# Visual Perception of Thick Transparent Materials

Roland W. Fleming, Frank Jäkel and Laurence T. Maloney
*Psychological Science* 2011 22: 812 originally published online 19 May 2011
DOI: 10.1177/0956797611408734

Research Article

Psychological Science
22(6) 812–820

DOI: 10.1177/0956797611408734
http://pss.sagepub.com

# Visual Perception of Thick Transparent Materials

**Roland W. Fleming[1,2], Frank Jäkel[1,3], and Laurence T. Maloney[4]**
[1]Max Planck Institute for Biological Cybernetics, Tübingen, Germany; [2]Department of Psychology, Justus-Liebig-Universität Gießen; [3]Institute of Cognitive Science, University of Osnabrück; and [4]Department of Psychology and Center for Neural Science, New York University

## 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. Recognizing transparent materials is especially challenging. Previous work on the perception of transparency has focused on objects composed of flat, infinitely thin filters. In the experiments reported here, we considered 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 signaling 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 derived from the distortion field induced by thick transparent objects, and we provide experimental evidence that cues arising from the distortion field predict both the successes and the failures of human perception in judging refractive indices.

## Keywords

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

Received 9/15/10; Revision accepted 1/27/11

Transparent objects, such as gemstones and glass, can be fascinating and beautiful to look at. However, they pose a formidable problem for the visual system. Under typical viewing conditions, light passing through such objects refracts and reflects several times before reaching the eye; this process yields 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 (Gerbino, 1994; Metelli, 1970, 1974a, 1974b; Robilotto, Khang, & Zaidi, 2002; Robilotto & Zaidi, 2004; Singh & Anderson, 2002; von Helmholtz, 1867/1962). Almost all research to date models transparent objects as ideal thin, neutral-density filters perpendicular to the line of sight (Fig. 1a). The primary effect of such transparent surfaces on the retinal image is to modify the intensities or chromaticities of patterns visible through them. Specific photometric and geometric preconditions for the perception of transparency have been derived from this assumption. For example, it is commonly argued that transparency creates *contour junctions* in the image (i.e., locations at which a number of contours meet in the image; Adelson & Anandan, 1990; Anderson, 1997; Beck & Ivry, 1988; Beck, Prazdny, & Ivry, 1984). A type of contour junction known as an X junction occurs when an edge of a background layer passes behind a transparent foreground object and thus undergoes a reduction in contrast (Figs. 1a and 1c).

However, many transparent objects that we encounter on a daily basis, such as ice cubes or chunks of glass, are thick, refractive bodies (Fig. 1b). The physics of light transport through such real-world objects is markedly more complex than through 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 transform not only 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 portion of the light is reflected, and the remaining light refracts as it enters the body of the object (Fig. 1d).

**Corresponding Author:**
Roland W. Fleming, Justus-Liebig-Universität Gießen, FB06 – Psychologie, Otto-Behaghel-Str. 10/F2 -338, 53594 Gießen, Germany
E-mail: roland.w.fleming@psychol.uni-giessen.de

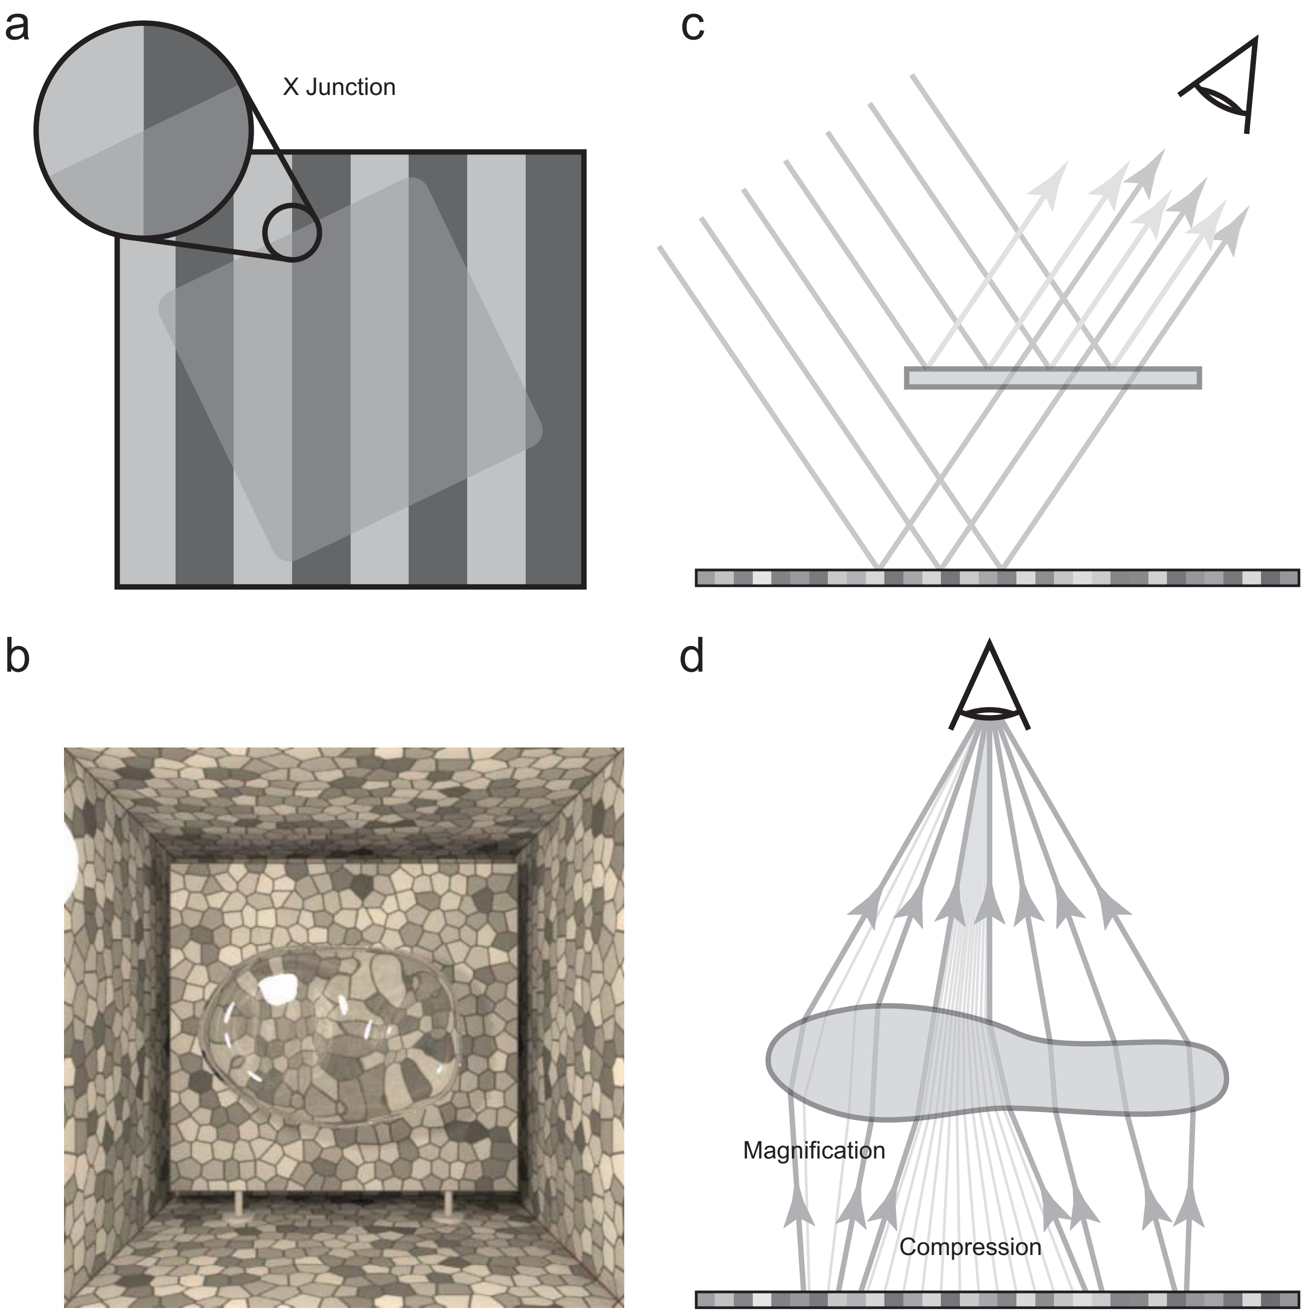


**Fig. 1.** Theoretical underpinnings of the study and an example stimulus used in the three experiments. An X junction (a) occurs at the boundary of a transparent foreground object when an edge of a background layer passes behind the object and thus undergoes a reduction in contrast. The experimental stimulus (b) consisted of a computer-generated image of a smooth, pebble-shaped, transparent blob of homogeneous refractive material inside a 20-cm × 20-cm × 20-cm textured box. The traditional approach to studying the perception of transparency has focused on thin neutral-density filters (c) that do not create distortions in the background behind the object. Refraction through thick transparent materials, such as our blob stimulus (d), leads to varying degrees of compression and magnification of the transmitted background patterns.

The angle of refraction depends on (a) the local geometry of the surface and (b) an intrinsic property of the material: its *refractive index*. Light exiting the objects is similarly transformed.

Using a computer-graphics simulation of refraction, it is possible to systematically modify an object's refractive index while holding constant other scene variables, such as lighting and object shape. For example, the refractive index of the object shown in Figures 2a through 2c varies from 1.1 to 2.3, and this variation causes the observer to experience a concomitant change in the apparent material properties of the object. As its refractive index increases, the object appears more lustrous and substantial. Refraction can evidently play a key role in the appearance of many transparent objects.

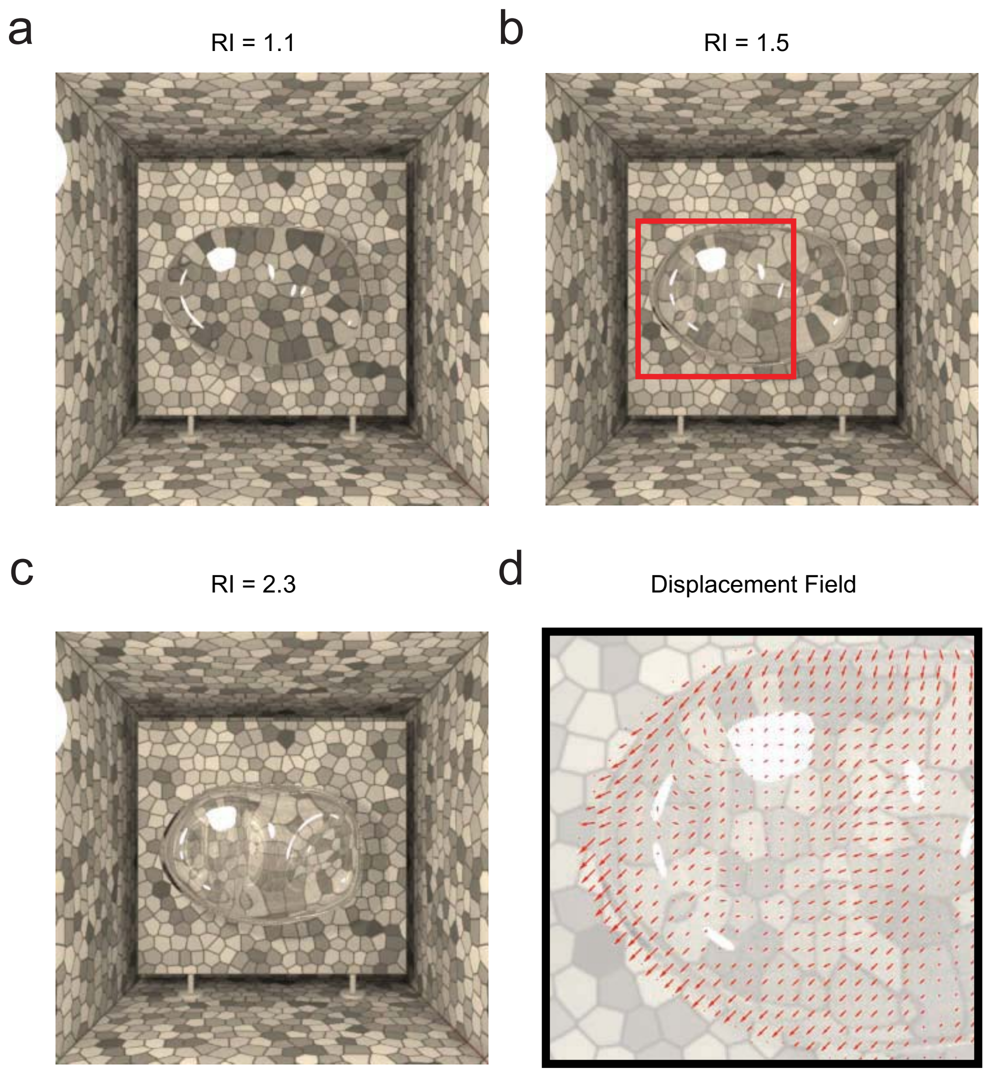


**Fig. 2.** Characteristics of the blob stimulus used in Experiments 1 through 3. The stimulus is shown at three different levels of refractive index (RI; a–c), each of which creates a different perception of the object's material appearance. The vectors (arrows) visible in the close-up (d) of the region outlined in (b) show the continuous variation in the stimulus's displacement field as a result of the object's irregular shape.

As Figure 1d shows, 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, observers 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 an object's refractive index. Researchers 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 $p_i = (x, y)$ and its position when refracted through the transparent object is $p_r = (x', y')$, then a vector field can be defined in image space as $D(x, y) = p_r - p_i$, which measures the displacement of all features in the background when seen through the transparent object. This is called the *displacement field* (Fig. 2d).

Given that the visual system does not have access to the initial, nonrefracted positions of features in the background ($p_i$), 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 with its local context. The relative magnitude of compression can be captured mathematically as the divergence of the displacement field, $d(x, y) = \nabla \cdot D(x, y)$, which is called the *distortion field*. If $d$ is positive, the displacement field is diverging, leading to a local magnification of the refracted pattern; if $d$ 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, Torralba, & Adelson, 2004). We suggest that a related strategy could provide the brain with information about refractive index.

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 derived from the distortion field, then we should expect judgments of the material appearance of transparent objects 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 changed as a function of changes in its physical refractive index, while all other properties of the scene were held constant. We used maximum-likelihood difference scaling (MLDS) to estimate the function relating a physical parameter (in this case, refractive index) with its perceptual correlate (Knoblauch & Maloney, 2008; Maloney & Yang, 2003). In MLDS, the subject is presented with two pairs of images (a total of four images) and asked to report which pair appears to have a greater within-pair difference. 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 refractive index. We asked subjects to base their judgments on the apparent material that the objects were made from. No reference was made to refractive index in the instructions.

### Method

***Stimuli.*** Stimuli consisted of computer-generated images of a smooth, pebble-shaped, transparent blob of homogeneous refractive material inside a 20-cm × 20-cm × 20-cm textured box, as shown in Figure 2. We included the box frame around the blob to enhance the perception of three-dimensional distance and to provide a plausible means for the object's support.

The object shape was created with 3ds Max modeling software (Autodesk, London, England) by generating a unit Geosphere primitive with approximately $10^5$ triangular faces and applying various modifiers to create a flattened but curvaceous blob. The scene was illuminated by two 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, provided by Henrik Wann Jensen, University of California, San Diego. Trace depth (i.e., the number of times a traced ray can spawn additional rays because of 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 .jpg files of 320 × 320 pixels. The images were monocular (i.e., they had no stereoscopic depth information).

***Subjects.*** Six naive observers recruited from the Max Planck Institute subject database were paid to participate in the experiment. All participants had normal color vision (as determined using the Ishihara Test) and normal or corrected-to-normal visual acuity. Subject age ranged from 20 to 40 years.

***Procedure.*** Subjects viewed all stimuli in a dark room on a 24-in. Sony Trinitron CRT monitor (1,280 × 1,024 resolution) at a distance of 55 cm (using a chin rest). The monitor gamma was calibrated to 1.8, leading to natural-looking contrasts. Subjects viewed the screen with both eyes.

On each trial, subjects were presented with two pairs of stimuli simultaneously (i.e., four images, or a "quadruple") and asked to indicate in which pair the material composition of the objects appeared to have a greater within-pair difference. The refractive indices of the four objects on each trial were drawn without replacement from ten possible values, such that all four objects had different refractive indices, but the physical intervals between pairs varied from trial to trial. Subjects viewed all combinations of refractive index in random order across 210 trials. 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

Figure 3a plots the estimated perceptual scales for the 6 subjects, along with the mean across subjects. There are two notable aspects of the results. First, all subjects displayed a pronounced positively bowed perceptual scale; this implies that a small change in refractive index had a larger effect on material appearance than a large change did. This suggests that the visual system does not directly represent physical refractive index, but rather some perceptually transformed quantity related to refractive index (much as perceptual brightness is a transformed representation of intensity). Second, there were substantial individual differences in the data, suggesting that subjects based their judgments on several cues but weighted each cue differently. Plotting the mean distortion-field magnitude as a function

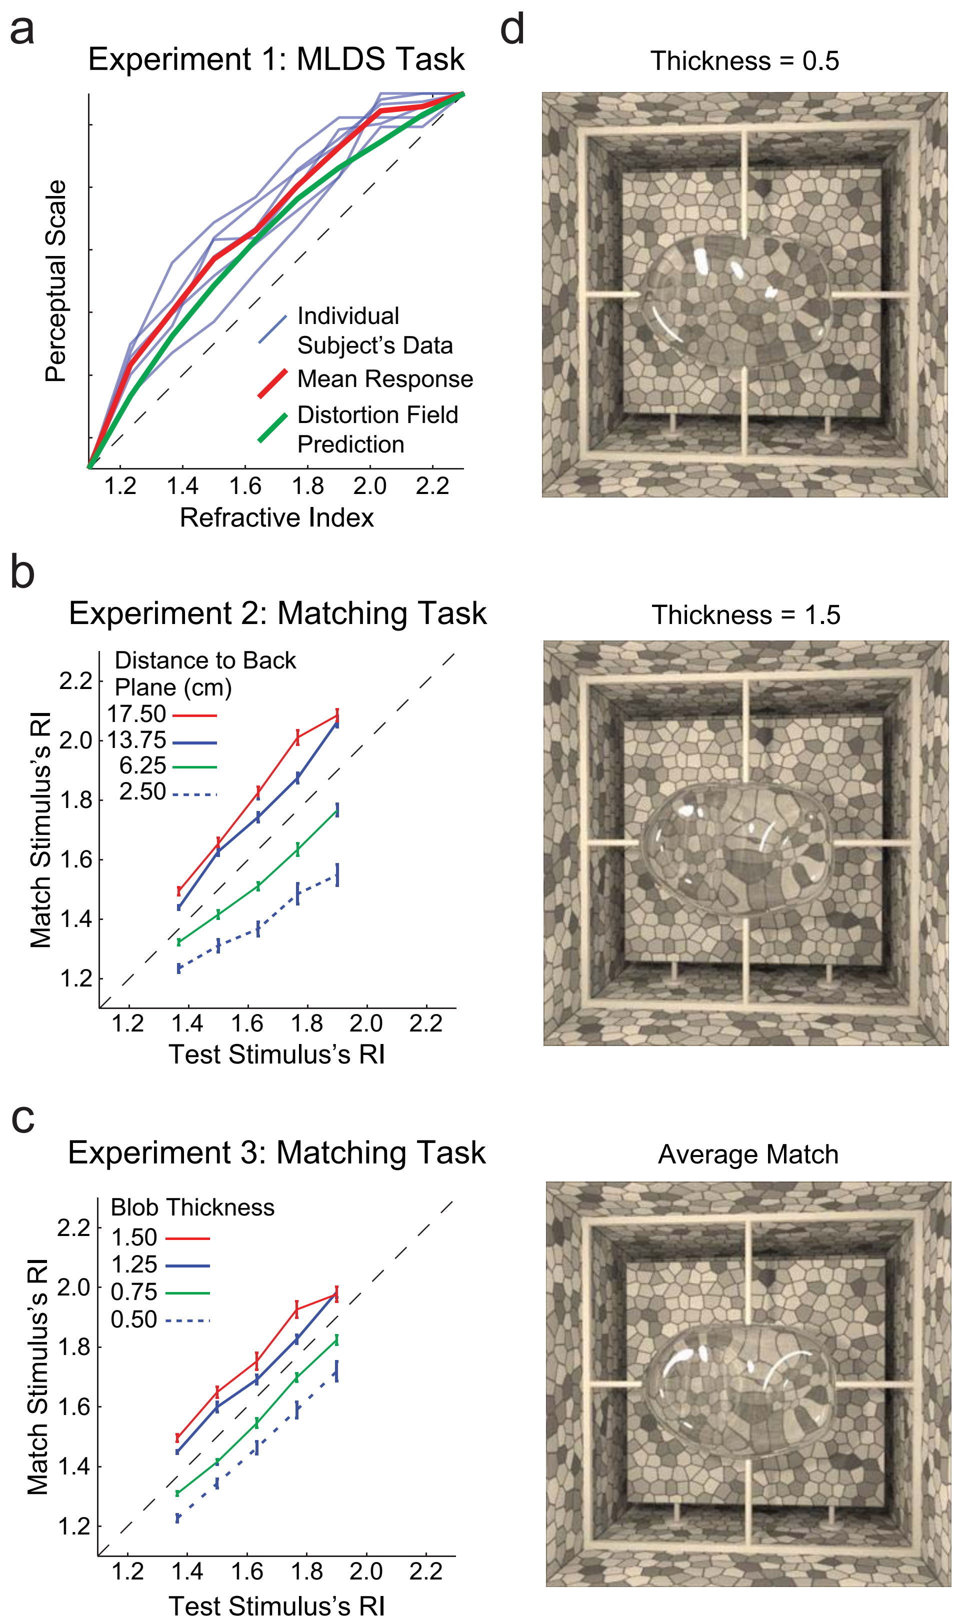


**Fig. 3.** Results from the three experiments (left column) and example stimuli and average selected match stimulus from Experiment 3 (right column). Estimated perceptual scales as a function of the stimulus's refractive index (RI) are shown for the 6 subjects in the maximum-likelihood-difference-scaling (MLDS) task of Experiment 1 (a). The mean across subjects and the prediction derived from the distortion fields are also plotted. Mean matching results for 6 observers are shown separately for (b) Experiment 2 and (c) Experiment 3. In each graph, the RI selected for the match stimulus is plotted as a function of the RI of the test stimulus. The dashed diagonal shows where the true matches would lie. Results of Experiment 2 are shown for four sample distances between the test stimulus and the back plane; the distance of the match stimuli was kept constant at 10 cm. Results of Experiment 3 are shown for four sample thicknesses of the test stimuli; the thickness of the match stimuli was kept constant at 1.0. Error bars show standard errors of the mean. The stimuli used in Experiment 3 (d) had a constant RI, but varied in thickness. The example stimuli shown have thicknesses of 0.5 and 1.5. The average RI that subjects selected for the match stimulus when the test stimulus had a thickness of 1.5 is shown in the bottom illustration.

of physical refractive index (Fig. 3a) revealed 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 refractive index.

## 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 refractive index. This is not hard to do. Distortion fields are affected by other attributes of the scene than refractive index, most notably (a) the shape of the refracting object and (b) 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 back plane is moved farther from the object, rays that diverge as they exit the transparent object strike the back plane at a greater distance from one another, and this leads 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 farther the rays travel through the refractive medium, and this leads to greater degrees of distortion.

These principles of refraction provide 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, the distortion fields can be modified without affecting the physical refractive index. This procedure should induce subjects to misestimate refractive index if they do not correctly compensate for these changes in the distortion field. We put this hypothesis to the test in two psychophysical experiments using an adjustment task, in which subjects adjusted the refractive index of a match stimulus to make it appear the same as a test stimulus, despite differences in other properties of the scene.

### *Method*

***Stimuli.*** The stimuli in Experiments 2 and 3 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 2.5 cm to 17.5 cm (i.e., 2.5 cm, 6.25 cm, 10 cm, 13.75 cm, and 17.5 cm, of which the test stimuli were set to the first, second, fourth, or fifth values, and the match stimuli were always set to the third value). 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; see Fig. 3d for examples). Again, the test stimuli were set to the first, second, fourth, or fifth value from this sequence, and the match stimuli were always set to the third value. In comparison, all of the stimuli in Experiment 1 used the middle value of these two ranges. Additionally, instead of 10 values of refractive index between 1.1 and 2.3, we used 128, thus allowing much finer variations in refractive index than was possible in Experiment 1.

***Subjects.*** The same subjects were used as in Experiment 1. They participated in Experiments 2 and 3 immediately after Experiment 1 under the same viewing conditions.

***Procedure.*** The order of Experiments 2 and 3 was randomly counterbalanced across subjects. On each trial, subjects were presented with two renderings simultaneously: a test stimulus, whose parameters were selected by the computer, and a match stimulus, which was adjusted by the subjects. Subjects continuously varied the refractive index of the match stimulus by moving a mouse (in practice, the mouse position specified which of the 128 prerendered images was presented on the screen). 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. Refractive index 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.

It is important to note that either the distance to the back plane (Experiment 2) or the thickness of the refractive object (Experiment 3) was set to different values for the test and match stimuli. This allowed us to measure the subject's ability to ignore or discount the contribution of this extrinsic scene variable to the appearance of the object. Specifically, on each trial, the test stimulus was set at one of the four different values of back-plane distance or thickness, and 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 illumination or context (Krantz, 1968).

### *Results*

The mean responses of the 6 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 had been able to perfectly discount the effect of the distance to the back plane or the thickness of the object, then all four curves should lie on the diagonal. However, the data exhibit systematic biases. When the object was thin, refractive index was judged to be consistently lower than when the object was thick. Similarly, when the back plane was near, refractive index was judged to be lower than when the back plane was far, even though the layout of the scene obviously had nothing to do with the intrinsic material properties of the transparent object. In both cases, these results are consistent with the effects of thickness and back-plane distance on the magnitude of the distortion field. Reducing thickness, or bringing the back plane closer to the object, reduced distortions, and this led to an underestimation of refractive index. In contrast, larger values led to greater distortions and a concomitant overestimation of refractive index.

There were substantial intersubject differences in the extent of the misperception. One observer exhibited only a very weak effect in Experiment 2, but the effects were much stronger for other subjects. As in Experiment 1, this probably reflects the fact that subjects can rely on several cues to make their judgment. The extent of the misperception probably depends on how much subjects 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 refractive index.

## Discussion

We argue 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 the suggestions of other recent researchers on material perception, who have suggested that human vision relies on a range of simple but imperfect image measurements that correlate with material attributes (Fleming & Bülthoff, 2005; Fleming, Dror, & Adelson, 2003; Ho, Landy, & Maloney, 2008; Motoyoshi, Nishida, Sharan, & Adelson, 2007; Nishida & Shinya, 1998; although see Anderson & Kim, 2009, and Kim & Anderson, 2010, for a challenge of recent claims about the role of image statistics in material perception).

This proposal can be contrasted with theories in which it is argued that the visual system effectively estimates and discounts the contribution of extrinsic scene variables (such as illumination) to the image data (Boyaci, Maloney, & Hersh, 2003; D'Zmura & Iverson, 1993; Maloney & Wandell, 1986; von Helmholtz, 1867/1962). 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., using 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 midlevel cue that involves comparing the relative scale of texture elements seen through the transparent object with the elements 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 refractive index. 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 then 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 three-dimensional 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 refractive index. In the experiments reported here, we simply took the arithmetic mean of the local estimates within the image region belonging to the object. However, it seems plausible that there may be some nonlinear 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 do regions in which the visual system can estimate the magnitude of distortion with high reliability. The fact that we found 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 emphasize that it is clearly not the only cue that the brain uses to estimate refractive index. 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 2, in which 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. As with distortion-field cues, such cues are also mid to high level, in the sense that before the visual system can measure glossiness, it must separate the image features that are visible through the object from those that are specularly reflected from its surface. How this might be done is still poorly understood.

In contrast with such mid- and high-level cues, lower-level image measurements, such as average contrast and luminance, are poor predictors of the objects' 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 the stimuli we used in the present experiments, contrast (measured as either normalized or nonnormalized pixel variance within the region of the blob) varied nonmonotonically as a function of refractive index. This was 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 predicted the positively bowed results of Experiment 1.

By contrast, average pixel intensity does vary in a positively bowed function with variations in refractive index. Thus, at first sight, it might appear that we cannot rule out the possibility that subjects based their responses in Experiment 1 on image intensity. However, the results of Experiment 2 are not consistent with this interpretation. When we moved the back plane away from the blob, it got dimmer (as it was farther from the light source). This caused the average intensity of the blob to decrease as a function of back-plane distance. If subjects based their matches of refractive index on the average image intensity, then the blob viewed against a more distant back plane should have appeared less refractive than it really was. This is the exact opposite of what we actually found: Perceived refractive index increased as the distance to the back plane increased. Thus, the results of Experiment 2 rule out mean image intensity as the main cue that subjects use when asked to judge refractive index.

In all likelihood, there are many additional photometric and geometric cues to discover. For colored transparent materials, such as amethyst or bottle glass, spatial variations in color 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 allowing us to use his Dali software.

## Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

## Funding

Roland W. Fleming was supported by German Research Foundation (DFG) Grant FL 624/1-1. Laurence T. Maloney was supported in part by the Alexander-von-Humboldt Stiftung.

## Note

1. Because of Helmholtz reciprocity, all light rays can be traced in either direction. In this example, 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 farther the background, the more the rays diverge or converge, and this leads to a concomitant increase in the magnitude of distortion.

## References

Adelson, E. H., & Anandan, P. (1990). Ordinal characteristics of transparency. In P. Kahn (Ed.), *Proceedings of the AAAI-90 Workshop on Qualitative Vision* (pp. 77–81). San Mateo, CA: Morgan Kaufmann.

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

Anderson, B. L., & Kim, J. (2009). Image statistics do not explain the perception of gloss and lightness. *Journal of Vision*, *9*(11), Article 10. Retrieved from http://www.journalofvision.org/content/9/11/10

Beck, J., & Ivry, R. (1988). On the role of figural organization in perceptual transparency. *Perception & Psychophysics*, *44*, 585–594.

Beck, J., Prazdny, K., & Ivry, R. (1984). The perception of transparency with achromatic colors. *Perception & Psychophysics*, *35*, 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*(8), Article 2. Retrieved from http://www.journalofvision.org/content/3/8/2.full

D'Zmura, M., & Iverson, G. (1993). Color constancy. I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America*, *10*, 2148–2165.

Fleming, R. W., & Bülthoff, H. H. (2005). Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception*, *2*, 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), Article 3. Retrieved from http://www.journalofvision.org/content/3/5/3

Fleming, R. W., Torralba, A., & Adelson, E. H. (2004). Specular reflections and the perception of shape. *Journal of Vision*, *4*(9), Article 10. Retrieved from http://www.journalofvision.org/content/4/9/10

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

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

Kim, J., & Anderson, B. L. (2010). Image statistics and the perception of surface gloss and lightness. *Journal of Vision*, *10*(9), Article 3. Retrieved from http://www.journalofvision.org/content/10/9/3.full

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

Krantz, D. (1968). A theory of context effects based on cross-context matching. *Journal of Mathematical Psychology*, *5*, 1–48.

Maloney, L. T., & Wandell, B. A. (1986). Color constancy: A method for recovering surface spectral reflectance. *Journal of the Optical Society of America*, *3*, 29–33.

Maloney, L. T., & Yang, J. N. (2003). Maximum likelihood difference scaling. *Journal of Vision*, *3*(8), Article 5. Retrieved from http://www.journalofvision.org/content/3/8/5

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

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

Metelli, F. (1974b). The perception of transparency. *Scientific American*, *230*, 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. *Journal of the Optical Society of America*, *15*, 2951–2965.

Robilotto, R., Khang, B. G., & Zaidi, Q. (2002). Sensory and physical determinants of perceived achromatic transparency. *Journal of Vision*, *2*(5), Article 3. Retrieved from http://www.journalofvision.org/content/2/5/3.full

Robilotto, R., & Zaidi, Q. (2004). Perceived transparency of neutral density filters across dissimilar backgrounds. *Journal of Vision*, *4*(3), Article 5. Retrieved from http://www.journalofvision.org/content/4/3/5.full

Singh, M., & Anderson, B. L. (2002). Toward a perceptual theory of transparency. *Psychological Review*, *109*, 492–519.

von Helmholtz, H. (1962). *Helmholtz's treatise on physiological optics* (J. P. C. Southall, Ed.). New York, NY: Dover (Original work published 1867)