Specular motion and 3D shape estimation

Dicle N. Dövencioğlu Department of Psychology, Bilkent University, Ankara, Turkey; National Magnetic Resonance Research Center, Bilkent University, Ankara, Turkey; Department of Psychology, Justus-Liebig-University Giessen, Giessen, Germany

Ohad Ben-Shahar Computer Science Department, Ben-Gurion University of the Negev, Beer-Sheva, Israel

Pascal Barla INRIA Sud-Ouest, Bordeaux, France

Katja Doerschner Department of Psychology, Bilkent University, Ankara, Turkey; National Magnetic Resonance Research Center, Bilkent University, Ankara, Turkey; Department of Psychology, Justus-Liebig-University Giessen, Giessen, Germany

Dynamic visual information facilitates three-dimensional shape recognition. It is still unclear, however, whether the motion information generated by moving specularities across a surface is congruent to that available from optic flow produced by a matte-textured shape. Whereas the latter is directly linked to the first-order properties of the shape and its motion relative to the observer, the specular flow, the image flow generated by a specular object, is less sensitive to the object's motion and is tightly related to second-order properties of the shape. We therefore hypothesize that the perceived bumpiness (a perceptual attribute related to curvature magnitude) is more stable to changes in the type of motion in specular objects compared with their matte-textured counterparts. Results from two two-interval forced-choice experiments in which observers judged the perceived bumpiness of perturbed spherelike objects support this idea and provide an additional layer of evidence for the capacity of the visual system to exploit image information for shape inference.

Introduction

Motion signals provide important information that facilitates perception and interaction with the environment. Specifically, image motion critically contributes to the recognition of object shape, its material properties, and its three-dimensional (3D) motion characteristics. Humans are often able to infer properties about the shape, material, and motion of an object instantly and effortlessly. However, this is not a trivial accomplishment because all of these properties are—to a certain extent—derived from the very same set of image information. Indeed, simultaneous estimations of shape, material, and motion constitute a classic underconstrained problem. Because all three object attributes contribute to the resulting optic flow in a dynamic scene, it is not surprising that they frequently interact in a given perceptual task.

Interactions of shape, surface material, and object motion in dynamic scenes

Object motion affects surface material estimation

It is now well established that image motion enhances perceived shininess or glossiness (Hartung & Kersten, 2002; Sakano & Ando, 2008; Wendt, Faul, Ekroll, & Mausfeld, 2010). In particular, Doerschner, Kersten, and Schrater (2011) showed that image motion characteristics predict whether an object would appear shiny or matte and how these cues, which are closely related to object curvature, can be used in

Citation: Dövencioğlu, D. N., Ben-Shahar, O., Barla, P., & Doerschner, K. (2017). Specular motion and 3D shape estimation. *Journal of Vision*, *17*(6):3, 1–15, doi:10.1167/17.6.3.

doi: 10.1167/17.6.3 Received December 16, 2016; published June 5, 2017 ISSN 1534-7362 Copyright 2017 The Authors

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

Downloaded from jov.arvojournals.org on 07/24/2020

machine vision applications (Doerschner, Fleming, et al., 2011; Yilmaz & Doerschner, 2014).

Surface reflection affects local shape estimation

Dövencioğlu, Wijntjes, Ben-Shahar, and Doerschner (2015) showed that surface reflectance affects the perceived local curvature sign of a computer-rendered moving object. Specifically, they found the shape estimations of complex *matte* objects to be closer to the "ground truth" (3D model of the object) compared with *specular* ones. Hurlbert, Cumming, and Parker (1991) showed that the velocity of a specular feature alters the perceived local 3D curvature. Using real (rather than rendered) moving objects and a gauge figure task, Vota, Dovencioglu, Ben-Shahar, Doerschner, and Wijntjes (2015) later demonstrated that local tilt and slant estimations for rotating specular objects were rather different from the ground truth 3D shape of the object. Finally, Norman et al. (2016) showed that deforming specular highlights facilitate 3D shape perception.

Surface reflection affects estimation of object motion

Results by Doerschner, Yilmaz, Kucukoglu, and Fleming (2013) indicate that the estimated rotation axis of an object varies with surface material. They found larger estimation errors for specular objects compared with matte-textured ones. Interestingly, this effect was strongly modulated by shape complexity, with simpler specular 3D shapes being more prone to larger errors in rotation axis estimation. This dependency on shape complexity was not found for matte-textured objects.

Taken together, these examples suggest that specular and matte-textured 3D shapes are perceived differently in several ways. In dynamic scenes, these differences might be related to the information that their respective optic flows convey. We review these differences next.

Peculiarities of specular and matte optic flow

Consider the object in Figure 1a (adapted from Adato, Vasilyev, Zickler, & Ben-Shahar, 2010), whose Gaussian curvature map is shown (from above) in Figure 1b. Figure 1c shows a corresponding optic flow pattern that would be generated by this object if it had been rotated about a horizontal in-plane axis, and Figure 1d approximates the optic flow pattern that would be generated had the object been made of a diffusely reflecting textured surface. Arrows indicate the flow direction whereas colors denote the flow magnitude. Comparing Figure 1c to the curvature map in Figure 1b nicely illustrates how optic flow generated by a specular object (a.k.a. the *specular flow*) carries information about second-order shape properties. In particular, characteristic for specular objects are the singularities, in both magnitude and direction, generated by the shape's parabolic lines (i.e., 0-isolines of Gaussian curvature; yellow-red colors in Figure 1c). Not only does the flow magnitude in these locations blow up to infinity, but its direction also constitutes sinks or sources in the field (Adato et al., 2010). At the same time, relatively little information is carried by specular flows about surface slant and tilt or object motion (Doerschner, Kersten et al., 2011; Koenderink & Van Doorn, 1980).

In contrast to specular flows, optic flow generated by a moving matte-textured object (hereafter simply "optic flow") conveys visual information about the object's motion and its first-order shape properties such as local slant and tilt (Figure 1d). Specifically, optic flow direction is dominated by the motion of the object, and its magnitude is directly modulated by surface slant.

What is the effect of motion characteristics?

Next, we vary the axes about which the object rotates in the scene. Adapted again from Adato et al. (2010), Figures 1e and f depict a rotation about the vertical axis whereas Figures 1g and h depict rotation about the viewing axis (perpendicular to the image plane). Two observations are immediate. First, the specular and optic flows are consistently and dramatically different from each other. Second, although optic flows appear dependent on object motion, considerable portions of the specular flow remain highly robust and agnostic to the particular way the object moves. Perhaps most obvious is the case of the viewing axis rotation (Figure 1g, h), in which the specular flow remains informative about the shape of the object (e.g., about its singularities), whereas the matte flow is still dominated by the object motion but, in this case, marginally modulated by surface slant.

One might speculate that the informal observations made above will manifest themselves in the ability to perceive the shape of the object, particularly in the perceptual judgements of 3D curvature magnitude. The experiments in this article were designed to examine this very idea.

Does specular flow facilitate or interfere with 3D shape estimation?

Figure 1 suggests that specular flow carries information about surface curvature magnitude more consistently than matte optic flow. Unfortunately, the perceptual evidence to that effect is inconclusive. For example, studies by Doerschner et al. (2013), Dövencioğlu et al. (2015), and Vota et al. (2015) suggest that estimated curvature, slant, tilt, and object rotation axis are quite different for moving specular objects

Downloaded from jov.arvojournals.org on 07/24/2020

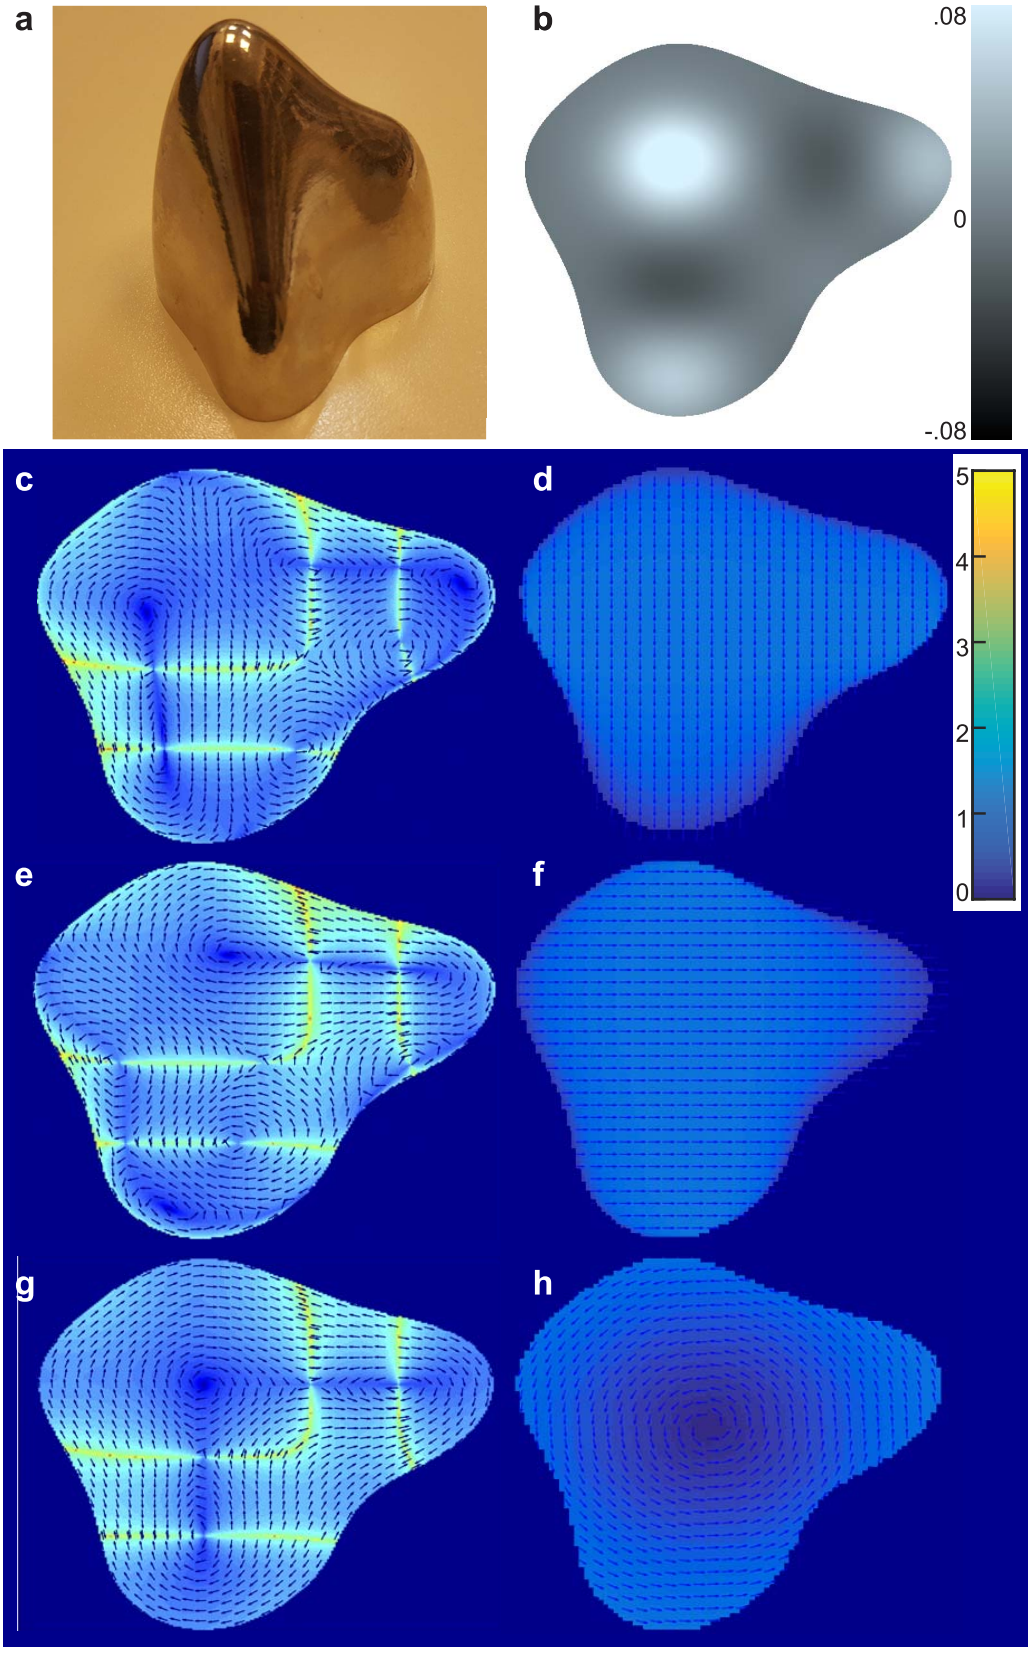

Figure 1. Specular flow and optic flow. (a) Specular object of simple parametric shape that we use to illustrate the two different types of flows. (b) Corresponding Gaussian curvature map for a top view. (c, e, g) Specular flow generated by different rotations of the environment (adapted from Adato et al., 2010). (d, f, h) Corresponding optic flows generated by rotations of the same parametric shape, this time with an ideal matte-textured surface material. Arrows on the plots show direction of flow where the colors show the magnitude. Each row shows optic flow for one type of rotation: horizontal (c, d), vertical (e, f), around the viewing axis (g, h). Arrows give the direction of the flow vectors, and color corresponds to flow magnitude (blue values denote lower magnitudes).

Downloaded from jov.arvojournals.org on 07/24/2020

compared with their matte-textured counterparts (Mazzarella, Cholewiak, Phillips, & Fleming, 2014). But even if motion is taken out of the equation (Blake & Bulthoff, 1990, 1991; Mazzarella et al., 2014), findings about the contribution of specular reflections to estimated 3D shape are conflicting. For example, although Fleming, Torralba, and Adelson (2004) suggested that specularities aid and constrain local shape perception in images, Savarese, Fei-Fei, and Perona (2004) reported the opposite.

In this study, we focus on the interaction of motion and material and investigate whether specular flow (i.e., the image optic flow generated by specular moving objects) can provide better information on 3D shape than optic flow (i.e., the image flow generated by matte-textured objects). As mentioned above, specular flows are directly related to 3D curvature and seem to be less sensitive to the particular motion of the object, whereas optic flows vary more substantially with the latter. Thus, if a perceptual task required observers to make judgments about an object's 3D curvature structure, for example, its bumpiness, one would expect more consistent shape perception across changes in object rotation axis. We next describe two experiments that put this idea to the test.

Methods

Overview

Observers compared the bumpiness of two rotating objects presented successively in each trial. Objects varied in curvature magnitude and surface reflectance properties. A critical manipulation in the experiments was the rotation axis of the objects, because we predicted that curvature-related judgments for matte-textured objects would be more susceptible to changes in the object's rotation axis In Experiment 1, the reference object was always specular, whereas in Experiment 2, it was always matte textured. In every other respect, the two experiments were identical; therefore, we combined the respective methods sections into one.

Stimuli

3D Models

Stimuli were images of bumpy spherical objects (see Figure 2a). The objects were created in 3DS MAX Autodesk with a custom script. We perturbed each sphere with sine wave modifiers, where the phase of the waves varied randomly between $-\pi$ and π and the wavelengths varied between $\pi/6$ and $\pi/4$. There were five amplitude conditions. For each, the amplitude of sine wave modulations were chosen randomly between 0 and a fixed number. For example, the amplitude of the sine wave modulations of the object with bumpiness Level 5 were selected between 0 and 6, whereas those for the object with bumpiness Level 3 were chosen between 0 and 3. Note that these manipulations generated variations in curvature that were not parametric and did not increase in uniform steps (maximum radius for each amplitude condition: 7.86, 8.20, 8.57, 8.65, and 9.09 for bumpiness Level 1–5, respectively, where the minimum radius was always 7.5). We describe these stimulus levels in terms of the difference between maximum and minimum Gaussian curvatures averaged over all vertices $mean(C_{max} - C_{min}) = [0.0081, 0.0180, 0.0208, 0.246, 0.0324]$, respectively, for each stimulus level. We will refer to the curvature magnitude of our stimuli also as *bumpiness* in the remainder of the article. Note that with monotonically increasing bumpiness, the amount of curvature magnitude $mean(C_{max} - C_{min})$ also increased (Figure 2a). Stimuli subtended approximately 9.5° visual angle.

Rotation axes

Objects rotated back and forth 20° at a rate of 12°/s (10° each direction, 0.2°/frame with 60 frames per second) around one of three axes passing through the objects' center. Two of these were in-depth rotations around vertical or horizontal axes, and one was a rotation around the viewing axis (respective sample movies can be found at https://vimeo.com/169838093 and https://vimeo.com/169837852).

Surface materials

Mirror-like *specular* shapes were rendered using *environment mapping* (Debevec, 2002) with two *light probes*; Debevec's *Grace* probe (see Figure 2; Debevec, 2002) and, for variety in the stimuli, a desaturated and phase scrambled version of it (Doerschner, Fleming, et al., 2011).[1] Diffusely reflecting *matte-textured* objects were generated with a "stuck-on" effect. The stimuli in the specular case are specular (i.e., black, except for highlights), as would be a mirror. In contrast, the matte stimuli are diffuse, and reflectance variations are entirely due to textures. Only the mixed stimuli show a combination of diffuse and specular reflectance. This is detailed in the Appendix A. In other words, when viewed statically in their initial viewpoint configuration, rendered images of matte-textured and specular objects were identical (Figure 2b). A "mixed" material condition was generated by using a weighted combination of specular and matte-textured materials. These stimuli were tone mapped to ensure that their overall luminance was matched to the other two materials (Appendix A).

Downloaded from jov.arvojournals.org on 07/24/2020

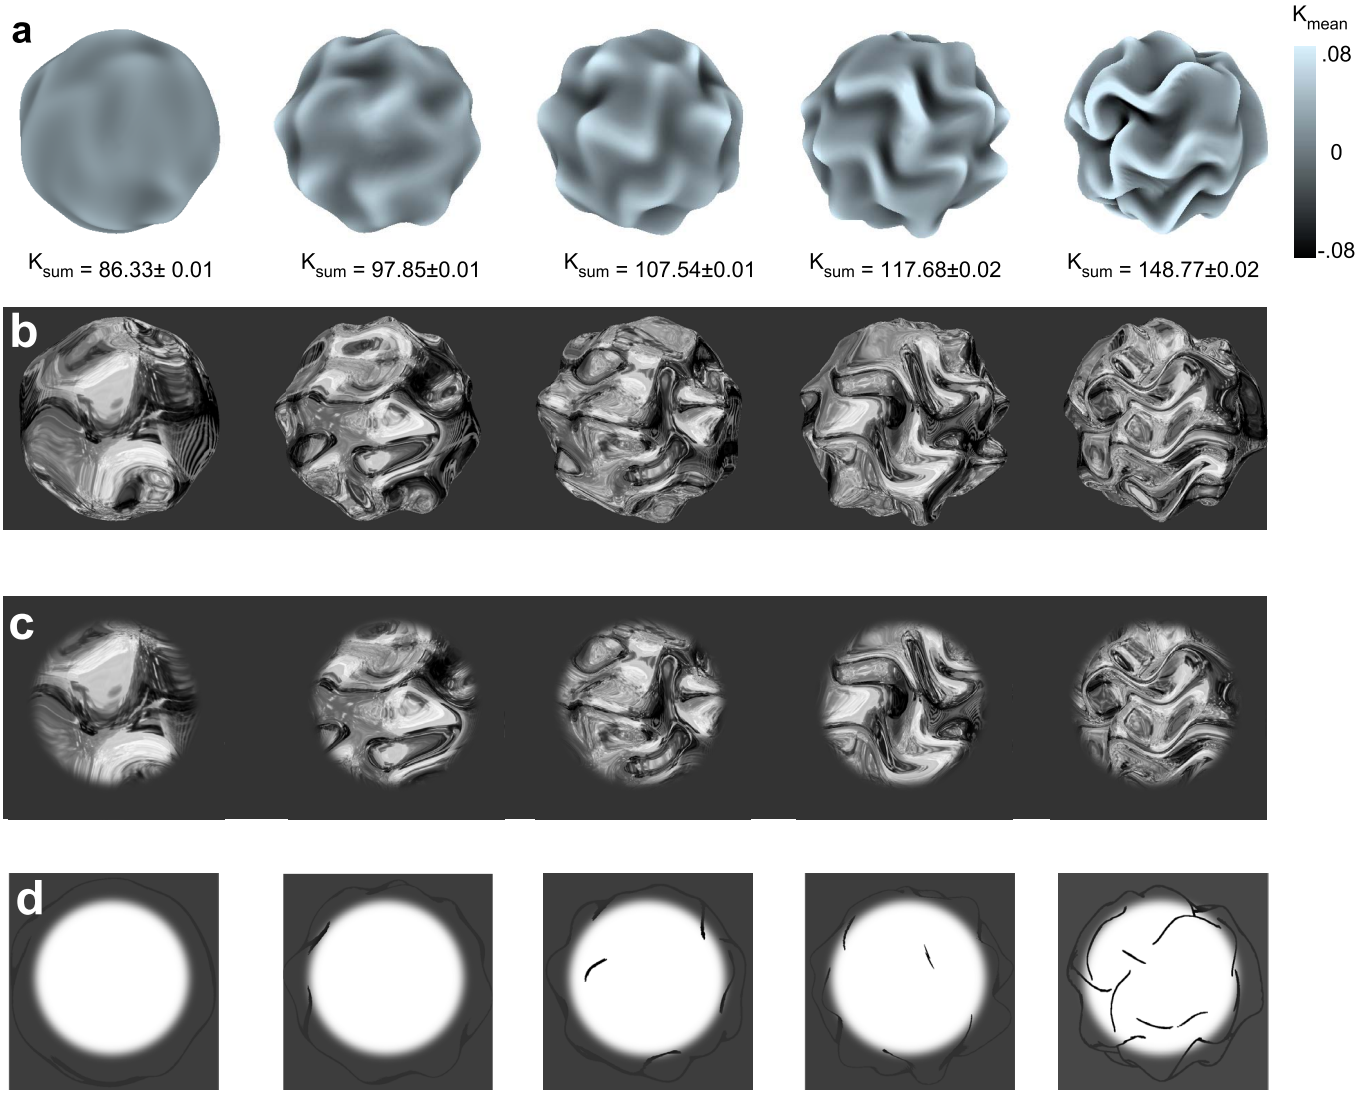

Figure 2. Stimuli in Experiments 1 and 2. The bumpiness level of the objects varied at five levels from flattest (first column) to bumpiest (last column). The reference object was always at mid-level (third column). In the first row (a), we present the mean curvature of individual vertices as overlaid on the 3D mesh. Below each object, we report the sum of the absolute mean curvature values summed over all vertices (8,066 for each object). In panel b, screen shots of the 3D models rendered with the environment map are shown, including the object boundary. Note that this information was never available during the experiments. Observers always saw stimuli through a Gaussian aperture in order to exclude the object boundary (c). (d) Corresponding viewpoint plots containing the best possible view onto self-occluding contours for each stimulus. The white area corresponds to unmasked parts of the object.

Object boundaries

To eliminate information to 3D shape from the object's occluding boundaries, stimuli were always presented to observers through a Gaussian aperture. Figure 2c shows rendered images of stimuli at all levels of bumpiness with the superimposed Gaussian aperture. Figure 2d depicts self-occlusion boundaries of the stimuli, because the motion of these might act as additional cues to 3D shape. However, as Figure 2d shows, there is basically no information available from the self-occlusions below bumpiness Level 5 (rightmost column). Note that even if we determined the self-occlusion information, these regions were not masked or blurred in the experiments.

Reference object

In both experiments, observers compared objects to a reference object of bumpiness level 3. The surface reflectance of this reference object was always specular in Experiment 1 and always matte textured in Experiment 2. Although it would have been theoreti-

Downloaded from jov.arvojournals.org on 07/24/2020

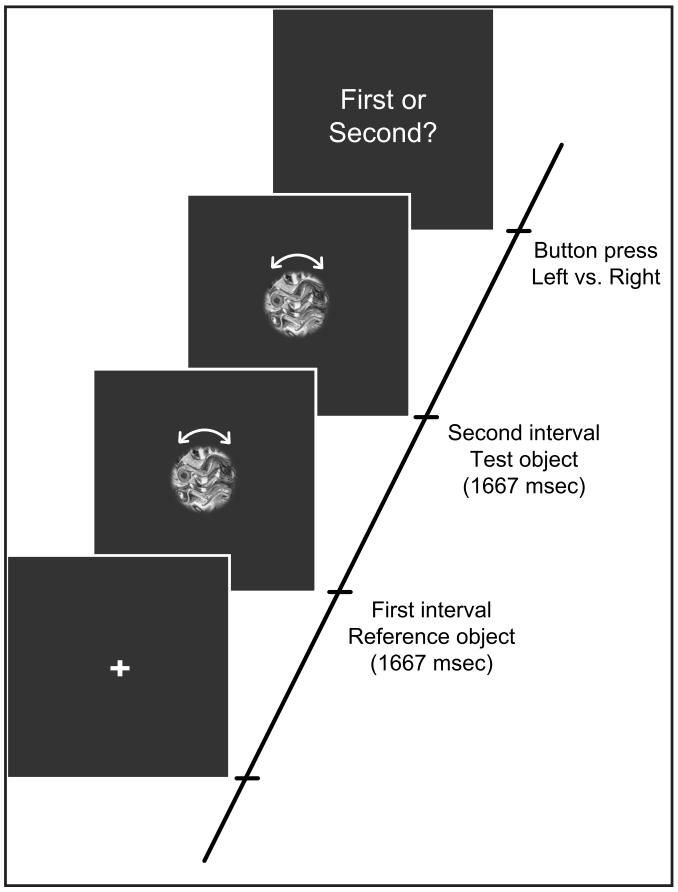

Figure 3. Example trial in Experiments 1 and 2. Reference and test objects were displayed sequentially in randomized order. After the two intervals, a fixation cross was displayed until the observer made a keyboard response.

cally possible to run a full factorial design to investigate our experimental question, we opted for a more economical design (i.e., just one reference object) to facilitate experiments with more reasonable length. However, because. with this design we introduce an asymmetry in the exposure to matte and specular materials, we needed—at the very least—to check whether and how this asymmetry influences the results. Details follow in the Results section.

Apparatus

Frames were rendered using dedicated GLSL shaders in the Gratin system (Vergne, Ciaudo, & Barla, 2014), and movies were displayed on a CRT monitor (HP P1230, 22 in., 1,024 × 1,280 resolution, 60 Hz) using a custom software written with Psychtoolbox v3.0.11 routines (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997). Observers viewed the monitor binocularly at a distance of 57 cm where their heads were stabilized with a chin rest. They pressed left and right arrow keys on a keyboard to indicate their response for the first or second object, respectively.

Tasks

Observers viewed two objects successively and indicated which stimulus appeared bumpier (or less bumpy): first or second (Figure 3). On every trial, one interval contained *a reference object*. The other interval contained a test object of one of the three materials (specular, matte textured, mixed) and one of the five bumpiness levels. Figure 3 shows the order of events in a typical trial. The rotation axes and environment maps of both objects in a trial varied together (i.e., in a given trial); both intervals showed the same environment map and motion characteristics, but bumpiness and material changed according to the experimental condition. To prevent pure image-based comparisons, the orientation of the two objects to be compared varied randomly on every trial.

Half of the observers judged which of the two objects appeared *bumpier* ("Which object has greater bumps?"), whereas the other half of the observers judged which of the two objects appeared *flatter* ("Which object has smaller bumps?"). We reasoned that by presenting the task in two different ways about the same object property, we would get an estimate of observers understanding of the task and robustness of their internal concept of *bumpiness* and how it is related to our manipulation of curvature. Ideally, we would hope to find no differences between the two task types.

Procedure

All experimental conditions were interleaved in a single block. Observers completed this block in one session but were allowed to take breaks and were instructed to complete the task at their own pace. In Experiment 1, observers completed at least 1,350 trials[2] (five objects × three materials × three rotation axes × two light probes × 15 repetitions)[3] in under 3 hr on average.

Prior to the experiment, observers went through an informal task familiarization procedure, in which they were shown two Lambertian objects (with uniform albedo) having our highest and lowest bumpiness levels, respectively, and were asked—as in the experiment—to report which object had greater or smaller bumps. No feedback was given. In some cases in which observers asked for further description of bumpiness, they were instructed to pay attention to the amplitude not the frequency: *Bumpier* means bumps with larger amplitude rather than a greater number of total bumps. The objects used in these practice trials were different from those

Downloaded from jov.arvojournals.org on 07/24/2020

used in the actual experiments but were constructed in the same way as described in the Stimuli section.

Observers

Twenty-eight naive observers from Bilkent University volunteered to participate in the study; 14 of these participated in Experiment 1 and 14 in Experiment 2. In each of the two experiments, seven observers made "bumpier" judgments and the seven others "flatter" judgments. All observers had normal or corrected-to-normal vision. Experiments were approved by the Bilkent University Ethics Review Board and were in agreement with the Declaration of Helsinki. Observers gave written informed consent prior to the experiment and were paid 15 Turkish Liras (about 3.8 Euros) per hour.

Analysis

For each observer in each condition of Experiments 1 and 2 and each task type (bumpier and flatter), we calculated the proportions of the number of trials in which test objects were judged bumpier (or flatter) than the respective reference object. We then estimated slopes and intercepts of lines fitted to the "judged bumpier/flatter" proportions (using "polyfit" in Matlab; MATLAB, 2014). The slopes and intercepts essentially provide a measure of how discriminable reference and test objects were. Although differences in slopes might be related to differences in observers' sensitivity to manipulations of bumpiness (larger slope indicates higher sensitivity), differences in intercepts (e.g., an upward shift) might indicate a perceptual bias and help determine if certain materials or rotations around certain axes may *always* look bumpier (but are not necessarily more discriminable).

Employing also an analysis of variance (ANOVA), we investigated the effects of surface material (specular, mixed, matte textured) and rotation axis (horizontal, vertical, viewing axis) and task type (bumpier, flatter) on bumpiness discrimination ability (slopes and intercepts of the proportion data). We expected significant interactions between surface material and rotation axis. All analyses were conducted in Matlab (MATLAB, 2014) and SPSS (IBM Corp., 2010).

Results

Overall

To determine whether the fits are optimal for the data, we used a nested hypothesis test to test the hypothesis that the observers' bumpiness discrimination is a constant (Mood, Graybill, & Boes, 1974). The constant (p0), linear (p1), and quadratic (p2) models fitting the data are shown in Figure 4. The case of the constant hypothesis would suggest that the observer would pick the test object at the same frequency regardless of the object's bumpiness level. The log likelihood of the unconstrained model (p1) was fitted to observers' test-judged-bumpier proportions using the maximum likelihood for the group data in each experimental condition (six panels in Figure 4). In the constrained model, we forced the fitted model to be a constant, and we compared log likelihood ratio to the relevant chi-square distribution (χ^2_1). For all experimental conditions, proportions were significantly different from zero ($p < 0.000001$), which means that we can reject the null hypothesis. We then tested whether quadratic models fit the data better (p2 vs. p1), but linear fits were not different from quadratic fits (for all conditions, $p > 0.43$). Parameters for all models are presented in Table 1. This analysis shows that (a) the observers can do the task and (b) we reject the hypothesis that quadratic fits are significantly better than linear fits. Hence, we carry on reporting the rest of the results of linear fits.

Experiment 1: Specular reference object

Slopes: Results shown in Figure 5 are explained here in detail. The three (material) × three (rotation axis) × two (task-type) ANOVA on the bumpiness discrimination ability revealed a main effect of material, $F(2, 24) = 3.971$, $p < 0.05$; see Table 3 for mean slopes per condition); a main effect of rotation axis, $F(2, 24) = 10.348$, $p < 0.05$; and no main effect of task type. The two-way interaction between material and rotation axis was significant, $F(4, 48) = 10.224$, $p < 0.0001$. The remaining two-way and the three-way interaction yielded no significant results.

Following up the main effect of *material*, a post hoc analysis (pairwise comparisons in the same ANOVA) revealed that the slopes were not significantly different between materials.

Figure 5 illustrates the nature of the *two-way interaction between material and rotation axis*. Although for specular and mixed materials there is no difference in slopes between in-depth and viewing axis rotations, slopes for matte-textured objects were significantly higher for in-depth rotations than viewing axis rotations.

Overall, slopes were positive, suggesting that participants can in fact do the task. Notably, comparing a specular test object to a specular reference object yielded more consistent results (smaller error bars in Figure 5) than between material comparisons.

Downloaded from jov.arvojournals.org on 07/24/2020

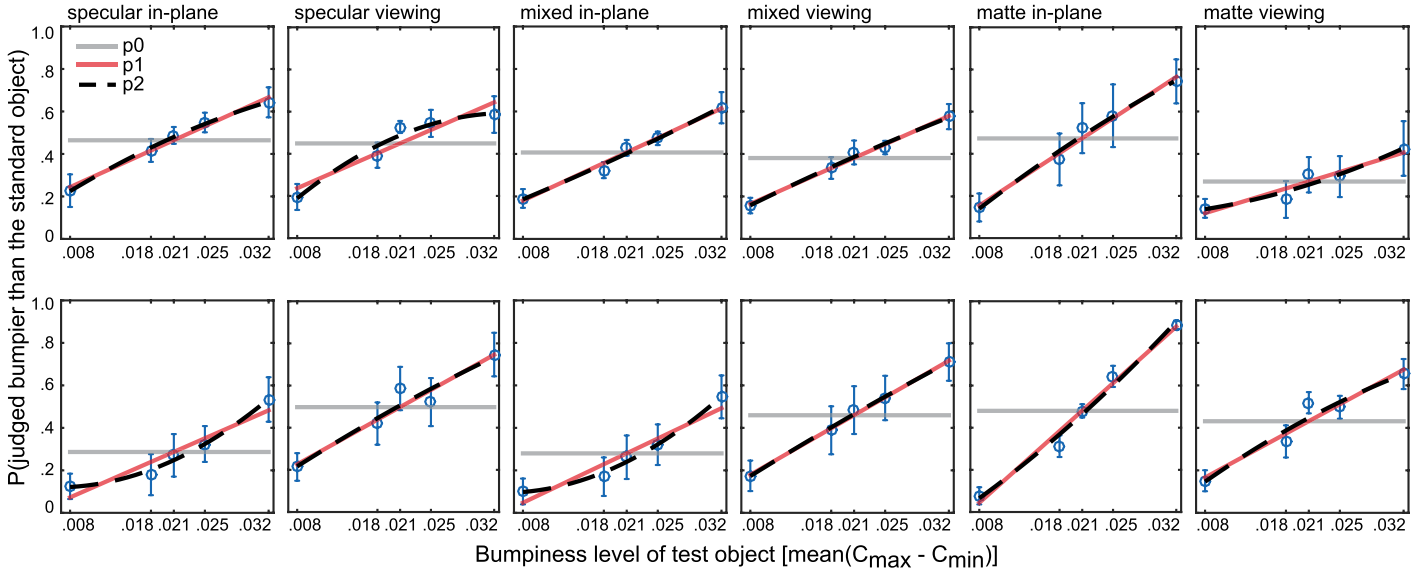

Figure 4. Group results from Experiments 1 and 2. Shown are mean proportions of test objects judged bumpier than the *specular* (first row) and the *matte* (second row) reference object. The group data presented here are averaged over task type and in-depth rotation axes. Model parameters are listed in Table 1. Results for each material type are shown in separate panels as indicated by titles above plots. Blue icons represent group means, where error bars are ±2 *SEM*. In each panel, the gray line indicates a constant model, and the red line and black dashed line fit a linear and a nonlinear model, respectively.

Intercepts: The three (material) × three (rotation axis) × two (task-type) ANOVA on the intercepts of lines in Figure 5 revealed a main effect of material, $F(2, 24) = 7.780$, $p < 0.001$ (see Table 2), no main effect of task type, and no main effect of rotation axis. None of the two-way and three-way interactions were significant. Pairwise comparisons (post hoc in same ANOVA) showed that intercepts for specular conditions were significantly higher than for those in the mixed condition ($\Delta\mu_{specular-mixed} = 0.067$, $p < 0.004$); differences of the two materials with the matte condition remained nonsignificant.

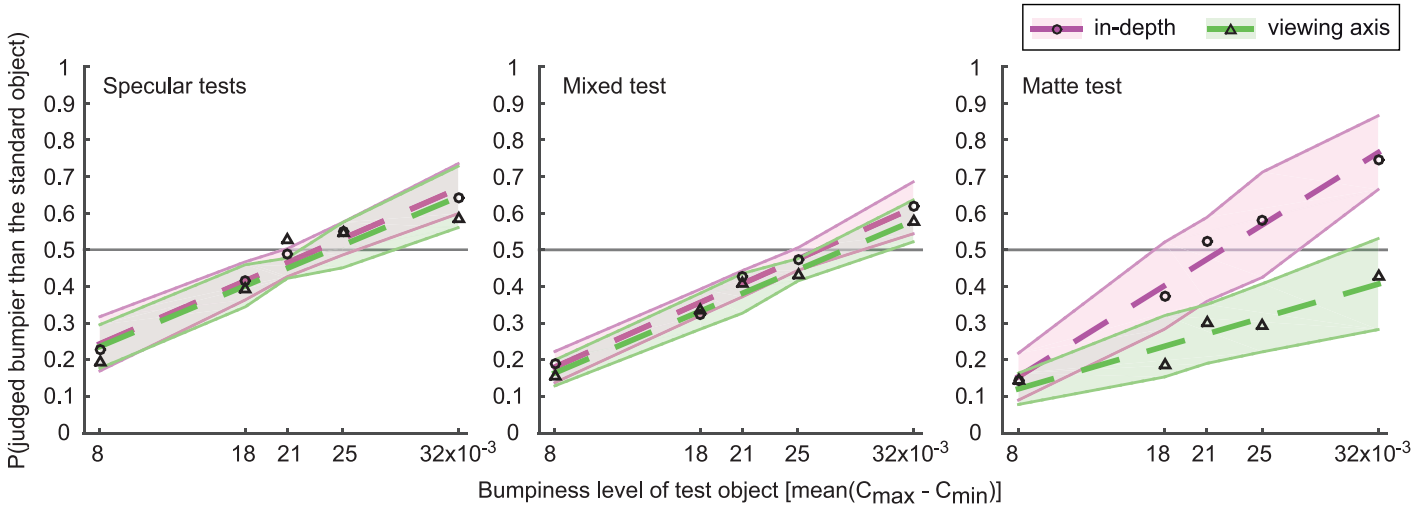

Figure 5. Results from Experiment 1. Shown are mean proportions of test objects judged bumpier than the *specular* reference object. For simplification, we plot results over the two task types combined: To do so, we combined the *judged flatter* data with the *judged bumpier* proportions. We also combined the data for the two in-depth rotation axes, because we found no difference in slopes. Results for each material type are shown in separate panels. In each panel, black circles and pink dashed lines correspond to results for in-depth rotations and corresponding line fits, respectively. Black triangles and green dashed lines correspond to results for viewing axis rotations and corresponding line fits, respectively. Shaded error regions are 2 *SEM*.

Downloaded from jov.arvojournals.org on 07/24/2020

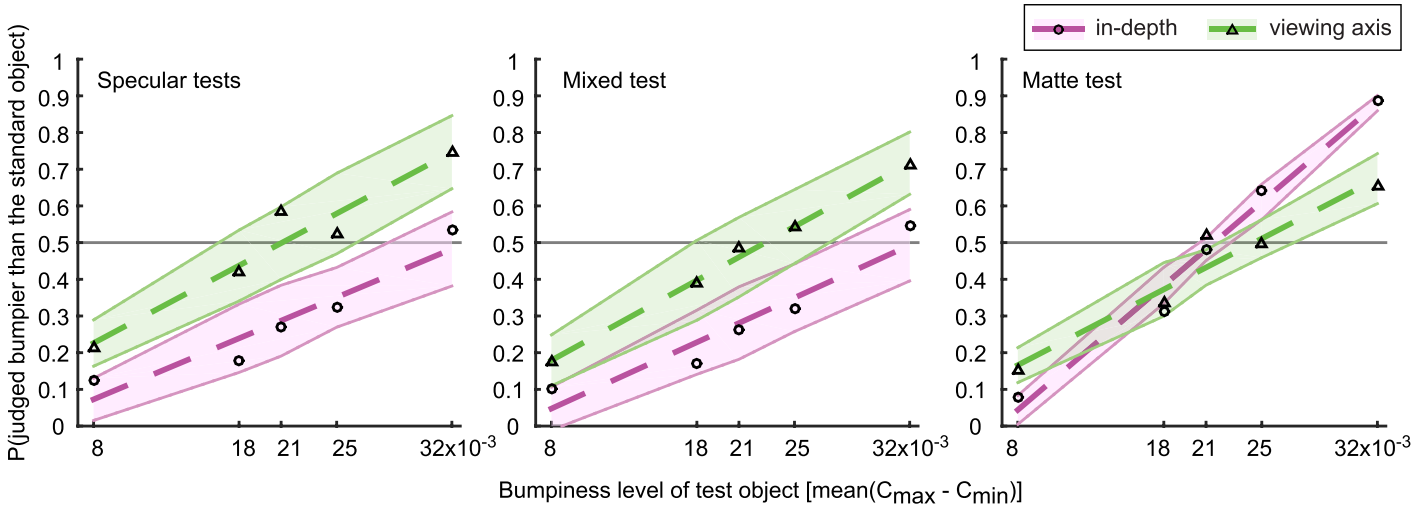

Figure 6. Results from Experiment 2. Shown are mean proportions of test objects judged bumpier than the *matte* reference object. As in Figure 5, the data presented here are averaged over task type and in-depth rotation axes. Results for each material type are shown in separate panels. In each panel, black circles and pink dashed lines correspond to results for in-depth rotations and corresponding line fits, respectively. Black triangles and green dashed lines correspond to results for viewing axis rotations and corresponding line fits, respectively. Shaded error regions are 2 *SEM*.

Experiment 2: Matte reference object

In contrast to Experiment 1, the reference object in Experiment 2 was matte textured. All other aspects were the same, and the analysis followed the same steps. The results shown in Figure 6 are explained below.

Slopes: The three (material) × three (rotation axis) × two (task-type) ANOVA on the bumpiness discrimination ability revealed a main effect of material, $F(2, 24) = 41.944$, $p < 10^{-6}$ and no main effect of rotation axis and no main effect of task type (see Table 5 for mean slopes per condition and see Table 4 for the ANOVA results). As in Experiment 1, the two-way interaction between material and rotation axis was significant, $F(3.15, 37.85) = 37.211$, $p < 10^{-6}$. The three-way interaction was significant, $F(4, 48) = 37.211$, $p < 0.05$, and the remaining two-way interactions yielded no significant results.

Post hoc analysis on the main effect of *material* revealed that slopes for specular and mixed materials were smaller than the slopes of the matte-textured material ($\Delta\mu_{specular–matte} = 0.066$, $p < 0.003$, $\Delta\mu_{mixed–matte} = 0.059$, $p < 0.003$, Bonferroni corrected). Slopes of the specular and mixed materials ($\Delta\mu_{specular–mixed} = –0.007$) were not significantly different.

Post hoc analysis following up the main effect of *rotation axis* revealed no significant differences in the pairwise comparisons of rotation axes ($\mu_{vertical} = 0.130$, $\mu_{horizontal} = 0.134$, $\mu_{viewing} = 0.119$).

Figure 6 illustrates the two-way interaction between material and rotation axis. Whereas for specular and mixed materials there were no differences in slopes between in-depth and viewing axis rotations, for the matte-textured material, slopes in the in-depth rotation condition were higher than for the viewing axis rotations.

As in Experiment 1, slopes were positive, suggesting that participants could do the task. Also in Experiment 2, within-material comparisons yielded more consistent results (smaller error bars in Figure 6) than between-material comparisons.

Intercepts: The three (material) × three (rotation axis) × two (task-type) ANOVA on the intercepts of lines in Figure 5 revealed a main effect of material, $F(2, 24) = 10.837$, $p < 10^{-5}$; a main effect of rotation axis, $F(2, 24) = 17.748$, $p < 10^{-5}$; and no main effect of task type. The two-way interaction between material and rotation axis was significant, $F(4, 48) = 3.976$, $p < 0.001$. No other two-way or three-way interactions were significant.

Overall, results of Experiments 1 and 2 were highly consistent, and we discuss their implications next.

Discussion

Summary

The specific dependencies of specular flow on the 3D curvature structure of an object lead us to hypothesize that the perceived bumpiness of specular objects should be more invariant across changes in rotation axis than the perceived bumpiness of a corresponding matte-textured shape. The results from our two experiments supported this idea, showing significant changes in bumpiness comparisons across changes in rotation axis

Downloaded from jov.arvojournals.org on 07/24/2020

	*p*0	*p*1	*p*2
Experiment 1			
Specular in-line	$c = 0.46$	$y = 0.10 + 1.74x$	$y = 0.02 + 2.7x - 2.38x^2$
Specular viewing	$c = 0.45$	$y = 0.10 + 1.68x$	$y = -0.10 + 4.09x - 6x^2$
Mixed in-line	$c = 0.41$	$y = 0.04 + 1.79x$	$y = 0.06 + 1.48x + 0.76x^2$
Mixed viewing	$c = 0.38$	$y = 0.03 + 1.71x$	$y = 0.00 + 2.02x - 0.78x^2$
Matte in-line	$c = 0.47$	$y = -0.05 + 2.51x$	$y = -0.11 + 3.26x - 1.86x^2$
Matte viewing	$c = 0.27$	$y = 0.03 + 1.17x$	$y = 0.11 + 0.24x - 2.34x^2$
Experiment 2			
Specular in-line	$c = 0.29$	$y = -0.06 + 1.69x$	$y = 0.14 - 0.78x + 6.13x^2$
Specular viewing	$c = 0.50$	$y = 0.05 + 2.14x$	$y = 0.01 + 2.63x - 1.23x^2$
Mixed in-line	$c = 0.28$	$y = -0.10 + 1.83x$	$y = 0.11 - 0.71x + 6.31x^2$
Mixed viewing	$c = 0.46$	$y = 0.00 + 2.21x$	$y = -0.03 + 2.51x - 0.75x^2$
Matte in-line	$c = 0.48$	$y = -0.24 + 3.44x$	$y = -0.14 + 2.36x - 2.69x^2$
Matte viewing	$c = 0.43$	$y = 0.00 + 2.09x$	$y = -0.09 + 3.07x - 2.43x^2$

Table 1. Model parameters for Experiments 1 and 2. *Notes:* Shown are the parameters for each model fitted to six experimental conditions in Experiments 1 (first six rows) and 2 (last six rows). The first column of each experiment gives the fitted value for a constant model (*p*0), the second column gives two values for a linear fit (*p*1), and the last column gives three values for a quadratic fit.

for matte-textured shapes but not for objects that were specularly reflecting. Specifically, matte stimuli were less discriminable in terms of their bumpiness when rotating around the viewing direction. This effect was independent of whether a matte object was compared with a specular or a matte reference shape. Surprisingly, overall specular objects tended to be less discriminable than matte-textured objects.

What does perceived bumpiness really depend on?

There was a systematic relationship between the curvature characteristics of the object and its perceived bumpiness, even though the manipulation of overall curvature magnitude was not parametric for our stimuli. In addition to the image motion generated by different surface materials, the 3D structure of a moving object could also be conveyed by other sources of information. The object boundary is, for example, a very prominent cue to 3D shape (Barrow & Tenenbaum, 1981; Ramachandran, 1988; Schofield, Rock, & Georgeson, 2011; Todorović, 2014; Wagemans, Van Doorn, & Koenderink, 2010), but here we eliminated its contributions by masking it. The motion of self-occlusions might be another cue to 3D shape, particularly at higher bumpiness levels (see Figure 2; Karsch, Liao, Rock, Barron, & Hoiem, 2013). Although these might have contributed to the overall bumpiness estimate, they could certainly not account for the systematic differences in our data between specular and matte-textured objects. It would be an interesting question to investigate how all three cues interact in a shape estimation task in future experiments.

Matte-textured bumpiness

Given that there is relatively less information about 3D curvature in the optic flow of matte objects for rotations around the viewing axis, it is surprising how well observers can do the task. What exactly are observers comparing here? One possibility is that they judge local image cues, such as the compression patterns of texture (Fleming et al., 2004). Note that we

	Slopes			Intercepts		
Experiment 1	*F*	*df*	η^2	*F*	*df*	η^2
Material	3.971*	2, 24	0.249	7.780**	2, 24	0.393
Rotation	10.348*	2, 24	0.463	0.402	2, 24	—
Material × Rotation	10.224***	4, 48	0.460	1.285	4, 48	—
Material × Task	1.384	2, 24	—	0.032	2, 24	—
Rotation × Task	0.814	2, 24	—	0.875	2, 24	—
Three-way interaction	0.099	4, 48	—	0.177	4, 48	—

Table 2. ANOVA results for Experiment 1. *Notes:* Shown are main effects, two- and three-way interactions for slopes and intercepts data in Experiment 1. *$p < 0.05$. **$p < 0.01$. ***$p < 0.001$.

Downloaded from jov.arvojournals.org on 07/24/2020

	Slopes		Intercepts	
Experiment 1	μ	*SE*	μ	*E*
Specular	17.15	2.62	0.103	0.057
Mixed	17.54	1.85	0.034	0.039
Matte	20.61	2.18	−0.023	0.038
Horizontal	19.50	2.33	0.030	0.044
Vertical	20.65	2.41	0.032	0.050
Viewing	15.15	1.90	0.052	0.038
Bumpier	17.37	2.41	0.049	0.059
Flatter	19.49	1.90	0.026	0.059

Table 3. Mean values for Experiment 1. *Notes:* Shown are the group means and standard deviations for slopes and intercepts data for each material and rotation in Experiment 1.

	Slopes		Intercepts	
Experiment 2	μ	*SE*	μ	*SE*
Specular	18.33	1.98	−0.024	0.044
Mixed	19.53	1.97	−0.066	0.045
Matte	29.83	1.27	−0.156	0.028
Horizontal	22.63	1.59	−0.122	0.039
Vertical	22.64	1.71	−0.142	0.044
Viewing	21.42	1.71	0.018	0.036
Bumpier	22.89	2.24	−0.054	0.051
Flatter	22.25	2.24	−0.110	0.051

Table 5. Mean values for Experiment 2. *Notes:* Shown are the group means and standard deviations for slopes and intercepts data for each material and rotation in Experiment 1.

used "stuck-on" reflections for matte objects that essentially mimic the distortions of specular reflections at a given frame. Thus, bumpier matte objects would have had more compressed texture regions than flatter shapes. This implies that for the matte-matte comparisons, we might have obtained similar results if we had used static stimuli, which would be consistent with our previous findings (Dövencioğlu et al., 2015). Nevertheless, the fact that judgments of the same matte-textured object depended on its rotation axis orientation shows that motion does play a role in bumpiness perception beyond texture compression.

What makes matte objects rotating around the viewing axis much harder to discriminate in terms of their bumpiness than those rotating in depth? In-depth rotations generate quite powerful motion parallax: Points closer to the observer (peaks) move faster with respect to the observer than points that lie more inwardly on the shape (troughs). The larger the difference between peaks and troughs, the stronger the motion parallax. Thus, parallax serves as a cue to relative surface depth and the object's bumpiness. When an object rotates around the viewing axis, the parallax cue is much less indicative of 3D structure because the viewpoint with respect to the object remains stable; thus, the distance of points from the camera does not change during rotation, unlike in the in-depth case. Moreover, the closer object parts are to the rotation axis, the less image motion they produce (see Figure 1). As a result, the view-axis rotation is actually compatible with a 2D rotation, where motion alone cannot convey 3D shape. Thus, an observer in this condition is left with rather little information about 3D structure from motion. Therefore, we suggest that for viewing axis rotations, observers substantially relied on the texture compression information when the comparing bumpiness of objects.

Specular bumpiness

As expected, the comparisons of bumpiness of specular objects turned out to be more invariant across changes in rotation axis. Specular flow patterns for a given curvature magnitude remain more similar across changes in rotation axes, at least around singularities, and if an observer was to interpret changes in local specular flow magnitudes as corresponding changes in curvature magnitude, she or he would be quite successful in the task. However, for in-depth rotations, discriminability of specular objects was overall lower than that of matte shapes, and for viewing axis rotations, it never exceeded that of matte objects. Why? Points on a specular rotating object do not move consistently and do not generate powerful motion

	Slopes			Intercepts		
Experiment 2	*F*	*df*	η^2	*F*	*df*	η^2
Material	41.944***	1.22, 14.71	0.778	10.837***	2, 24	0.475
Rotation	2.855	2, 24	—	17.748***	2, 24	0.597
Material × Rotation	37.211***	3.15, 37.85	0.756	3.976**	4, 48	0.249
Material × Task	1.317	2, 24	—	0.348	2, 24	—
Rotation × Task	1.314	2, 24	—	0.279	2, 24	—
Three-way interaction	2.634*	4, 48	0.180	1.238	4, 48	—

Table 4. ANOVA results for Experiment 2. *Notes:* Shown are main effects, two- and three-way interactions for slopes and intercepts data in Experiment 1. *$p < 0.05$. **$p < 0.01$. ***$p < 0.001$.

Downloaded from jov.arvojournals.org on 07/24/2020

parallax information. In fact, some participants reported that the specular objects sometimes appeared fluidlike or nonrigid.[4] This is not surprising given that, *under a rigidity assumption,* there are many violations of correspondence in specular flow patterns (i.e., around singularities in the optic flow caused by sinks and sources). Although these violations could make it theoretically quite hard to construct the 3D structure of a specular surface (Doerschner, Fleming et al., 2011), it might also be that identification of these singularities may be a cue for shape. Whether or not the visual system has a rigidity prior has been a topic of a debate (Jain & Zaidi, 2011).

Our results show that observers appear to use parallax and texture information to evaluate 3D structure more readily than cues provided by specular flow. Nevertheless, they are able to use specular flow information successfully. And in some conditions (e.g., rotations around the viewing axis), specular flow should provide better information about 3D structure than optic flow generated by matte-textured objects. Although we see a trend in the data supporting this latter idea, results did not reach statistical significance. Remember, however, that we defined invariance essentially as the similarity of slopes (sensitivity) across conditions. Although the slopes for specular objects under all conditions remain the same, they do change significantly for the matte-textured objects with changes in rotation axis.

In our experiments, we found no difference in the results between fully and partly (mixed) specular objects. This implies that specular flow cues to bumpiness clearly dominated the percept and somehow rendered the matte-textured optic flow information inaccessible to the observer. If we had chosen a much lower weight for the specular component, we might have found interactions between two cues. At what weight value an interaction of the two types of flow information might occur could be a question for future study.

Task difficulty

The starting point of our investigation was the observation that a moving object generates quite different optic flow patterns depending on whether its surface material is specular or matte textured; therefore, it is not necessarily a surprising finding that observers are more consistent when comparing bumpiness levels for shapes of the same material (smaller error regions). To make the task nontrivial (especially for within-material comparisons), we prevented any image-based matching strategy by randomly jittering an object's initial orientation on each trial. When comparing bumpiness of shapes with different surface reflectance, the task seemed to be much harder for observers.

The reference object

Changing the reference object from specular (Experiment 1) to matte textured (Experiment 2) caused a significant change in bias for specular and mixed materials for in-depth rotations. Compared with the matte reference, these objects appeared much less bumpy (downward shift of the line), although their discriminability remained the same across all conditions. How can these shifts be explained? We suggested above that parallax information may allow observers to construct 3D depth information more robustly than specular flow. Seeing a matte reference object on every trial might thus affect the observer's internal reference point for bumpiness, and consequently, any (not so robust and possibly nonrigid) specular object would be (down) scaled with respect to that reference point. This is in line with previous reports of scaling between matte and glossy static objects (Nefs, 2008; Todd, Norman, Koenderink, & Kappers, 1997). It is interesting, however, that we observed this "scaling" only for in-depth rotations, whereas no such shift occurs for rotations around the viewing axis. This indicates that the matte-textured object that is rotating in depth is perceived as the bumpiest and the easiest discriminable object in our set of stimuli. One possible explanation for this is that this condition is the only one in which both first-order information (slant, conveyed through motion) and second-order information (curvature conveyed through texture compression) are available.

Because the slopes of specular and mixed materials are never significantly larger than those of matte materials, one may be tempted to conclude that even under the worst conditions (i.e., viewing axis rotations), specular flow can never contribute anything additional that is useful in estimating 3D shape.[5] However, the goal of this study was to assess changes in discrimination sensitivity in perceived bumpiness and not the magnitude differences of perceived bumpiness as a function of surface reflection. Our data clearly show that discrimination sensitivity for specular objects under all conditions remains the same.

Integration of specular and matte-textured flow cues in dynamic scenes

In the Introduction, we raised the question whether specular flow provides information that supplements shading and texture information of diffusely reflecting objects. For mixed materials (i.e., weighted combinations of matte-textured and specular components), we found that specular flow information completely overruled motion parallax information. This is quite

Downloaded from jov.arvojournals.org on 07/24/2020

interesting because, theoretically, the combination of both cues could make shape estimations more robust. Disagreement between shading/texture cues and specular highlight geometry can have quite dramatic effects on the perception of surface material (Anderson & Kim, 2009; Kim, Marlow, & Anderson, 2011), suggesting that these cues are combined by the visual system. However, in our experiment, we find that this is not the case. Instead, specular motion appears to capture all shape estimation resources, maybe because of the higher motion energy it produces, for example, along parabolic lines (Adato & Ben-Shahar, 2011). We do not know, however, whether this pattern would persist if specular flow and motion parallax cues were set into stronger conflict. In the specular stereo literature, we have seen demonstrations of how the disparity of specular highlights affects perceived 3D shape (Blake & Bulthoff, 1990; Muryy, Welchman, Blake, & Fleming, 2013). In a potential cue-conflict experiment, we could explore to what extend specular motion could override the 3D shape information generated by matte-textured optic flow.

Conclusion

Motion information generated by moving specularities across a surface is used by human observers when judging the bumpiness of 3D shapes. In the presence of specular motion, observers tend to not rely on the motion parallax information generated by the matte-textured surface reflectance component. This study further highlights how 3D shape, surface material, and object motion interact in dynamic scenes.

Keywords: surface reflectance and 3D shape, *structure from motion*, *shape from specular flow*

Acknowledgments

DND, PB, and KD were supported by an EU Marie Curie Initial Training Network "PRISM" (FP7 - PEOPLE-2012-ITN, Grant Agreement 316746). DND was also supported by a TUBITAK BIDEB 2232 Postdoctoral Reintegration Fellowship (21514107-232.01-9150) and KD by a grant of the Scientific and Technological Research council of Turkey (TUBITAK 1001 Grant 112K069) and a Turkish Academy of Sciences (TUBA) GEBIP Award for Young Scientists. KD was also supported by a Sofja Kovalevskaja Award from the Alexander von Humboldt Foundation, endowed by the German Ministry of Education. Support for OBS was provided in part by the Israel Science Foundation (ISF individual grant no. 259/12 and BIKURA grant 1274/11) and by the Frankel Fund, the ABC Robotics initiative, and the Zlotowski Center for Neuroscience at Ben-Gurion University.

Commercial relationships: none.
Corresponding author: Dicle N. Dövencioğlu.
Email: dicle@bilkent.edu.tr.
Address: Otto Behaghel Strasse 10F, Justus-Liebig-University Giessen, Giessen, 35394, Germany.

Footnotes

[1] This type of manipulation created a very visually different light probe but kept the luminance histogram and spatial frequency distribution of the light probe intact.

[2] Repetitions could vary between observers but were at least 15.

[3] Note that 'light probe' was not a factor in the analysis.

[4] While specular objects appeared overall less rigid than matte-textured ones, there was no systematic relationship to their perceived bumpiness (See Supplementary Material 1).

[5] I.e., a floor effect, where the performance is best predicted merely by the compression of texture patterns in static versions of the stimuli.

References

Adato, Y., & Ben-Shahar, O. (2011). Specular flow and shape in one shot:. In J. Hoey, S. McKenna, & E. Trucco (Eds.), *Proceedings of the British machine vision conference* (pp. 24.1–24.11). Durham, UK: BMVA Press, http://dx.doi.org/10.5244.C.25.24.

Adato, Y., Vasilyev, Y., Zickler, T., & Ben-Shahar, O. (2010). Shape from specular flow. *IEEE Transactions Pattern Analysis and Machine Intelligence*, *32*, 2054–2070.

Anderson, B. L., & Kim, J. (2009). Image statistics do not explain the perception of gloss and lightness. *Journal of Vision, 9(11)*:10, 1–17, doi:10.1167/9.11.10. [PubMed] [Article]

Barrow, H. G., & Tenenbaum, J. M. (1981). Interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence*, *17*, 75–116.

Blake, A., & Bulthoff, H. (1990). Does the brain know the physics of specular reflection? *Nature*, *343*, 165–168.

Blake, A., & Bulthoff, H. (1991). Shape from specularities: Computation and psychophysics. *Philo-*

Downloaded from jov.arvojournals.org on 07/24/2020

sophical Transactions of the Royal Society of London. Series B: Biological Sciences, *331*, 237–252.

Brainard, D. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

Debevec, P. (2002). Image-based lighting. *IEEE Computer Graphics and Applications*, *22*(2), 26–34, doi: 10.1109/38.988744.

Doerschner, K., Fleming, R. W., Yilmaz, O., Schrater, P. R., Hartung, B., & Kersten, D. (2011). Visual motion and the perception of surface material. *Current Biology*, *21*, 2010–2016.

Doerschner, K., Kersten, D., & Schrater, P. R. (2011). Rapid classification of specular and diffuse reflection from image velocities. *Pattern Recognition*, *44*, 1874–1884.

Doerschner, K., Yilmaz, O., Kucukoglu, G., & Fleming, R. W. (2013). Effects of surface reflectance and 3d shape on perceived rotation axis. *Journal of Vision*, *13*(11):8, 1–23, doi:10.1167/13.11.8. [PubMed] [Article]

Dövencioğlu, D. N., Wijntjes, M. W., Ben-Shahar, O., & Doerschner, K. (2015). Effects of surface reflectance on local second order shape estimation in dynamic scenes. *Vision Research*, *115*, 218–230.

Fleming, R. W., Torralba, A., & Adelson, E. H. (2004). Specular reflections and the perception of shape. *Journal of Vision*, *4*(9):10, 798–820, doi:10.1167/4.9.10. [PubMed] [Article]

Hartung, B., & Kersten, D. (2002). Distinguishing shiny from matte. *Journal of Vision*, *2*(7):551, doi: 10.1167/2.7.551. [Abstract]

Hurlbert, A., Cumming, B., & Parker, A. (1991). Recognition and perceptual use of specular reflections. *Investigative Ophthalmology & Visual Science*, *32*, 105.

IBM Corp. (2010). *IBM SPSS Statistics for Mac OSX*. Version 19.0. Armonk, NY: IBM Corp.

Jain, A., & Zaidi, Q. (2011). Discerning nonrigid 3d shapes from motion cues. *Proceedings of the National Academy of Sciences, USA*, *108*, 1663.

Karsch, K., Liao, Z., Rock, J., Barron, J. T., & Hoiem, D. (2013). Boundary cues for 3D object shape recovery. In: *2013 IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2163–2170). New York: IEEE. Retrieved from http://ieeexplore.ieee.org/document/6619125/

Kim, J., Marlow, P., & Anderson, B. L. (2011). The perception of gloss depends on highlight congruence with surface shading. *Journal of Vision*, *11*(9):4, 1–19, doi:10.1167/11.9.4. [PubMed] [Article]

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? *Perception*, *36*(ECVP Abstract Supplement 14).

Koenderink, J., & Van Doorn, A. (1980). Photometric invariants related to solid shape. *Optica Acta*, *27*, 981–996.

MATLAB (2014). Version 8.4.0 (R2014b). Natick, MA: MathWorks Inc.

Mazzarella, J., Cholewiak, S., Phillips, F., & Fleming, R. (2014). Limits on the estimation of shape from specular surfaces. *Journal of Vision*, *14*(10):721, doi: 10.1167/14.10.721. [Abstract]

Mood, A. M., Graybill, F. A., & Boes, D. C. (1974). *Introduction to the theory of statistics* (3rd ed.). Singapore: McGraw-Hill.

Muryy, A. A., Welchman, A. E., Blake, A., & Fleming, R. W. (2013). Specular reflections and the estimation of shape from binocular disparity. *Proceedings of the National Academy of Sciences, USA*, *110*, 2413–2418.

Nefs, H. T. (2008). Three-dimensional object shape from shading and contour disparities. *Journal of Vision*, *8*(11):11, 1–16, doi:10.1167/8.11.11. [PubMed] [Article]

Norman, J., Phillips, F., Cheeseman, J., Thomason, K. E., Ronning, C., Behari, K., ... Lamirande, D. (2016). Perceiving object shape from specular highlight deformation, boundary contour deformation, and active haptic manipulation. *PloS One*, *11*, e0149058.

Pelli, D. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.

Ramachandran, V. S. (1988). Perception of shape from shading. *Nature*, *331*, 163–166.

Sakano, Y., & Ando, H. (2008). Effects of self-motion and stereo viewing on perceived glossiness. *Journal of Vision*, *10*(9): 1–14, doi:10.1167/10.9.15. [PubMed] [Article].

Savarese, S., Fei-Fei, L., & Perona, P. (2004). What do reflections tell us about the shape of a mirror? In: *Proceedings of the 1st Symposium on Applied Perception in Graphics and visualization* (pp. 115–118). New York: ACM.

Schofield, A. J., Rock, P. B., & Georgeson, M. A. (2011). Sun and sky: Does human vision assume a mixture of point and diffuse illumination when interpreting shape-from-shading? *Vision Research*, *51*, 2317–2330.

Todd, J. T., Norman, J. F., Koenderink, J. J., & Kappers, A. M. (1997). Effects of texture, illumination, and surface reflectance on stereoscopic shape perception. *Perception*, *26*, 807–822.

Todorović, D. (2014). How shape from contours affects shape from shading. *Vision Research*, *103*, 1–10.

Vergne, R., Ciaudo, G., & Barla, P. (2014). Gratin: A

Downloaded from jov.arvojournals.org on 07/24/2020

programmable node-based system for gpu-friendly applications. Retrieved from http://gratin.gforge.inria.fr/

Vota, R. M., Dovencioglu, D., Ben-Shahar, O., Doerschner, K., & Wijntjes, M. (2015). The contribution of motion to shape-from-specularities. *Perception*, *44*, 352–352.

Wagemans, J., Van Doorn, A. J., & Koenderink, J. J. (2010). The shading cue in context. *i-Perception*, *1*, 159–177.

Wendt, G., Faul, F., Ekroll, V., & Mausfeld, R. (2010). Disparity, motion, and color information improve gloss constancy performance. *Journal of Vision*, *10*(9):7, 1–17, doi:10.1167/10.9.7. [PubMed] [Article]

Yilmaz, O., & Doerschner, K. (2014). Detection and localization of specular surfaces using image motion cues. *Machine Vision and Applications*, *25*(5): 1333–1349.

Appendix A

Reflectance mapping

The "stuck-on" effect for matte-textured objects is obtained by choosing a special combination of material (BRDF) f_r and lighting (incoming radiance) L_i:

$$f_r(x,\ell,v) = \frac{1-\alpha}{\pi}\,\mathrm{Env}(r_{v_0}(x)) \;+\; \alpha\,\delta(r_v(x),\ell)$$
$$L_i(\ell) \;=\; (1-\alpha) \;+\; \alpha\,\mathrm{Env}(\ell),$$

where x is a surface point; ℓ and v are the light and view directions, respectively; Env is the distant environment lighting, $r_v(x)$ is the reflection of v about the normal $n(x)$; and $\delta(r, \ell)$ is a dirac that equals 1 iff $r = \ell$. When $\alpha = 0$, the material f_r becomes pure diffuse with albedo variations controlled by $\mathrm{Env}(r_{v_0}(x))$ (matte-textured), where v_0 corresponds to the initial viewing direction and the lighting L_i is equal to 1. When $\alpha = 1$, the material becomes a pure mirror, and the specular object reflects light by sampling the environment Env in the direction $r_v(x)$.

In the general case, the BRDF and incoming radiance are used in the rendering equation to obtain the reflected radiance L_r:

$$L_r(x,v) = \int_\Omega f_r(x,\ell,v)\,L_i(\ell)\,\ell\cdot n(x)\,d\ell.$$

It is easy to show that for $v = v_0$, we get $L_r(x, v_0) = Env(r_{v_0}(x))$ both for $\alpha = 0$ and $\alpha = 1$. In other words, the rendered images are identical for the pure matte and specular cases in the initial viewpoint configuration. This is the basis of the stuck-on effect used for the matte-textured objects here: when $\alpha = 0$, the object appears as made of a mirror-like material in a static image, but reflections stick to its surface when the camera moves away from its initial position (i.e., when $v \neq v_0$).

Unfortunately, the stuck-on effect does not occur with $0 < \alpha < 1$, because it can be shown that in the general case, we have

$$L_r(x,v_0) = \alpha(1-\alpha) + (\alpha(1-\alpha)(\mathrm{Diff}-2)+1)\mathrm{Env}(r_{v_0}(x)),$$

where $\mathrm{Diff} = \frac{1}{\pi}\int_\Omega \mathrm{Env}(\ell)\,\ell\cdot n(x)\,d\ell$ is the diffuse-filtered environment, which we assume to be constant over the object. This is a valid assumption only when Env is made of a stationary noise pattern.

We also rendered objects with a material in between where $\alpha = 0.9$. We used *tone mapping* to make sure that the mixed material objects appeared identical to other materials if they were to be viewed static. To obtain a stuck-on effect for any value of a, we introduce a specific-purpose tone-mapping operator:

$$T_\alpha(L) = \frac{L-\alpha(1-\alpha)}{1+\alpha(1-\alpha)(\mathrm{Diff}-2)}.$$

It can be verified that we now have $T_\alpha(L_r(x,v_0)) = Env(r_{v_o}(x))$, for all values of α as required. Note in particular that $T_0(L) = T_1(L) = L$; hence, the tone-mapping operator has no effect in the matte-textured ($T_0(L) = L$) and specular cases ($T_1(L) = L$).

Additional details about the rendering can be found in Dövencioğlu et al. (2015).

Downloaded from jov.arvojournals.org on 07/24/2020