unlimited time to respond to each trial and entered their response by pressing one of two keys on the keyboard. Details of how the perceptual scale is inferred from the pattern of responses are provided in Maloney and Yang (2003).

RESULTS

In figure 3A we plot the estimated perceptual scales for six subjects, along with the mean across subjects.

There are two notable aspects to the curves. First, all subjects displayed a pronounced positively bowed perceptual scale, implying that a given change in RI has a larger effect on material appearance for small values of RI than for larger values of RI. This suggests that the visual system does not directly represent physical RI, but rather some perceptually transformed quantity related to RI (much as perceptual brightness is a transformed representation of intensity). Second, there are substantial individual differences, suggesting that subjects probably base their judgments on several cues that are weighted differently by different subjects.

If we plot the mean distortion field magnitude as a function of physical RI (Figure 3A), we see that it is also positively bowed for this range of stimuli. This suggests that some measure of distortion may be among the cues that subjects rely on to judge RI.

EXPERIMENTS 2 AND 3

A critical test of the role of distortion fields in human perception would be to find a case in which the cue predicts erroneous judgments of RI. This is not hard to find. Distortion fields are not only affected by refractive index but also by other attributes of the scene, most notably: (i) the shape of the refracting object and (ii) the distance in depth between the object and the background that is visible through it. This is a straightforward consequence of the ray optics of refraction. For example, when the backplane is moved further from the object, rays that diverge as they exit from the transparent object strike the backplane at a greater distance from one another, leading to a greater degree of compression in the image (or a greater degree of magnification for converging rays)[1]. Similar arguments apply when the thickness of the object is varied. In brief, the thicker the object, the further the rays travel through the refractive medium, leading to greater degrees of distortion

This provides us with a way to measure whether human vision relies to some extent on the degree of distortion when judging refractive properties of objects. By changing these extrinsic scene variables we can modify the distortion fields without affecting the physical RI. This predicts that subjects should misestimate refractive index if they do not correctly compensate for these changes in the distortion field. We put this to the test in two psychophysical experiments using an adjustment task.

---

[1] Due to Helmholtz reciprocity, all light rays can be traced in either direction. Here we use the convention of tracing rays from the eye into the scene (rather than from the light source). Thus, rays start at the eye, strike the near surface of the transparent object, refract and emerge from the rear surface. The angle between the emerging rays determines how much compression or rarefaction of the background pattern occurs for any given distance to the background. The further the background, the more the rays diverge or converge, leading to a concomitant increase in the magnitude of distortion.

METHODS

Stimuli. The stimuli were the same as in Experiment 1 except that two additional scene parameters were varied. In Experiment 2, the distance of the rear wall of the box varied linearly in five steps from 2.5cm to 17.5cm (i.e., 2.5, 6.25, 10, 13.75, 17.5cm, of which the Test stimuli was set to the first, second, fourth and fifth value, while the Match stimuli was always set to the third value—see Procedure subsection for clarification). In Experiment 3, the thickness of the object was varied by linearly scaling it along the z-axis (line of sight) by a factor ranging from 0.5 to 1.5 (i.e., 0.5, 0.75, 1, 1.25, 1.5). Again, the Test stimulus was set to the first, second, fourth or fifth value from this sequence, while the Match stimulus was always set to the third value. For comparison, the stimuli in Experiment 1 also used the middle value of these two ranges. Additionally, instead of 10 values of RI between 1.1 and 2.3, we used 128, allowing much finer variations in RI when the subjects adjusted RI.

Subjects. The same subjects were used as in Experiment 1. They conducted Experiments 2 and 3 immediately after Experiment 1 under the same viewing conditions. The order of Experiments 2 and 3 was randomly counterbalanced across subjects.

Procedure. On each trial, subjects were presented with two renderings simultaneously: a Test stimulus, whose parameters were selected by the computer, and a Match stimulus, whose refractive index could be continuously varied by the subject by moving the mouse (in practice, the mouse position specified which pre-rendered image was presented on the screen). The subjects were instructed to adjust the position of the mouse until the Match stimulus appeared to be made of the same material as the Test stimulus, while ignoring any additional perceived differences between the two images. RI was not explicitly mentioned or described to the subjects, although all participants agreed that moving the mouse changed something about the intrinsic appearance of the material. Importantly, either the distance to the backplane (Experiment 2) or the thickness of the refractive object (Experiment 3) was clamped at different values for the Test and Match stimuli. This allows us to measure the subject's ability to "ignore" or "discount" the contribution of this extrinsic scene variable from the appearance of the object. Specifically, on each trial, the Test stimulus was set at one of the four different values of backplane distance or thickness, while the Match stimuli for all conditions was the middle value of both parameters. The experimental method is analogous to asymmetric matching in surface color perception but with distance or thickness playing the role of the illumination or context (Krantz, 1968).

RESULTS

The mean responses of six observers are shown in Figure 3B (Experiment 2) and Figure 3C (Experiment 3). Each plot shows the data for four different levels of the extrinsic scene variable. If subjects were able to perfectly discount the effect of the distance to the backplane or the thickness of the object, then all four curves should lie on the diagonal. However, the data exhibits systematic biases. When the object is thin, RI is judged to be consistently lower than when the object is thick. Similarly, when the backplane is near, RI is judged to be lower than when the backplane is far, even though the layout of the scene obviously has nothing to do with the intrinsic material properties of the transparent object. In both cases, this is consistent with the effects of thickness and backplane distance on the magnitude of the distortion field. Reducing thickness, or bringing the backplane close to the object reduces distortions, leading to an underestimate of RI, while larger values lead to greater distortions, and a concomitant overestimation of RI.

Insert Figure 3 Here

There are substantial inter-subject differences in the extent of the misperception. One observer exhibited only a very weak effect in Experiment 2, while for others the effects were much stronger. As in Experiment 1, this probably represents the fact that subjects can rely on several cues to make their judgment. The extent of the misperception probably depends on how much they rely on the distortion field cue. Nevertheless, the pattern of results across subjects clearly suggests that when distorted refracted patterns are salient—as in our stimuli—subjects readily rely to some extent on the pattern of distortions to equate the material appearance, even if this leads to incorrect estimates of RI.

DISCUSSION

We have argued that the pattern of image distortions that occurs when a textured background is visible through a refractive object provides a key source of information that the brain can use to estimate an object’s intrinsic material properties. In spirit, this proposal is similar to other recent research on material perception that suggests that human vision relies on a range of simple but imperfect image measurements that correlate with material attributes [Nishida and Shinya, 1998; Fleming et al. 2003; Fleming and Bülthoff, 2005; Motoyoshi et al. 2007; Ho et al. 2008; although see also Anderson and Kim, 2009; Kim and Anderson, 2010, which challenge recent claims about the role of image statistics in material perception]. This can be contrasted with cases in which it is argued that the visual system effectively estimates and discounts the contribution of extrinsic scene variables (such as the illumination) to the image data [von

Helmholtz, 1867; Maloney and Wandell, 1986; D’Zmura and Iverson, 1993; Boyaci et al, 2003]. In all probability the brain uses a range of strategies depending on the available data and the difficulty of the ‘inverse optics’ computation. However, when computing the physically correct solution involves knowledge of the scene that cannot be readily estimated from the image (*e.g.*, the shape of the rear surface of the transparent object to estimate refractive index), the brain must make do with heuristics.

The distortion field is a ‘mid-level’ cue that involves comparing the relative scale of texture elements seen through the transparent object with those seen directly. There are at least two theoretical challenges to understanding how the visual system extracts this information from the image and uses it to derive a heuristic proxy for RI. First, we need to explain how the outputs of lower-level image measurements are combined to measure the relative local spatial scale of the texture. This might involve identifying individual texture elements and making a local estimate of their average size, which can be compared with the surroundings. Alternatively, the visual system might estimate spatial scale using the outputs of filters tuned to different spatial frequencies. Either way, it is worth noting that this computation is theoretically similar to measuring texture compression for the estimation of 3D shape from texture and thus similar mechanisms may play a role.

The second challenge is to explain how the local estimates of distortion magnitude are pooled into a global estimate of RI. Here, we have simply taken the arithmetic mean of the local estimates within the image region belonging to the object. However, it seems plausible that there may be some non-linear transformation of the local estimates of distortion, or that not all locations in the image might be given equal weight in the pooling operation. For example, if some locations yield unreliable estimates of the local texture scale (e.g. when the amount of compression is very extreme), these may contribute less to the global estimate than regions where the visual system can estimate the magnitude of distortion with high reliability. The fact that we find relatively large differences between subjects suggests that there are several cues that subjects weight differently, or that the pooling function may vary from subject to subject.

Although we believe the distortion field is an important source of information about solid transparent objects, we should emphasise that it is clearly not the only cue. From the Fresnel equations, we know that the extent of specular reflection varies as a function of optical density. This can be seen in Figure 1, where the more refractive objects also appear glossier than the less refractive ones. Thus, the visual system could also use gloss-related cues in the interpretation of transparent objects. Like the distortion field, such cues are also mid- to high-level, in the sense that, before the visual system can measure glossiness, the visual system must separate the image features that are visible through the object from those that are specularly reflected from the surface. How this might be done is still poorly understood.

By contrast to such mid- and high-level cues, we find that lower level image measurements, such as average contrast and luminance are poor predictors of the subjects' settings. This makes intuitive sense, because such quantities are strongly influenced by factors that are unrelated to the object of interest, such as the illumination in the scene. When taken in isolation luminance and contrast are known to be poor predictors of other material properties, such as surface albedo. In particular, for our stimuli, contrast (measured as either normalized or non-normalized pixel variance within the region of the pebble) varies non-monotonically as a function of refractive index. This is due to a combination of many factors, including increases in the amount of specular reflection and changes in the relative size of dark and light features from the background when distorted through the object. No simple transformation of contrast predicts the positively-bowed results of Experiment 1.

By contrast, average pixel intensity does vary in a positively bowed function with variations in RI. Thus, at first sight it might appear that we cannot rule out the possibility that subjects base their responses in the MLDS experiment on image intensity. However, the results of Experiment 2 are not consistent with this interpretation. When we move the backplane backwards, it gets dimmer (as it is further from the light source). This causes the average intensity of the pebble to *decrease* as a function of backplane distance. If subjects based their matches of RI on the average image intensity, this would predict that the pebble viewed against a more distant backplane should appear *less* refractive than it really is. This is the exact opposite of what we find in our experiments: perceived RI increases with the distance to the backplane. Thus, the results of Experiment 2 rule out mean image intensity as the main cue that subjects use when asked to judge RI.

In all likelihood there are many additional photometric and geometric cues to discover. For coloured transparent materials, such as amethyst or bottle glass, spatial variations in colour saturation across the object could also provide information about the shape and material of the object. Using more realistic physical models of transparency will help to reveal these additional sources of image information.

Finally, we should also note that it is possible to enjoy a vivid impression of light passing through an object even when no patterns are visible through the object. Translucent objects such as jelly, wax and cheese, and certain transparent objects such as intricately faceted crystals appear transparent even when we cannot see through them. Explaining how the visual system identifies that a given image gradient is caused by transmitted rather than reflected light, is surely one of the great outstanding challenges in the perception of material properties.

ACKNOWLEDGMENTS

We are deeply grateful to Henrik Wann Jensen for giving us access to his software DALI. RWF was supported by DFG Grant FL 624/1-1. LTM was supported in part by the Alexander-von-Humboldt Stiftung.

REFERENCES

Adelson, E. H. (1999). Lightness perception and lightness illusions. In *The New Cognitive Neurosciences*. 2nd edn. (ed. Gazzaniga, M.S.) pp. 339–351 (MIT Press, Cambridge, Massachusetts, 1999).

Adelson, E. H. (2001). On Seeing Stuff: the Perception of Materials by Humans and Machines. Proceedings of the. SPIE Vol. 4299, pp. 1–12, Human Vision and Electronic Imaging VI, Bernice E. Rogowitz; Thrasyvoulos N. Pappas; Eds.

Adelson, E. H., & Anandan, P. (1990). Ordinal characteristics of transparency. *Proceedings of the AAAI-90 Workshop on Qualitative Vision*, 77–81.

Anderson B. L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: the role of contour junctions. *Perception*. **26**(4): 419-453.

Anderson B.L., & Kim J (2009) Image statistics do not explain the perception of gloss and lightness. *Journal of Vision*, 9(11):10, 1-17.

Beck J, & Ivry R. (1988). On the role of figural organization in perceptual transparency. *Percept Psychophys*. **44**(6): 585-594.

Beck J, Prazdny K, & Ivry R. (1984). The perception of transparency with achromatic colors. *Percept Psychophys*. **35**(5): 407-422.

Boyaci, H., Maloney, L. T. & Hersh, S. (2003), The effect of perceived surface orientation on perceived surface albedo in three-dimensional scenes, *Journal of Vision*, **3**, 541-553.

D'Zmura, M. & Iverson, G. (1993). Color constancy. I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces. *J. Opt. Soc. Am. A* **10**, 2148–2165.

Fleming, R. W. & Bülthoff, H. H. (2005) Low-Level Images Cues in the Perception of Translucent Materials. *ACM Transactions on Applied Perception* **2**(3), 346–382

Fleming, R. W., Dror, R. O., & Adelson, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision*, **3**(5), 347-368.

Fleming, R. W., Torralba, A. & Adelson E. H. (2004). Specular reflections and the perception of shape. *Journal of Vision* **4**(9), 798–820.

Gerbino, W. (1994). Achromatic transparency. In A.L. Gilchrist, Ed. *Lightness, Brightness and Transparency*. (pp 215-255). Hove, England: Lawrence Erlbaum.

von Helmholtz, H. (1867/1962). *Helmholtz's Treatise on Physiological Optics*. New York: Dover (edited by J. P. C. Southall).

Ho Y. X., Landy M. S. & Maloney L. T. (2008). Conjoint measurement of gloss and surface texture. *Psychol Sci*. **19**(2):196–204

Kim J, Anderson B (2010) Image statistics and the perception of surface gloss and lightness. *Journal of Vision*, 10(9): 3.

Knoblauch, K. & Maloney, L. T. (2008), MLDS: Maximum likelihood difference scaling in R. *Journal of Statistical Software,* 25(2), 1-26.

Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt Brace.

Krantz, D. (1968), A theory of context effects based on cross-context matching, *J. Math. Psychol. 5*, 1-48.

Maloney, L. T. and Wandell, B. A. (1986). Color constancy: a method for recovering surface spectral reflectance. *J. Opt. Soc. Am. A* **3**, 29–33.

Maloney, L. T. & Yang, J. N. (2003). Maximum likelihood difference scaling. *Journal of Vision*, **3**(8):5, 573–585.

Metelli, F. (1970). An algebraic development of the theory of perceptual transparency. *Ergonomics*, **13**(1): 59-66.

Metelli, F. (1974a). Achromatic color conditions in the perception of transparency. In Eds. Macleod, R. B., & Pick, H. L. (Eds.) *Perception* (pp. 95–116). Ithaca, NY: Cornell University Press.

Metelli F. (1974b). The perception of transparency. *Scientific American*, **230**(4): 90–98.

Motoyoshi, I., Nishida, S., Sharan, L. & Adelson, E. H. (2007). Image statistics and the perception of surface qualities. *Nature* **447**, 206–209.

Nishida S. & Shinya, M. (1998). Use of image-based information in judgments of surface-reflectance properties," *J. Opt. Soc. Am. A* **15**, 2951–2965.

Robilotto R, Khang BG & Zaidi Q. (2002). Sensory and physical determinants of perceived achromatic transparency. *JOV*. **2**(5): 388-403.

Robilotto, R. and Zaidi, Q. (2004). Perceived transparency of neutral density filters across dissimilar backgrounds. *JOV*, **4**(3), 183–195.

Singh M, Anderson BL. (2002). Toward a perceptual theory of transparency. *Psychol Rev*. **109**(3): 492-519.

FIGURE LEGENDS

Figure 1. A. Most work on the perception of transparency uses a model inspired by neutral density filters, in which the image intensities are a linear combination of the reflectance of the transparent surface and the light transmitted through the filter. Note the X-Junction that occurs where the boundary of the filter crosses edge of the stripe in the background. **B**. Here we use a physically-based simulation of dielectric materials to study the perception of transparency. Note the distortion of the pattern visible through the transparent object. Despite the absence of X-junctions, the object appears vividly transparent.

Figure 2. Changing the refractive index of the object leads to changes in material appearance. Bottom right: a close-up of the region highlighted above, with vectors showing the displacement field for this object. The displacement field varies continuously due to the irregular shape of the object.

Figure 3. Data from the three experiments**. Left Column**: **A**. The estimated perceptual scales of RI for six subjects, along with the mean across subjects and the prediction derived from the distortion fields. **B**. Mean matching results of six observers, for four distances between the object and the backplane. The labels indicate the distances for the Test stimuli; Match stimuli were fixed at 10cm. **C.** Matching results for four thicknesses of the pebble. The labels indicate the relative thicknesses for the Test stimuli; Match stimuli were fixed at 1.0**. Right column**: Example stimuli from the matching experiments. The top and middle panels have the same RI, but different thicknesses. Note the illusory change in material appearance. Bottom panel is the average match that subjects made for the thickness = 1.5 condition, which is actually a significantly higher value of RI.

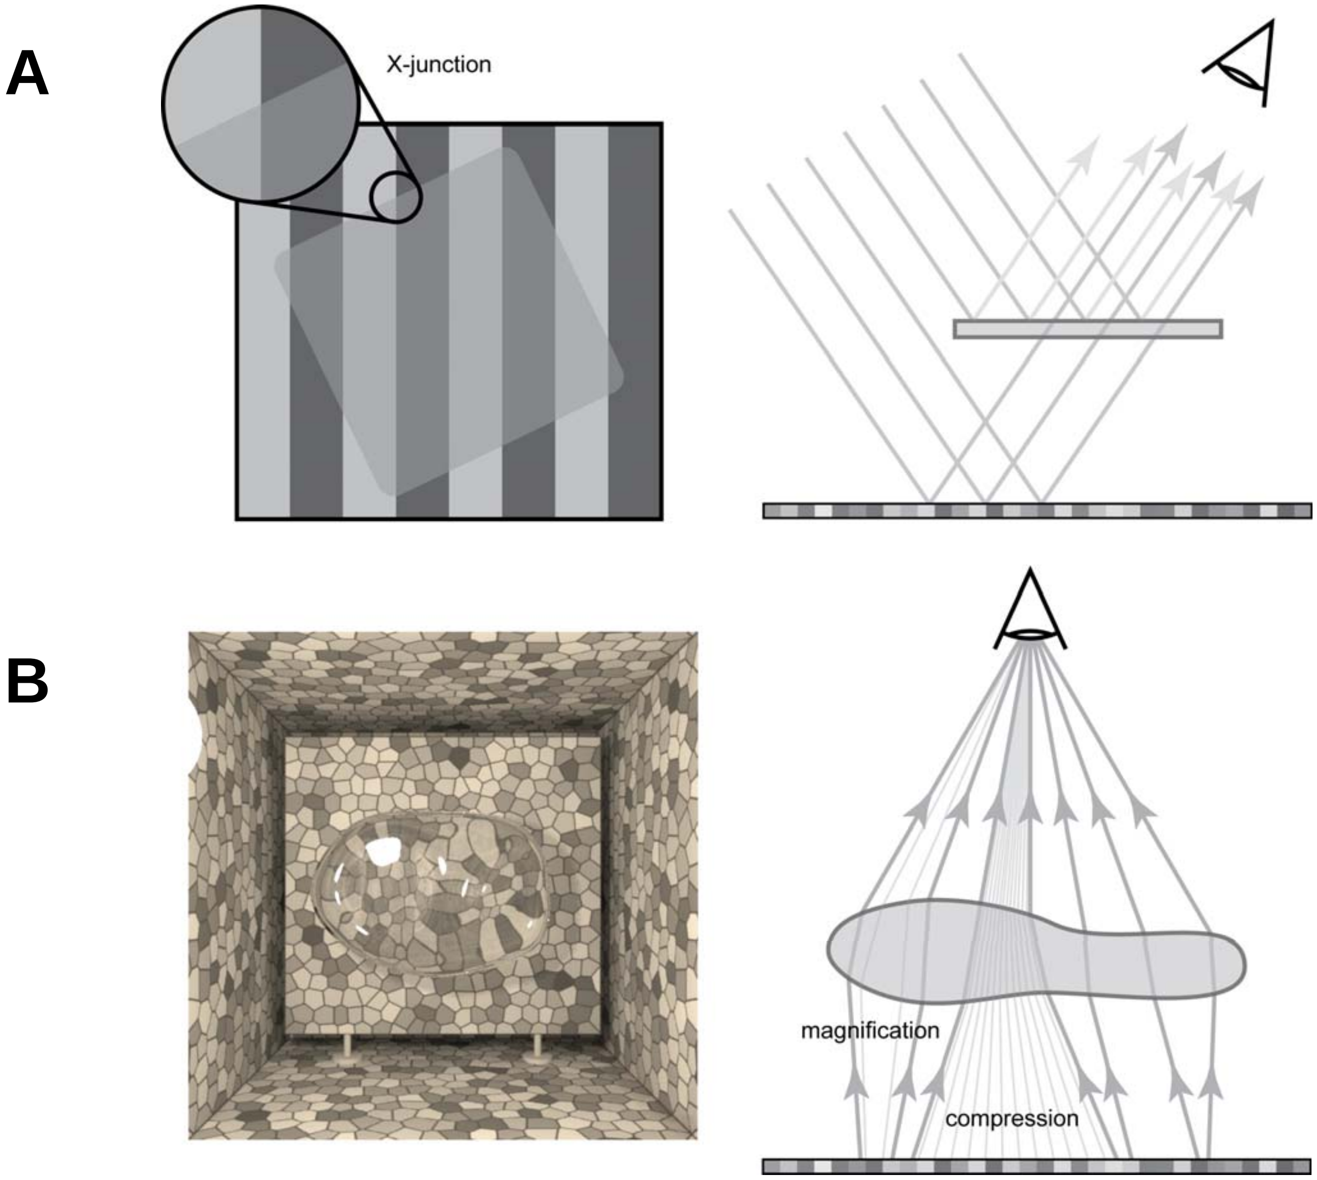

Fleming et al: Figure 1

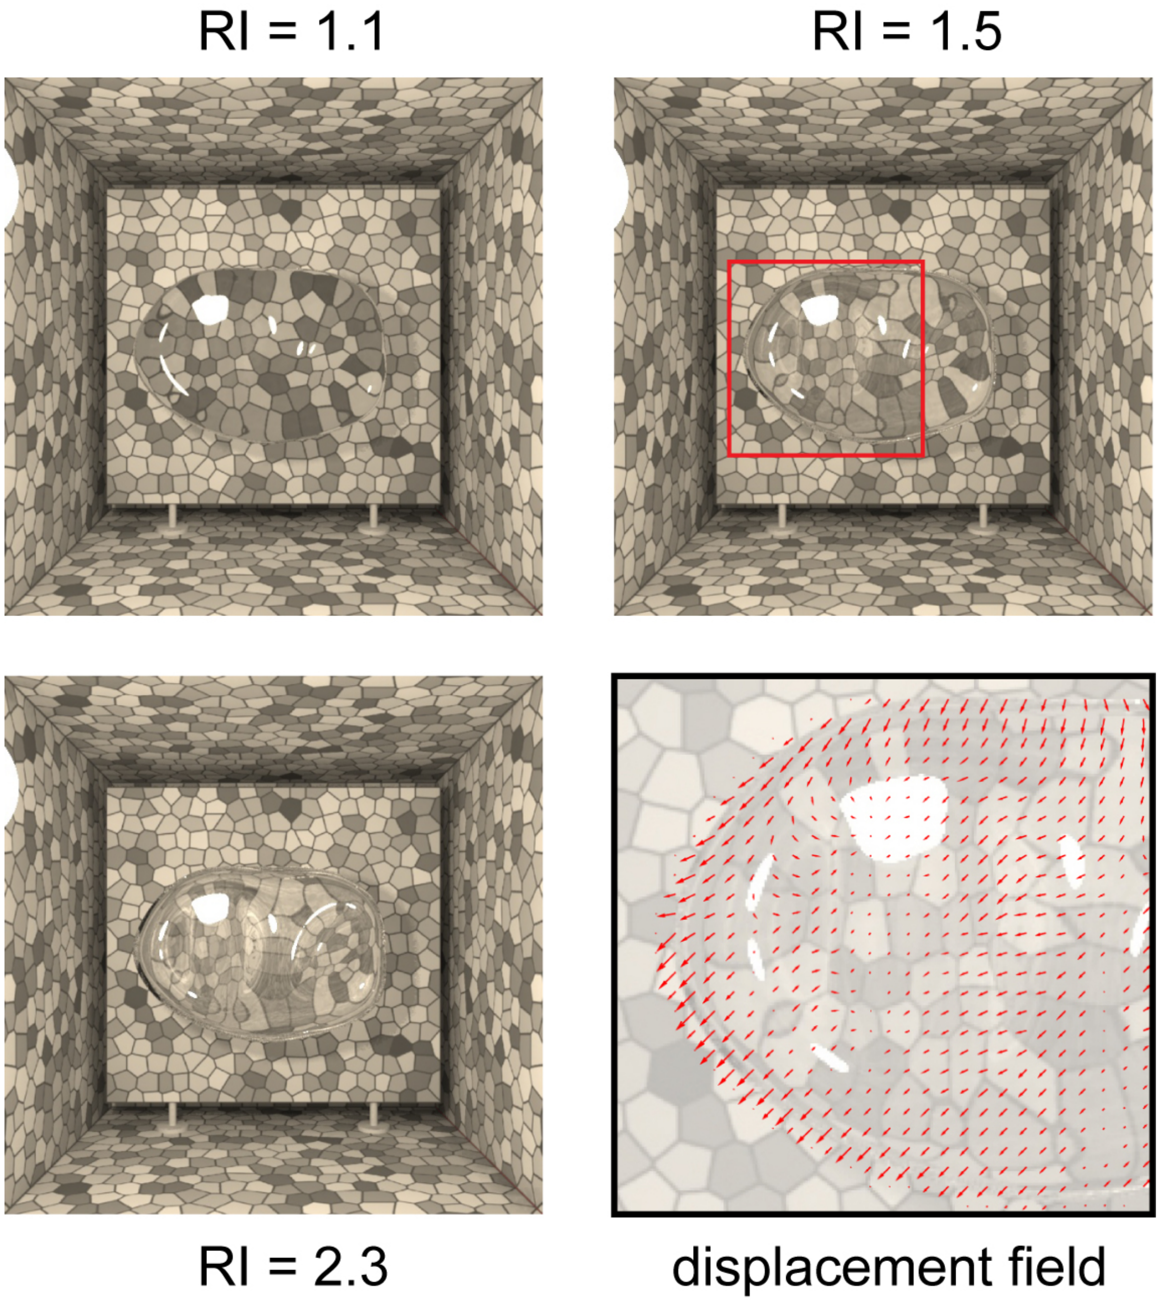

RI = 1.1

RI = 1.5

RI = 2.3

displacement field

Fleming et al: Figure 2

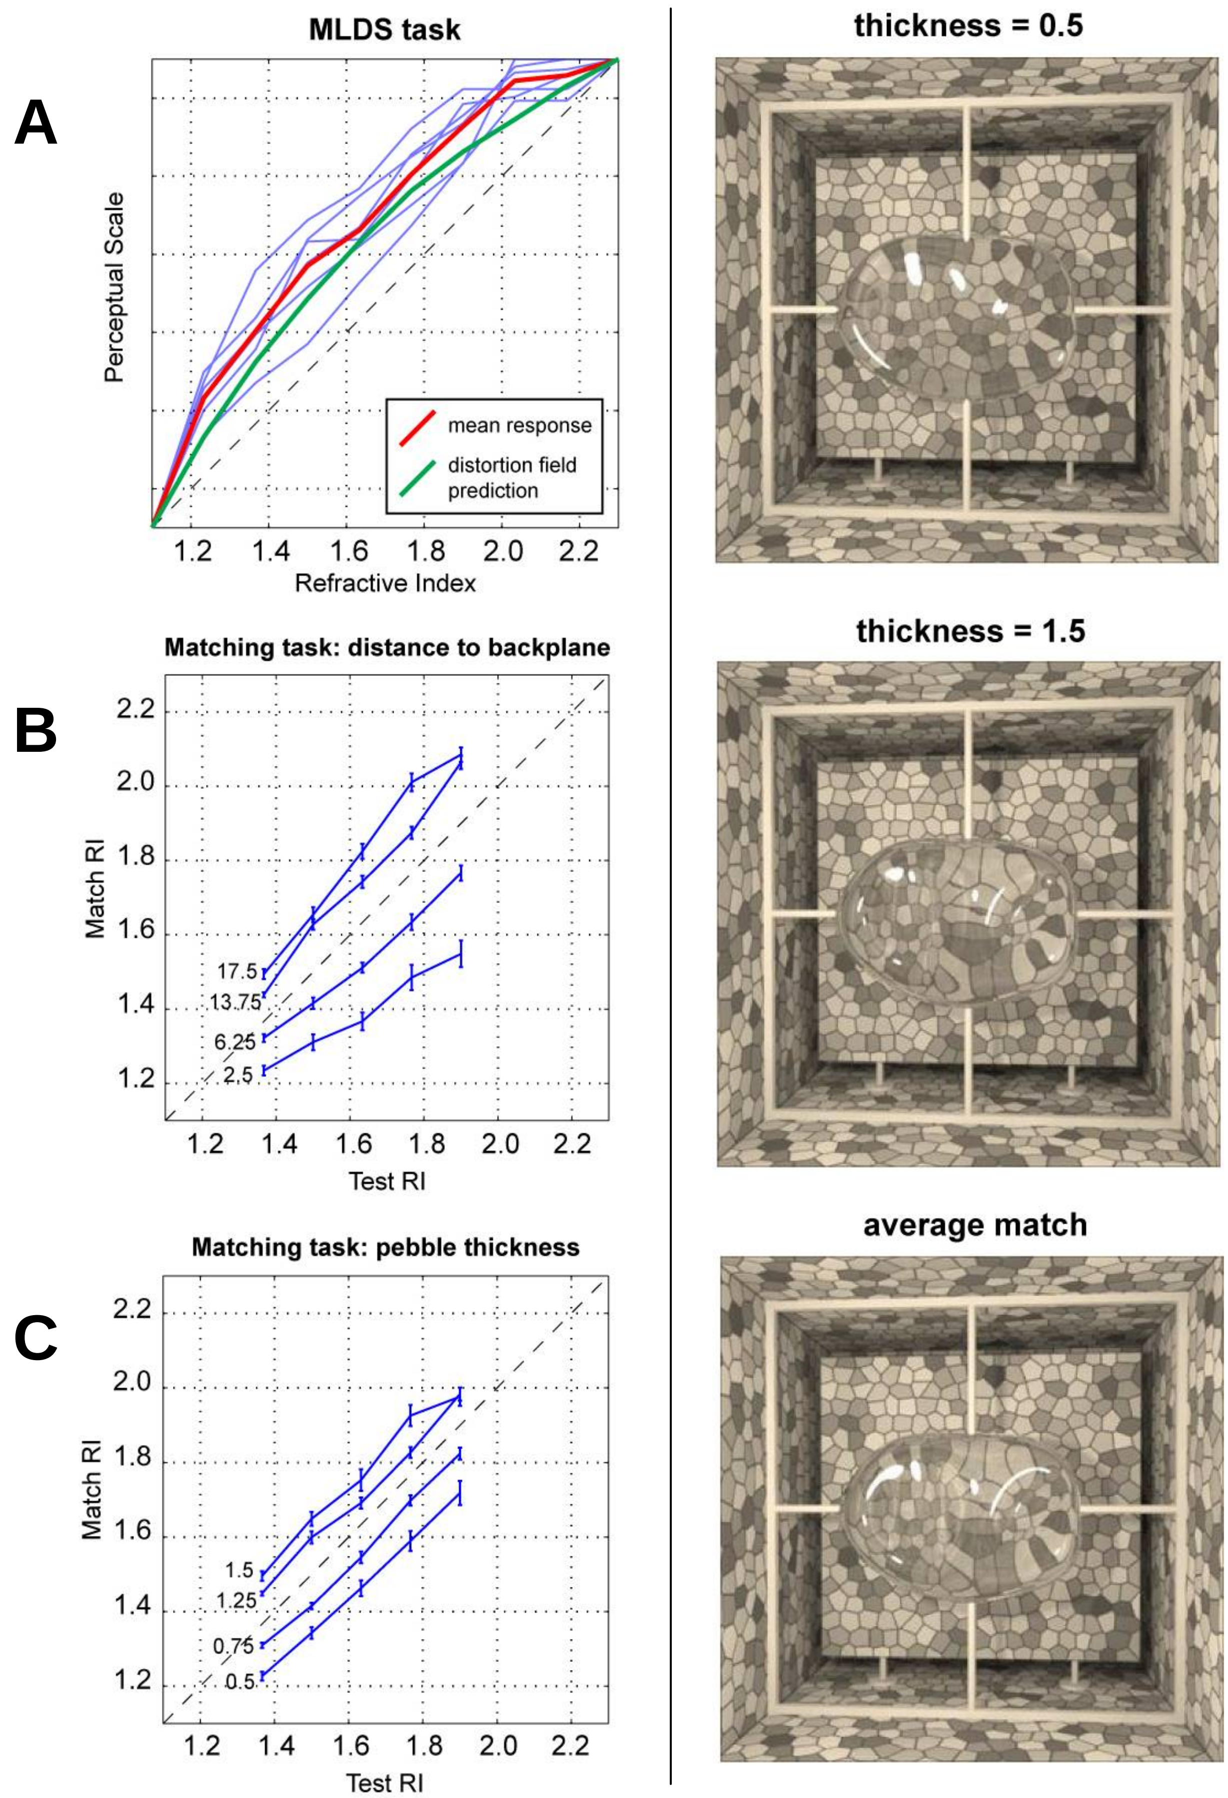

Fleming et al: Figure 3