Are optical distortions used as a cue for material properties of thick transparent objects?

Nick Schlüter
Institut für Psychologie,
Christian-Albrechts-Universität zu Kiel, Kiel, Germany

Franz Faul
Institut für Psychologie,
Christian-Albrechts-Universität zu Kiel, Kiel, Germany

Fleming, Jäkel, and Maloney (2011) asked subjects to match perceived material properties of thick, clear transparent objects in photo-realistic scenes by adjusting the refractive index. They found approximate correspondence between standard and test objects but also large systematic deviations. Nevertheless, they concluded that estimated refraction is used as an indicator for material properties of light-transmitting objects and emphasized the role of object-induced background distortions in this process. This, however, seems not plausible, because the necessary information for inferring the refractive index from distortions—for example, the object's exact shape and thickness, its background distance, and the undistorted background—was not available in their situation. A more plausible alternative explanation is that the subjects did not match estimated refractive indices, but instead performed simple similarity matches based directly on image attributes related to background distortions or specular reflections. We tested this hypothesis in a similar matching experiment in which it was possible to predict the refractive index for a similarity match based on background distortions and for a similarity match based on specular reflections. Our subjects always chose a value between these two predictions. The specular reflection tends to be the dominant factor as soon as it becomes clearly noticeable. Our findings are compatible with the assumption that the subjects tried to find a compromise between two image-based similarity criteria. They do not seem to be consistent with the assumption that the matches are made on the basis of internal estimates of refractive indices.

Introduction

Our daily experience shows that we are able to perceive material properties of objects solely on the basis of visual information. We can effortlessly tell if objects are soft or hard, stable or fragile, rough or smooth. Objects that transmit light, like amber or icicles, are especially fascinating, because they impress us with their play of colors, distortions, and reflections. From a theoretical point of view, such transparent objects are especially interesting because they present the visual system with problems that do not occur with opaque objects: The retinal images of opaque objects depend mainly on their intrinsic material properties and the prevailing illumination, whereas those of transparent objects are also influenced by properties of the background that is visible through them. This raises the question, to what extent the visual system uses information related to the background to infer material properties of transparent objects.

Historically, investigations of perceptual transparency arose from work on color perception. It is therefore not surprising that mainly proximal color and brightness regularities were examined as cues for transparent objects (Metelli, 1970; Singh & Anderson, 2002). In filter models of perceptual transparency, potential color-related regularities were derived from the spectral transmission properties of optical filters (Beck, 1978; Beck, Prazdny, & Ivry, 1984; Faul & Ekroll, 2002, 2011). To investigate transmission-related properties that lead to color changes of the background in isolation, highly simplified situations were used in which flat transparent filters were viewed frontally under a homogenous illumination (Figure 1a). If one considers more complex objects under more natural viewing conditions, then it is obvious that the transparency impression depends on additional material properties. This is, for instance, demonstrated by highly transmissive objects like a colorless wine glass that changes the spectral composition of the transmitted light only slightly but can nevertheless elicit strong impressions of transparency (Figure 1b). Potential cues that may be used to detect transparency

Schlüter, 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, http://www.journalofvision.org/content/14/14/2, doi:10.1167/14.14.2.

doi: 10.1167/14.14.2 Received May 26, 2014; published December 4, 2014 ISSN 1534-7362

Figure 1. Different types of transparent objects. (a) Flat transparent objects, viewed frontally under a homogeneous illumination are well-suited for investigating the perception of transmission-related material properties. (b) More complex transparent objects reveal that the impression of transparency does not only depend on transmission-related properties, but also on other material characteristics, for example, the refractive properties. (c) Massive, relatively thick objects with curved surfaces without internal absorption or self-occlusion are well-suited for isolating refraction-related properties.

in these cases may relate to the fact that light-transmitting objects distort the transmitted background texture and specularly reflect the environment on their surface. Both characteristics depend directly on a material property, the refractive index, and are therefore a natural basis from which properties of transparent objects may be inferred. It is important to note, however, that the refractive index is but one of several factors that influence distortions and reflections.

Fleming, Jäkel, and Maloney (2011) were the first to systematically investigate the material perception of such complex transparent objects. In particular, they tested the hypothesis that the optical distortion of the background is used to infer refractive material properties of thick, clear transparent objects. This hypothesis seems to be reasonable, because the refractive index is the only material-related property that influences the distortions of the background texture in the proximal image. To isolate refraction-related properties, they used simulated images of massive, relatively thick objects with curved surfaces that had an internal absorption of zero (see Figure 1c for an exemplary depiction of such objects). In each trial of their matching experiments, they presented two such objects in different static scenes and the subject's task was to adjust the refractive index of the test object until it appeared to be made of the same material as the standard object.

Such matches are nontrivial because the optical distortion of the background does not only depend on the refractive index of the transparent object but also on numerous context factors (e.g., object thickness, object shape, background distance, Figure 2). If the

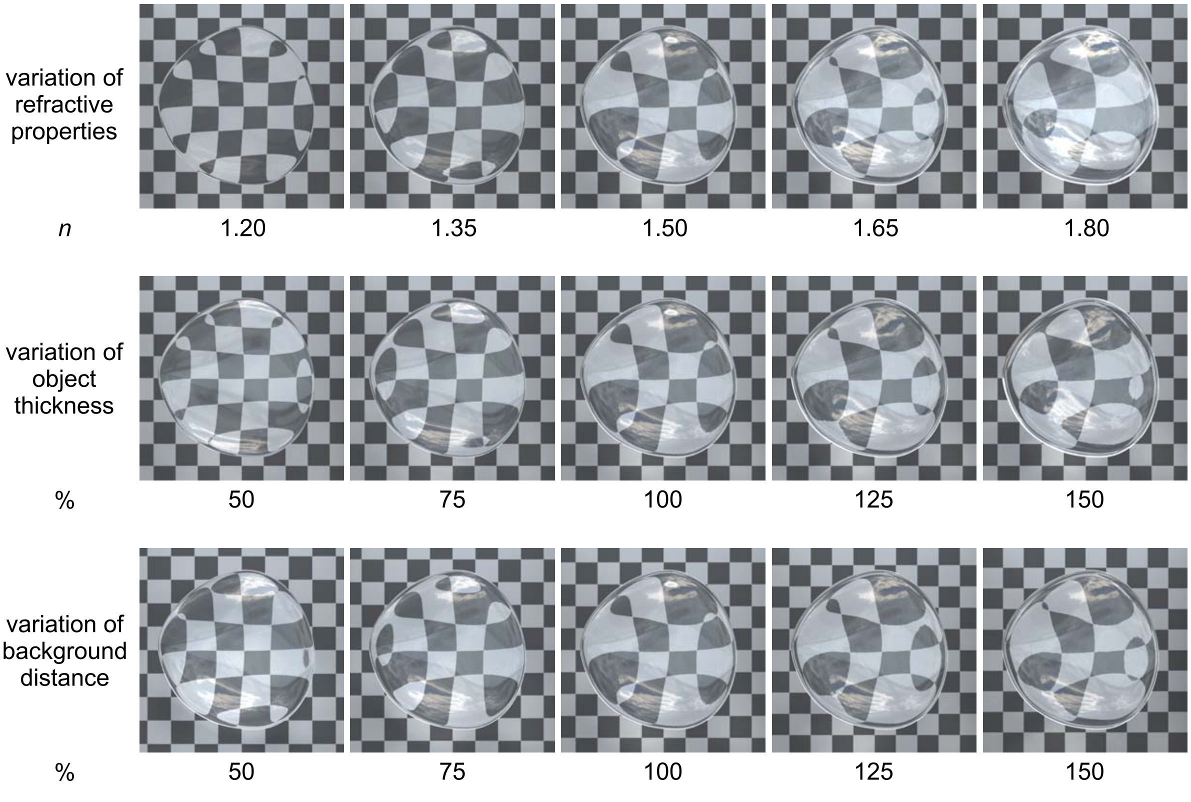

Figure 2. The optical distortion of the background is not only influenced by the refractive index of the object material (first row, object thickness and background distance are 100%, i.e., 6 mm and 60 mm, respectively), but also by context factors like the object thickness (second row, refractive index = 1.5, background distance = 100%) or the background distance (third row, refractive index = 1.5, object thickness = 100%).

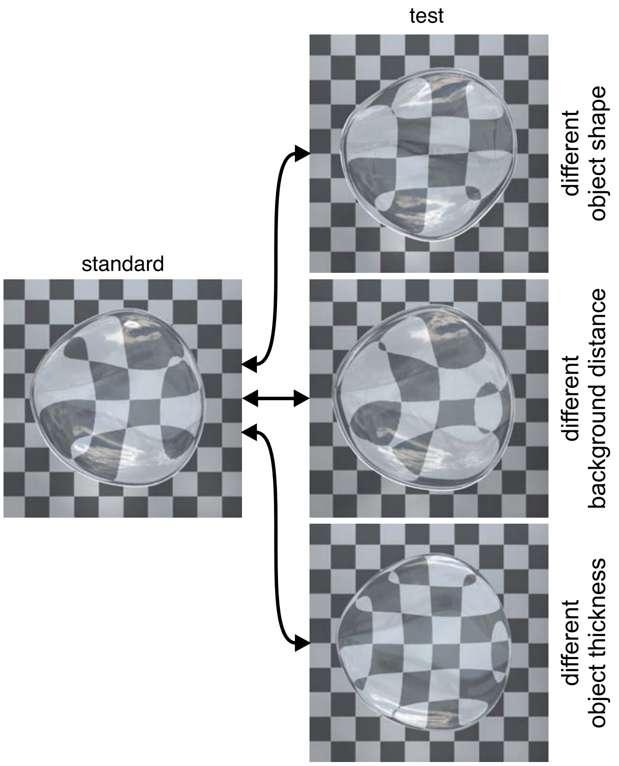

Figure 3. Matching situations in which standard and test object differ in either the object shape, the background distance, or the object thickness. If subjects were able to successfully match refractive material properties of transparent objects, the displayed objects should represent the best match (despite the clearly different background distortions), because they all have the same refractive index. Fleming et al. (2011) implemented only two of these types of context changes, namely changes in background distance and object thickness.

context of the two presented objects differs in at least one of these factors, then it is not possible to match the refractive properties on the basis of a simple proximal match. In their experiments, Fleming et al. (2011) varied the object thickness and the background distance (cf. Figure 3).

Fleming et al. (2011) found that the refractive indices, adjusted by the subjects, correlated highly with those simulated in the standard objects. From this the authors conclude that the optical distortion of the background is used by the visual system to infer refraction-related material properties of transparent objects. Accordingly, it is assumed that the subjects' settings reflect abstract matches between estimated refractive indices that are part of the mental representation of transparent materials.

However, if one takes into account that the optical distortion of such objects depends, in a highly complex way, on specific values of several context factors that are—at least in the situation that was simulated in the experiments—all unknown to the subjects, this interpretation seems not very plausible. In this work, we will deal in detail with the interpretation proposed by Fleming et al. (2011) and discuss alternative interpretations.

We will start by analyzing (a) which factors influence how the background texture is distorted, (b) which factors influence the detectability of an existing distortion, and (c) which context information and processing steps would be needed to use distortion as a cue for refraction-related material properties.

Our analysis starts from the well-known fact that optical distortions caused by thick, transparent objects are the result of a complex interplay of refraction-related material properties, object shape, parameters of the surrounding scene, and the viewing conditions. It seems nearly impossible to disentangle this interplay in order to determine the contribution of a single factor, for instance that of the refractive index, without exactly knowing the values of all other factors. This ambiguity can explicitly be demonstrated for the type of static situations that was used by Fleming et al. (2011). We show that in these scenes objects of different thicknesses and refractive properties can cause indistinguishable background distortions.

Refraction does not only influence background distortion, but also the amount of specular reflection at the surface of the object, a fact that Fleming et al. (2011) mentioned, but did not focus on. A brief analysis of specular reflection will reveal that it is, just like background distortion, not only influenced by the refractive index but also by several context factors. Thus, also in this case, it is unclear how it could be used as a cue for material-related properties.

We will also discuss an alternative interpretation of the findings of Fleming et al. (2011). Our central assumption is that the subjects did not compare internal estimates of material parameters but instead performed simple matches based directly on the similarity of image attributes. Such similarities on the image level can refer to properties of the background distortion and/or characteristics of the specular reflection component. In the following, we will call these image-based matches "similarity matches" to distinguish them from abstract matches on a representational level as they are assumed by Fleming et al. (2011).

We will present results from matching experiments similar to the ones used by Fleming et al. (2011), in which we determined how these two kinds of image attributes influence the subject's matching behavior. The results suggest that the subject's settings reflect a compromise between a similarity match based solely on background distortions and a similarity match based solely on specular reflections. The relative weight with which the two image attributes influence the matching behavior appears to depend on their particular salience.

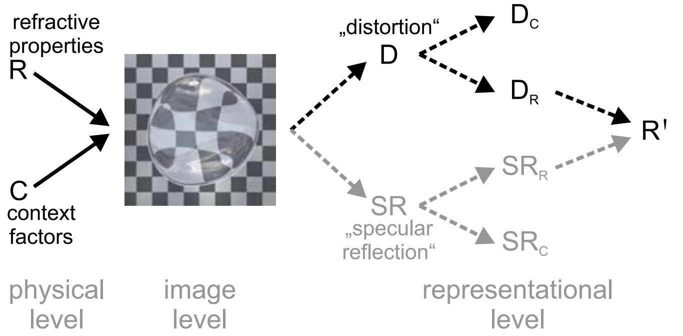

Figure 4. Illustration of the steps required for estimating refractive material properties of transparent objects. The black lines explicate the hypothesis proposed by Fleming et al. (2011) with respect to background distortions; the gray lines extend the principle to specular reflections. The starting point is a physical scene containing a transparent object. Besides the material-related refractive properties (R) of the object, it is determined by various context factors (C; e.g., the object shape and thickness, the illumination, the background distance, etc.). The visual system would first have to extract the components of the background distortion (D) and the specular reflection (SR) from the proximal stimulus. Afterwards, a separation of the two image attributes in refractive- and thus material-related components (D_R and SR_R) as well as context-related components (D_C and SR_C) would be required. The material-related components of both image attributes would then have to be used to generate a representation of refraction-related material properties (R′).

Background distortion and specular reflection as potential cues for estimating the refractive index

If one assumes that the subjects' settings in the experiment of Fleming et al. (2011) reflect a comparison between internal estimates of the refractive indices of the standard and test object, then this implies that relatively precise estimates are possible on the basis of the information provided in the stimulus. With respect to background distortions, this would require the visual system to first determine the background distortion from the proximal stimulus and then to isolate the part that is refraction-dependent (i.e., material-related) from the total distortion (Figure 4).

There are different ways of how the total distortion can be determined. One obvious possibility is to compare the distorted region of the proximal stimulus with an undistorted initial state. Since the undistorted state of the region behind the transparent object is naturally not directly available, it has to be actively generated and represented in an appropriate way. To this end, the visual system might relate to an undistorted part of the background, provided there is one in the proximal stimulus. In dynamic scenes (e.g., with changing object positions and orientations), in which the background is temporally visible in plain view, the undistorted state could be represented in memory. But even if undistorted regions are not available in the stimulus, the visual system might resort to a set of internalized assumptions about general regularities of the environment or simplicity rules as they are suggested by Gestalt theories to infer the undistorted reference texture. If, for instance, the distortion patterns in Figure 3 would be presented without the undistorted surround, then regularity considerations might nevertheless suggest that the undistorted texture is a checkerboard. Furthermore, if a fixed background is seen under slightly varying conditions, then the estimate of the undistorted background may improve over time due to an accumulation of partial information.

The extent to which the visual system can determine the refraction properties of the filter material from the distortion in the proximal stimulus is, in general, limited by the regularity and density of the background texture. It is more difficult to decide with irregular backgrounds than with regular ones, whether a proximal pattern originates from optical distortions elicited by a transparent object, or whether it is already contained in the undistorted background texture. Background texture density poses another limit, because an optical distortion can only be detected to the extent that it leads to noticeable texture changes in the proximal stimulus. Here, the worst case is a homogeneous background, where the proximal image remains completely unchanged under arbitrary optical distortions.

To derive refraction-related material properties from the distortion information generated in the second step, it would be necessary to isolate the refraction-dependent (i.e., material-related) part of the background distortion from the part caused by context factors. Context factors are, amongst others, the shape of the object and its position and orientation relative to the background and the observer. To be able to decompose the background distortion into a refraction-related and a context-related part, it would not only be necessary to possess exact information about the actual values of the parameters that characterize each context factor but also knowledge about the laws that describe their joint influence on the background distortion. It is not very plausible that all these requirements are fulfilled, especially for static stimuli such as those used by Fleming et al. (2011).

Compensation of thickness and refractive index

The ambiguity of background distortions with respect to the refractive properties of the transparent objects that caused them becomes even more apparent if one considers the influence of object thickness. An

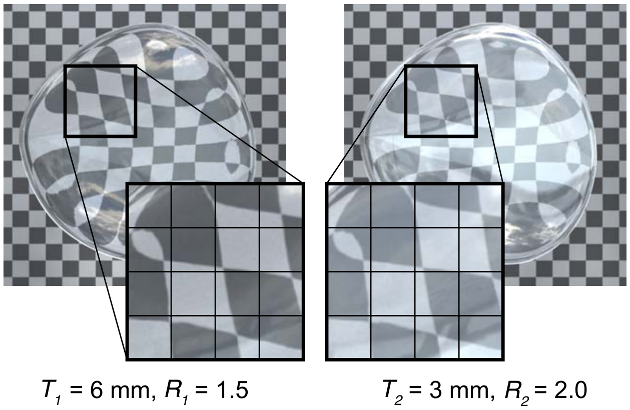

Figure 5. Two objects that differ in thickness and refractive index but lead to virtually identical background distortions. The left object has a maximum thickness T_1 of 6 mm and a refractive index R_1 of 1.5 mm. The right object is only half as thick ($T_2 = 3$ mm). If its refractive index R_2 is chosen to be 2.0, then the background distortion caused by both objects is virtually identical. Obviously, this does not hold for the degrees of specular reflection at the objects surface, which is, due to the larger refractive index, much higher on the right side.

interesting observation made in our simulations of transparent objects was that the optical distortion can be perceptually indistinguishable for objects that differ in thickness and refractive index. Figure 5 shows an exemplary pair of objects that produce virtually indistinguishable background distortions (an effect abbreviated below as "R-T compensation"). To examine the regularities of this compensation in a more systematic manner, we analyzed the simple case of objects with parallel surfaces (Figure 6a). To simplify things further, we first considered the optical displacement of a single ray. Figure 6b shows for a specific angle of incidence ($\theta = 5°$) and a reference object with refractive index R_1 and thickness T_1, how the refractive index R_2 of a second object with thickness ratio T_2/T_1, must be chosen to achieve the same optical displacement of the ray in both objects (see Appendix for a detailed derivation). The optical displacement D does not only depend on the object thickness T and the refractive index R, but also on the angle of incidence θ, with which a ray impinges upon the surface of the object. Thus, depending on the angle of incidence, different combinations of refractive indices and object thicknesses cause the same optical displacements of the corresponding ray. In general, if an observer looks at a transparent object, the angle of incidence may be different for each point on its surface (Figure 6c). The higher the viewing distance is and the smaller the object is, the smaller the differences in the angle of incidence are. For the complete optical distortion, which can be regarded as a combination of optical displacements of single rays, no simple solution exists, that is, a solution with a compensating refractive index that is constant across space. Figure 6b shows, however, that the influence of the angle of incidence θ onto the refractive index can be neglected for a large range of T_2/T_1.

Different combinations of refractive indices and object thicknesses that cause virtually indistinguishable background distortions can also be found for more complex scenes, like the one shown in Figure 1c. We used numerical methods to analyze the situation illustrated in Figure 5: Figure 8a shows (averaged across all angles of incidence) which refractive index R_2 of an object with thickness T_2 causes an optical distortion that is virtually indistinguishable to the one caused by a second object with refractive index R_1 and thickness T_1. The results are again plotted against the thickness ratio T_2/T_1. A comparison with Figure 6b reveals that the graphs are similar to those obtained for the simplified situation. Figure 8b shows that the residual differences between the distortions that were considered as being indistinguishable are indeed negligible for the range of values used in Figure 8. In summary, the analysis reveals that different combinations of thicknesses and refractive indices may cause virtually indistinguishable background distortions. In particular, this is the case for the parameter range used in Fleming et al.'s (2011) experimental paradigm ($2/3 < T_2/T_1 < 2$ and $1.1 < R_1 < 2.0$).

Specular reflection as a cue

Fleming et al. (2011) did not consider in detail that the refractive index of transparent objects also influences how much of the incident light is specularly reflected from their surfaces. The higher the refractive index of a transparent object is, the stronger these specular reflections are (Figure 2). Additionally, the higher the angle of incidence is, the more light is reflected (Fresnel effect).

In general, using specular reflections as a cue for the refractive properties of transparent objects leads to similar problems as the use of background distortions, because specular reflections also vary with several parameters of the scene that are not related to material properties. They depend, for instance, on the illumination or the viewpoint. Thus, like background distortions, specular reflections are also not uniquely related to the refractive index. To estimate refraction-related material properties from specular reflections, the visual system would first need to generate information about the specular reflection from the proximal stimulus. Afterwards, this information would have to be decomposed in a refraction-dependent (i.e., material-related) and a context-dependent part. Like in the case of background distortions, this decomposition would require exact knowledge about all context

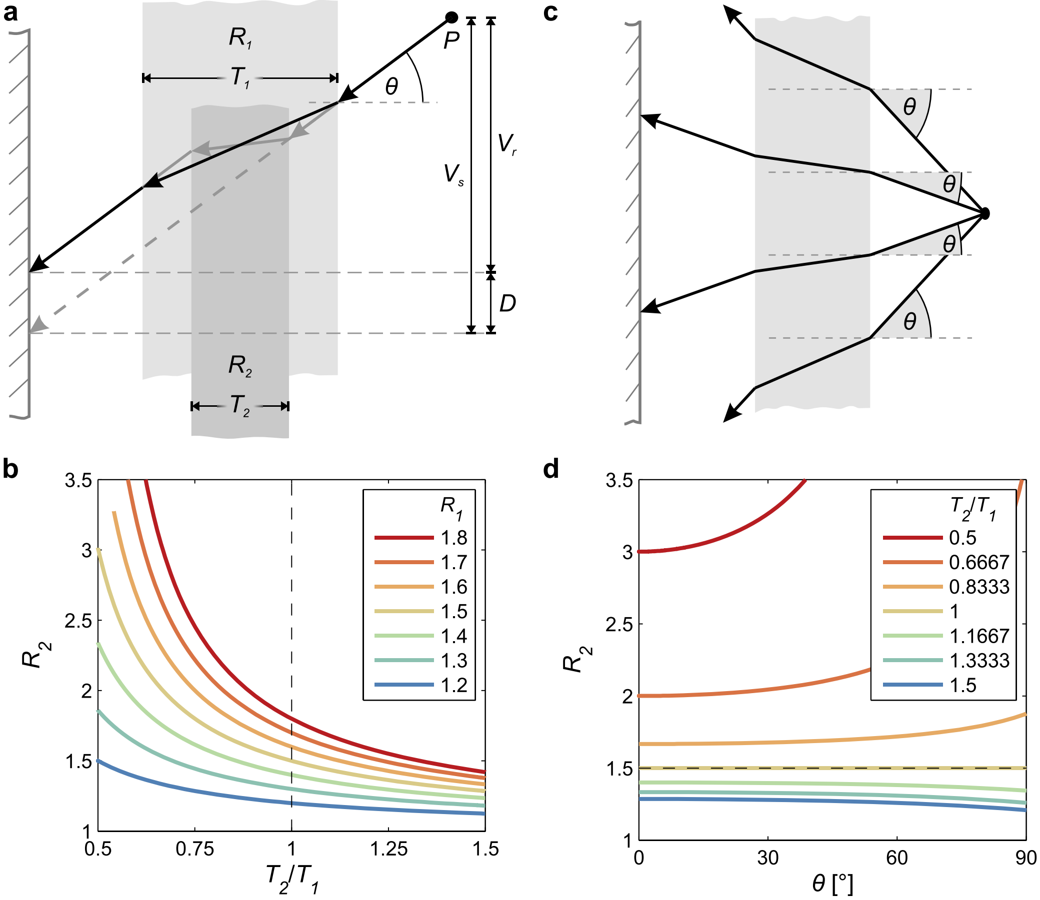

Figure 6. The influences of object thickness and refractive index on background distortions partly compensate each other. (a) Situation in which two objects with parallel surfaces (dark gray and light gray area) with different thickness ($T_1 > T_2$) and different refractive index ($R_1 < R_2$) cause the same optical displacement D, because here the larger refractive index R_2 compensates completely for the effect of the smaller thickness T_2. (b) For the situation illustrated in (a) and an angle of incidence of $\theta = 5^\circ$, the graphs show how the refractive index R_2 of an object with a thickness ratio of T_2/T_1 has to be chosen to create the same optical displacement as a second object with refractive index R_1 and thickness T_1. (c) The compensation of the effects of object thickness and refractive index depends on the angle of incidence θ. (d) The graphs show how the refractive index R_2 of an object with a thickness ratio of T_2/T_1 has to be chosen to create the same optical displacement as a second object with refractive index R_1 (shown here for $R_1 = 1.5$) and thickness T_1 depending on the angle of incidence. For many cases, these graphs are relatively flat, so that a constant value for R_2 represents a good approximation. In these cases, an approximate compensation of object thickness and refractive properties (R-T compensation) can be assumed. However, for large θ and thickness ratios less than 1, there are cases where either no solution for R_2 exists or where the values of R_2 would be higher than those typically found in nature.

factors and their specific influence on the specular reflection (Figure 4).

Discussion

The previous analysis revealed that the background distortion elicited by a transparent object is not uniquely related to refraction-related material properties, but instead depends on many additional context factors. Moreover, objects that differ in refractive index and thickness can cause virtually indistinguishable distortions. In order to use the background distortion to estimate refraction-related material properties, distortion information has to be extracted from the proximal stimulus. The detectability of the distortion strongly depends on context factors like the characteristics of the background itself. Furthermore, the total distortion information would have to be decomposed into a refraction-dependent (i.e., material-related) and a context-dependent part. This would require extensive knowledge about the actual values of all relevant context factors and how they influence background distortion. Putting all the facts together, it thus seems highly questionable whether the background distortions can be used, at all, to infer refraction-related material properties of transparent

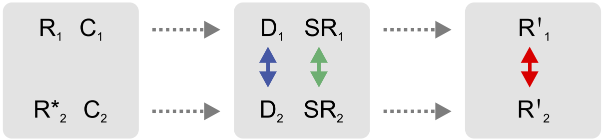

Figure 7. Schematic depiction of different possible matching behaviors in situations like the one used by Fleming et al. (2011), in which the material of a given object with fixed refractive index R_1 presented in context C_1 must be matched in a second object presented in context C_2, by adjusting its refractive index R^*_2. Fleming et al. (2011) proposed that the settings of R^*_2 were performed by comparing internal representations of refractive indices that were estimated from the proximal stimulus (R′ match, red arrow). Alternatively, subjects might have performed simple similarity matches based directly on image attributes related to either background distortions (D match, blue arrow) or specular reflections (SR match, green arrow) or a combination of these.

objects. In particular, this speaks against the interpretation that subjects compare internal representations of estimated refraction indices in matching tasks like the one of Fleming et al. (2011).

An alternative explanation is that subjects perform simple matches based directly on image attributes available in the proximal stimulus (Figure 7). This would be a reasonable strategy in the given experimental situation, because both objects are visible at the same time and can be compared directly. This explanation seems to be compatible with the results of Fleming et al. (2011): Even though the refractive indices chosen by the subjects for the test objects correlate with those of the standard objects, they nevertheless show systematic deviations. The specific form of these deviations suggests a matching behavior that was based on similarity matches of background distortions (abbreviated below as "D match"). Moreover, in their exemplary illustration of the mean settings made by their subjects, a highly similar background distortion in standard and test object can be observed (Fleming et al., 2011, figure 3d). Because background distortions are not the only image attribute available in the stimulus, such similarity matches could just as well refer to specular reflections (abbreviated below as "SR match"), or the subjects' settings might represent a compromise between both kinds of similarity matches.

It is presently unclear on which image criteria such similarity matches might be based. In principle, a large number of criteria is possible. For example, summary measures like the maximal brightness of the specular reflection or the average element size of the distorted background might be used. Alternatively, more complex scene statistics like the brightness histogram of specular reflections or histograms of texture element sizes could be used. Such similarity matches can be performed even if exact proximal identity matches are not possible.

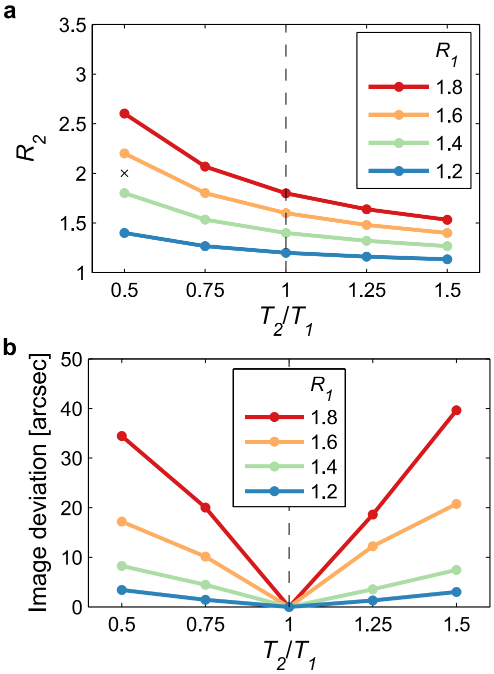

Figure 8. Compensation of refractive index and object thickness (R-T compensation) for the situation illustrated in Figure 5. (a) The graphs show how the refractive index R_2 of an object with thickness T_2 has to be chosen to achieve a background distortion that is virtually indistinguishable from the one caused by a second object with refractive index R_1 and thickness T_1. The values are plotted against the thickness ratio T_2/T_1 and represent averages across all angles of incidence. Because no simple analytic solutions exist for these situations, ray paths were simulated explicitly. The black cross represents the case shown in Figure 5. (b) The graphs show the mean difference between corresponding points in the distorted images obtained with the two objects, if the refractive index R_2 is chosen as illustrated in (a). The largest average deviation is less than 40 s of viewing angle. Thus, the corresponding background distortions are virtually indistinguishable and the influence of the angle of incidence θ is negligible. In this example, 0.01° viewing angle corresponds to only 0.14% of the object diameter in the image or to 0.26 px in an image of size 370 × 370 px.

It is important to point out that the results of similarity matches based on specular reflections (SR match) can easily be misjudged as resulting from successful abstract matches of estimated refraction indices ("R′ match"). Because the specular reflections are hardly influenced by the context factors varied in the experimental paradigm used by Fleming et al. (2011), SR matches would accidentally also lead to good matches of the refractive index. This means that any result that is consistent with an R′ match, can be explained more economically by a simple SR match.

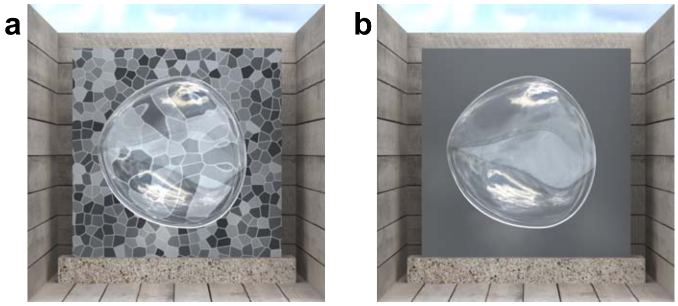

Figure 9. Illustration of the stimuli used in the first experiment. (a) In half of the trials we used a random Voronoi pattern as background texture that resembled the background textures used by Fleming et al. (2011). (b) To gain predictions for similarity matches based solely on the specular reflection (SR match), we used a homogeneous background texture in the other half of the trials.

A further interesting point is that similarity matches, like the ones outlined above, are still possible even if the stimuli do not elicit an impression of transparency. This clearly indicates that an explanation of the subjects' settings in terms of similarity matches is of a fundamentally different nature than the one proposed by Fleming et al. (2011) themselves, because the latter relates to estimates of refraction indices and therefore only appears reasonable if a transparency-specific mechanism of the visual system is assumed. The question on which basis the subjects actually made their settings is therefore of great theoretical importance.

Experiment 1

In order to investigate which strategy subjects use if they are asked to match the perceived material properties of two transparent objects, an experimental situation is required in which the use of different matching strategies leads to distinguishable results. The experimental paradigm of Fleming et al. (2011), in which the material of transparent objects of different thickness must be matched, seems well-suited for this purpose.

If subjects were able to successfully perform an abstract match of refraction-related material properties (R' match), then the refractive indices set in the adjustable test objects should be identical to those used in the corresponding standard objects. Contrary to this expectation, systematic deviations from the given refractive indices were observed in the results of Fleming et al. (2011). In principle, these deviations could at least in part be due to contradictory information in their stimuli. To reduce the likelihood of such negative influences in our study, we tried to further optimize the realism of the stimuli. Due to these changes, our stimuli differ slightly from those used in Fleming et al. (2011). All perspective parameters of our stimuli were chosen in correspondence with the actual settings and physical apparatus used in the experiment. Thus, with respect to geometry, the stimuli appeared virtually the same as a real instance of that scene would have appeared. Furthermore, we used complex natural environmental lighting instead of a few simple localized light sources, because this reduces high-contrast highlights that are difficult to handle in tonemapping. As a result, the salience of the specular reflections may be slightly higher in our stimuli.

If subjects would perform similarity matches based on background distortions (D match), then the adjusted refractive indices would differ from the ones given (as already mentioned, simple D matches lead in nontrivial matching situations as those realized by Fleming et al. (2011) to nonveridical refractive indices). The approximate R-T-compensation discussed above can be used to predict for standard and test objects of different thickness the settings of the refractive index that should result, if the subjects actually perform similarity matches of background distortions: It would be the refractive index that leads to virtually indistinguishable background distortions in standard and test, because this should always represent the best match, no matter which (potentially abstract) criteria are actually used in a similarity match of background distortions.

Predictions for similarity matches based on the specular reflection (SR match) were gained empirically by asking subjects to match objects with isolated specular reflections. For this purpose, we used homogeneous background textures that cannot map any background distortions. Thus, the specular reflection remains the only base for the matches.

Stimuli

The stimuli were computer-generated images created with the Mitsuba renderer (Jakob, 2013). The stimuli were created to closely resemble the ones used by Fleming et al. (2011) to ensure comparability. The stimulus images showed a thick transparent object in front of a background board that was located inside a box with front and top sides open (Figure 9). All scene elements were defined in real-world coordinates relative to a virtual projection plane, which represented the surface of the experimental screen. The shape of the transparent object was based on an icosahedron that was subdivided seven times. The resulting icosphere was deformed to a slightly warped ellipsoid with the computer graphics software Blender (Blender Foundation, 2013) by applying various shape modifiers to its mesh. The object size was about 50 × 50 mm (width, height). Depending on the experimental condition, the thickness of the standard object varied in four steps (T_S

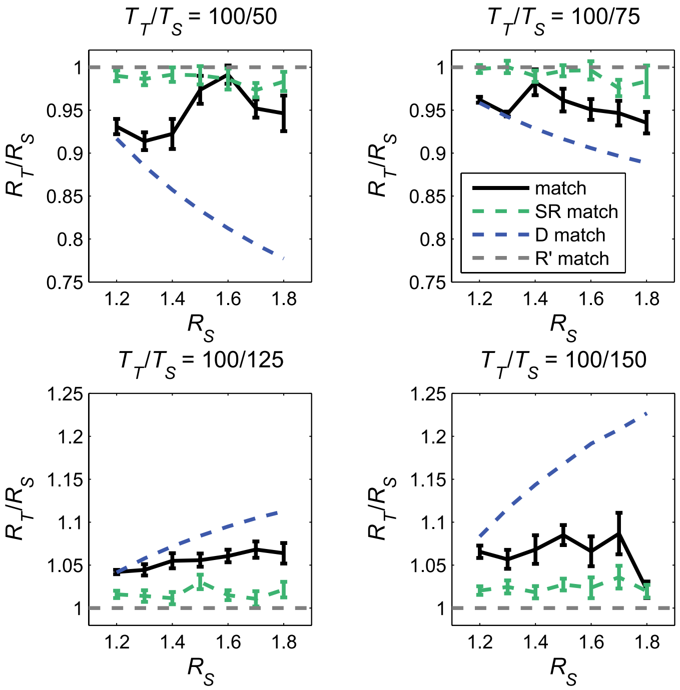

Figure 10. Results of the first experiment. Each plot corresponds to one object thickness ratio T_T/T_S. The black solid lines represent the mean setting of the refractive indices of the test object. The dashed blue lines represent the predictions for similarity matches based solely on the background distortions (D match), which were gained by simulations. The dashed green lines represent the predictions for similarity matches based solely on specular reflections (SR match), which were gained by matches with isolated specular reflections. The dashed gray lines represent the settings expected for abstract matches of estimated refractive indices (R′ match). All values are depicted as ratios of the refractive index of the test object to the refractive index of the standard object (R_T/R_S, ± SEM).

$\in$ {3, 4.5, 7.5, 9 mm}), while the thickness of the test object remained constant ($T_T = 6$ mm). The thickness was manipulated by applying scaling factors to the object mesh. The object was made of light transmitting material ("dielectric") without any internal absorption. The refractive indices of the standard stimuli varied according to the experimental condition in seven steps ($R_S \in \{1.20, 1.30, 1.40, 1.50, 1.60, 1.70, 1.80\}$) The refractive index of the test stimuli were adjusted by the subjects ($R_T \in \{1.010, 1.015, \ldots, 2.495, 2.500\}$, 299 steps). The refractive index of the medium surrounding the transparent object was set to 1. The transparent object was located at the center of the virtual projection plane. A textured 80 × 80 mm background board was placed behind the transparent object at a distance of 60 mm. The background textures were random Voronoi patterns created with Matlab (Mathworks, Inc., Natick, MA) that resembled the background textures used by Fleming et al. (2011). The individual faces of the pattern used in the rendering were separated by seams with a width of 0.32 mm and a color of R, G, B = 180. The colors of the faces were uniformly distributed between R, G, B = 75 and R, G, B = 175. Additionally, we used a homogeneous background texture (R, G, B = 125) to isolate the specular reflection component. The sole illumination was provided by an infinitely distant high dynamic range sphere emitter, containing a natural daylight outdoor scene with a partly cloudy sky. The camera settings (location and field of view) were chosen to correspond to the actual experimental setup. Thus, the stimuli appeared in exactly the same way as a corresponding real scene, and there were virtually no perspective distortions. Stimuli were rendered as 16-bit high dynamic range images (extended volumetric path tracer; low discrepancy sampler with 512 samples/px; Gaussian reconstruction filter with $SD = 0.5$) and subsequently tonemapped to 8-bit low dynamic range images (gamma = 1.6; exposure = 1.4). The final image size was 370 × 370 px which corresponded to 100 × 100 mm on the screen.

Procedure

In each trial, two stimuli were presented simultaneously on an Eizo ColorEdge CG243W LCD screen (display area 518.4 × 324.0 mm; resolution 1920 × 1200; color depth 8-bit per channel; 3.704 px/mm; Eizo Corporation, Hakusan, Japan) in a darkened room. The fixed standard stimulus was located at the top of the screen, the adjustable test stimulus at the bottom. The viewing distance was 400 mm. The subject's task was to adjust the refractive index of the test object until it appeared to be made of the same material as the fixed standard object. The settings were made with the arrow keys on a standard computer keyboard.

Each subject performed 168 trials in randomized order. In 84 of them, the Voronoi texture was used (7 R_S × 4 T_S × 3) and in the remaining 84 the homogenous texture (7 R_S × 4 T_S × 3).

Subjects

Six subjects, three of them female, participated in the experiment. Their ages ranged from 20 to 35. All subjects were naive as to the purpose of the experiment. They reported normal or corrected-to-normal visual acuity, and showed no color vision deficiency, as tested by Ishihara plates (Ishihara, 1969).

Results

Figure 10 shows the results of the first experiment. For all four object-thickness ratios (T_T/T_S) and all

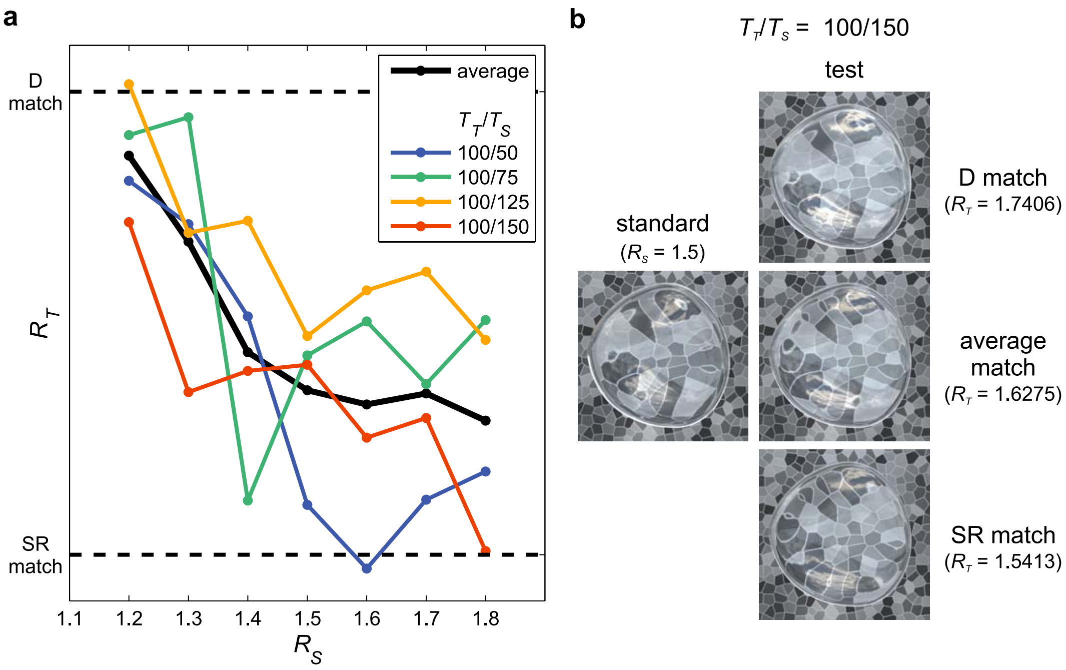

Figure 11. The settings for the refractive indices lie between the predictions for similarity matches of background distortions (D match) and similarity matches of specular reflections (*SR* match). (a) Relative position of the refractive indices of the test object (R_T) in the interval between the respective predictions of D matches and SR matches for all four thickness ratios (T_T/T_S) as a function of the refractive index of the standard object (R_S). The black line represents the average setting across all four thickness ratios. (b) Example of a fixed standard object (left column, $R_S = 1.5$) and a test object (right column, $T_T/T_S = 100/150$) with different settings for the refractive index R_T. The topmost object shows a similarity match based solely on background distortions (D match, gained by simulations using the R-T compensation); the lowermost object shows a similarity match of specular reflections (SR match, gained empirically). The refractive index of the center object corresponds to the subjects' mean setting. The images show only a part of the complete stimulus illustrated in Figure 9.

refractive indices of the fixed standard (R_S), the mean setting of the refractive indices of the test object lies between the predictions for a similarity match of background distortions (D match) and a similarity match of specular reflections (SR match). The settings of the refractive indices R_T averaged across all object thickness ratios T_T/T_S lie closer to the predictions for similarity matches of specular reflections (SR match) if the refractive indices R_S of the fixed standard is high (see black curve in Figure 11a) and tends towards the predictions for similarity matches of background distortions (D match) for low R_S. For an object thickness ratio of $T_T/T_S = 100/150$ and a standard object with a refractive index of $R_S = 1.5$, Figure 11b illustrates exemplarily how the test object would have looked like for a D match, a SR match, and for the subjects' mean setting for the refractive index.

Our findings closely resemble the corresponding results reported in Fleming et al. (2011), although this is somewhat hidden by the different way in which our results are presented: The refractive indices adjusted by the subjects correlate highly with the refractive indices of the standard objects, but show systematic deviations. This suggests that both experiments probe the same processes, despite the slight differences between the stimuli used in both experiments.

The settings of two subjects differ substantially from the results presented so far and are therefore considered separately. Regardless of the object thickness ratio (T_T/T_S) or the refractive index of the standard object (R_S), their settings always resemble the ones expected for a similarity match based solely on background distortions (D match; Figure 12).

Discussion

In our experiment, we used matching situations in which different matching strategies would lead to clearly discriminable settings. The range of the adjustable parameter of the test stimulus always enclosed the value predicted for similarity matches based solely on the background distortions (D match) and the value predicted for similarity matches based solely on the specular reflections (SR match). By comparing the actual settings with these two predictions, it is possible to gain information about the relative influence of both image criteria on the matching behavior. Our results

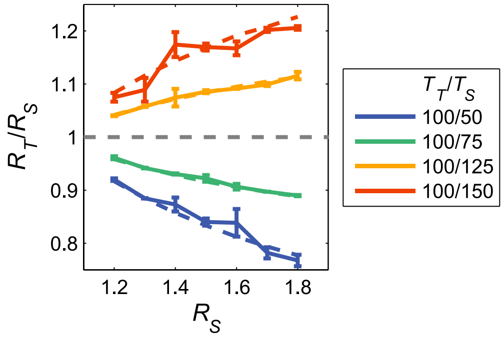

Figure 12. The settings of two of the subjects deviate systematically from those of the other subjects. Their adjusted refractive indices is very similar to the predictions made for similarity matches of background distortions (D match). The colored solid lines represent the ratios of the refractive indices of the adjustable test to the fixed standard (R_T/R_S, ± *SEM*) for a particular object thickness ratio (T_T/T_S). The dashed lines represent the D match prediction for each corresponding object thickness ratio.

suggest that the matches reflect a compromise of similarity matches based on background distortions and similarity matches based on specular reflections. The relative influence of the two image attributes seems to depend on their salience. With lower refractive indices of the standard object and thus weak specular reflections, the subjects tended to base their setting mainly on background distortions. With higher refractive indices of the standard object and correspondingly stronger specular reflections, they tend to rely more on specular reflections.

Apparently, two of our subjects always performed similarity matches of background distortions and completely ignored specular reflections. The fact that different subjects used different matching strategies suggests that subjects deliberately choose a matching criterion. This suggests that also the compromise between similarity matches of background distortions and similarity matches of specular reflections that was found for the other subjects might mainly be the result of a conscious decision and not the result of unconscious cue integration processes in visual perception.

Experiment 2

Our alternative interpretation of the results of Fleming et al. (2011) is that their subjects did not compare internal estimates of refractive indices, but instead performed simple similarity matches of image attributes in the proximal stimulus. It is presently unclear to which criteria such similarity matches actually refer. It might be possible that summary measures, such as the average size of texture elements or the maximum brightness of the surface reflection, play an important role. In any case, it seems clear that such similarity matches must rely on abstract similarity criteria and not on local proximal identity. To test the hypothesis that such similarity matches are not limited to matching situations in which exact proximal identity of an image attribute could be achieved (as was the case for the background distortion and, to a certain extent, also for the specular reflection in the first experiment), we conducted a matching experiment in which the object shape differed between the standard and test object. Thus, in contrast to the first experiment, neither the background distortion nor the specular reflection could be made proximal identical. Both the background distortions and the specular reflections were matched in isolation from each other.

If the results of this experiment are comparable to the ones of the first experiment, then the hypothesis that subjects perform abstract similarity matches based on the background distortion and the specular reflection is supported.

Stimuli

The stimuli were largely the same as in the first experiment. However, different object shapes were used for the standard stimuli by applying different shape modifiers. As in the first experiment, the refractive index of the standard object was varied in seven steps ($R_S \in \{1.20, 1.30, 1.40, 1.50, 1.60, 1.70, 1.80\}$). To isolate information given by the specular reflection component, we used the same homogenous background texture as in the first experiment (R, G, B = 125). In this condition, standard and test objects did not differ in any parameter, except in shape. To isolate information given by background distortions, we rendered stimuli without the specular reflection component by choosing a corresponding option in the Mitsuba renderer (Jakob, 2013). Because higher refractive indices lead to a higher degree of specular reflection, the deactivation of the specular reflection component darkened the image area of the object. To compensate for this, we adjusted the brightness of the object area accordingly. In this condition, the thickness of the standard object was varied in four steps ($T_S \in \{3, 4.5, 7.5, 9 \text{ mm}\}$). The object thickness was the same for all test stimuli ($T_T = 6$ mm), while their refractive properties varied according to the inputs made by the subject ($R_T \in \{1.010, 1.015, \ldots, 2.495, 2.500\}$, 299 steps).

Procedure

The procedure was the same as in the first experiment. Every subject was presented with 105 trials

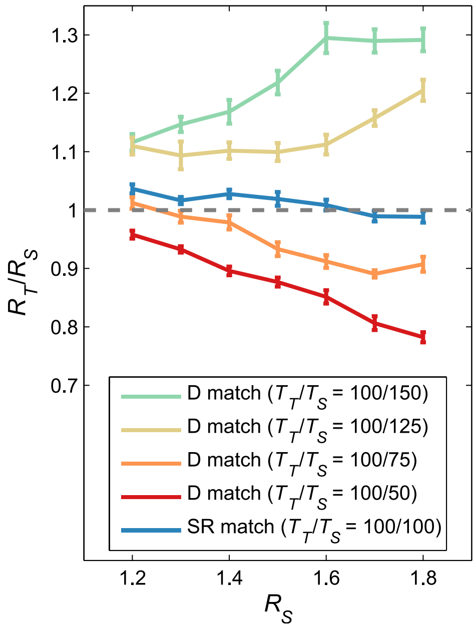

Figure 13. Results of the second experiment. Ratios of the mean refractive index settings to the refractive index of the corresponding standard object (R_T/R_S, ± *SEM*) for matches with isolated distortion (D match) under four object thickness ratios (T_T/T_S), and for matches with isolated specular reflection (SR match).

in randomized order. Twenty-one of the trials had an isolated specular reflection component (7 R_S × 3), and 84 had an isolated background distortion (7 R_S × 4 T_S × 3).

Subjects

The subjects were the same as in the first experiment.

Results

Figure 13 shows the mean settings of the refractive indices of the test object for the matches with isolated background distortions (D match) and for matches with isolated specular reflections (SR match). For situations with isolated background distortions, the settings are similar to the predictions for similarity matches of background distortions derived in the first experiment. Hence, these settings are also similar to those of the two subjects of the first experiment, who based their matches solely on background distortions.

The matches in situations with isolated specular reflections are similar to the predictions for similarity matches of specular reflection, which were gained in the first experiment in a similar way (but with the same object shapes in the standard and test stimulus).

Discussion

The data show that subjects are able to perform reliable similarity matches based on background distortion and specular reflection, even if the objects in the standard and match stimulus differ in shape so that neither the background distortion nor the specular reflection could be matched directly (i.e., by proximal identity). For both kinds of similarity matches, results similar to the ones of the first experiment were found. This suggests that subjects refer to abstract similarity measures, such as, for example, the average size of texture elements in the distorted background area or the maximal brightness of the specular reflection, and not to proximal identity.

General discussion

Fleming et al. (2011) tested the hypothesis that background distortions are used to perceive refraction-related material properties of thick transparent objects. In their experiments, they asked subjects to match the refractive indices of two transparent objects. From the fact that the refractive indices set by the subjects correlated highly with the given ones, they conclude that background distortions are indeed used to perceive material properties of transparent objects. They interpret the subjects' settings as the result of abstract matches of refractive properties—that is, they assume that the subjects compared internal representations of refractive indices that were estimated from information given in the stimulus.

Background distortions and specular reflections are two image attributes that depend on the refractive index, a material property of light transmitting objects, and may thus be regarded as potential cues that can be used to infer refraction-related properties from the proximal stimulus. However, an analysis of these image attributes reveals that they do not only depend on the refractive index but that they are also influenced by several unknown context factors. Thus, these image attributes can only serve as cues for the refractive index, if it would be possible to decompose them into a refraction-dependent (i.e., material-related) part and a context-dependent part. We argue that this assumption is not very plausible, especially if static stimuli are used as in Fleming et al. (2011).

We tested in two experiments the alternative explanation that subjects do not perform abstract matches of estimated refractive properties, but simple similarity matches of image attributes. In line with this hypothesis, our results indicate that the matches reflect a compromise between a similarity match of background distortions (D match) and a similarity match of specular reflections (SR match). The weight that is given to each image attribute in the match seems to partly depend on its salience. Our results indicate that specular reflections tend to be the dominant factor as soon as they become clearly noticeable. Background distortion dominated the settings only for low values of the refractive index of the standard object (and thus weak specular reflections). However, these findings should be interpreted with caution, because the additional finding that the relative contribution of the two image attributes differed considerably between subjects suggests that the compromise is mainly of a cognitive nature.

It is presently unclear which image-based criteria are used in such similarity matches. At any rate, the results of our second experiment, in which subjects matched perceived materials across objects of different shape, indicate that subjects use abstract similarity criteria in such matches and do not refer to local identity. This was found both for isolated background distortions and isolated specular reflections. It remains a goal of future work to find out to which criteria similarity matches of the background distortion or the specular reflection actually refer.

Our theoretical analysis revealed that it is highly implausible that background distortions caused by thick transparent objects can be used to gain an estimate of a refractive index that serves as an internal representation of refraction-related material properties. This judgment is primarily based on computational arguments, but there are several additional arguments that speak against the idea that the visual system tries to gain an estimate of the refractive index as part of an internal representation of light transmitting materials.

A first point is that most of the light-transmitting materials in our environment are either water ($R \approx 1.33$) or (acrylic) glass ($R \approx 1.50$ up to 1.70). Other exemplars are much more rare, like diamonds ($R \approx 2.42$). It is therefore not to be expected that a metrical representation of refractive properties would substantially improve the functional linking of the visual system to the world surrounding us. A second point is that the matching accuracy demonstrated in the experiment of Fleming et al. (2011) and in the experiments reported here is relatively low compared to the small range of relevant refractive indices found in the natural environment. Thus, in most real world cases one could probably reduce the estimation error considerably, if one would replace the estimate from the proximal stimulus with one of the values 1.33 or 1.5. A further point is that a representation of the refractive index would only be useful if its *absolute* value were exactly known. However, such absolute representations are at least very rare in vision; almost all representations are of a relational nature.

Clearly, this does not mean that background distortions and specular reflections of transparent objects are ignored by the visual system. To the contrary, it seems highly probable that this information is used in many different ways in transparency perception. For example, it seems obvious that specular reflections play an important role for shape perception, just as in the case of opaque objects (see e.g., Blake & Bülthoff, 1990, 1991; Fleming, Torralba, & Adelson, 2004; Norman, Todd, & Orban, 2004). Specular reflections of transparent objects are more complex because reflections at different surfaces (e.g., at the front and at the inner back of the object) can be superimposed in the retinal image. At present, it remains unclear if this higher complexity of specular reflections enhances the perception of the object shape (and maybe even allows perceiving the shape of the back surface, which is not directly visible) or if it makes shape perception more difficult.

With respect to background distortions, it appears plausible that they act as an unspecific cue, which plays an important role in identifying objects as being transparent, especially in situations where no other cues for transparency (e.g., color relations) exist. Background distortions can also potentially be used as a cue for the shape of the transparent object. For example, highly distorted areas could indicate significant changes in surface orientation. The use of background distortions as a cue for object shape seems more plausible than its use as a cue for the refractive index, because relative distortion information inside the boundary of the object may suffice in the former case, but not in the latter.

Keywords: perceptual transparency, optical background distortion, specular reflection

Acknowledgments

This research was supported by Deutsche Forschungsgemeinschaft (DFG) Grant FA 425/2-1. We want to thank Laurence T. Maloney for his comments that helped to improve the manuscript.

Commercial relationships: none.

Corresponding author: Franz Faul.
Email: ffaul@psychologie.uni-kiel.de.
Address: Institut für Psychologie, Christian-Albrechts-Universität zu Kiel, Kiel, Germany.

References

Beck, J. (1978). Additive and subtractive color mixture in color transparency. *Perception & Psychophysics*, *23*(3), 265–267, doi:10.3758/BF03204137.

Beck, J., Prazdny, K., & Ivry, R. (1984). The perception of transparency with achromatic colors. *Perception & Psychophysics*, *35*(5), 407–422, doi:10.3758/BF03203917.

Blake, A., & Bülthoff, H. (1990). Does the brain know the physics of specular reflection? *Nature*, *343*(6254), 165–168, doi:10.1038/343165a0.

Blake, A., & Bülthoff, H. (1991). Shape from specularities: Computation and psychophysics. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *331*(1260), 237–252, doi:10.1098/rstb.1991.0012.

Blender Foundation. (2013). *Blender—A 3D modelling and rendering package* (Version 2.66a) [Computer software]. Retrieved from http://www.blender.org

Faul, F., & Ekroll, V. (2002). Psychophysical model of chromatic perceptual transparency based on substractive color mixture. *Journal of the Optical Society of America A: Optics, Image Science, & Vision*, *19*(6), 1084–1095.

Faul, F., & Ekroll, V. (2011). On the filter approach to perceptual transparency. *Journal of Vision*, *11*(7):7, 1–33, http://www.journalofvision.org/content/11/7/7, doi:10.1167/11.7.7.

Fleming, R. W., Jäkel, F., & Maloney, L. T. (2011). Visual perception of thick transparent materials. *Psychological Science*, *22*(6), 812–820, doi:10.1177/0956797611408734.

Fleming, R. W., Torralba, A., & Adelson, E. H. (2004). Specular reflections and the perception of shape. *Journal of Vision*, *4*(9):10, 798–820, http://www.journalofvision.org/content/4/9/10, doi:10.1167/4.9.10.

Ishihara, S. (1969). *Tests for color blindness*. Tokyo: Kanehara Shuppan.

Jakob, W. (2013). Mitsuba rendering system (Version 0.4.4) [Computer software]. Retrieved from http://www.mitsuba-renderer.org

Metelli, F. (1970). An algebraic development of the theory of perceptual transparency. *Ergonomics*, *13*(1), 59–66, doi:10.1080/00140137008931118.

Norman, J. F., Todd, J. T., & Orban, G. A. (2004). Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychological Science*, *15*(8), 565–570, doi:10.1111/j.0956-7976.2004.00720.x.

Singh, M., & Anderson, B. L. (2002). Toward a perceptual theory of transparency. *Psychological Review*, *109*(3), 492–519, doi:10.1037/0033-295X.109.3.492.

Appendix

In addition to the parameters defined in Figure 6, let D_P be the distance of the object's center to the viewpoint P and D_B be the distance of the object's center to the background. To simplify calculations, the refractive index of the medium surrounding the object is assumed to be 1.

If no transparent object exists, a ray originating from the viewpoint P intersects the background at

$$V_s = (D_B + D_P)\tan\theta. \tag{1}$$

If the ray is distorted by the refractive properties of the transparent object, it intersects the background at

$$V_r = \frac{T\sin\theta}{\sqrt{R^2 - (\sin\theta)^2}} + (D_B + D_P - T)\tan\theta. \tag{2}$$

The difference of both intersections gives the optical displacement

$$D = \frac{T\sin\theta}{\sqrt{R^2 - (\sin\theta)^2}} - T\tan\theta. \tag{3}$$

The optical displacement depends on the object thickness T, the refractive index R, and the angle of incidence θ.

For two objects of different thickness (T_1 vs. T_2) and different refractive index (R_1 vs. R_2) but equal distances ($D_P = D_B$), the optical displacement is identical, if the thickness or the refractive index of one of the objects is appropriately chosen. For example, if the thickness of the second object is

$$T_2 = T_1 \frac{\sqrt{R_2^2 - (\sin\theta)^2}\left(-1 + \sec\theta\sqrt{R_1^2 - (\sin\theta)^2}\right)}{\sqrt{R_1^2 - (\sin\theta)^2}\left(-1 + \sec\theta\sqrt{R_2^2 - (\sin\theta)^2}\right)}, \tag{4}$$

then both objects cause the same optical displacement. The refractive index R_2 that leads to identical optical distortions can be described in a similar way by rearranging the formula.