



Rapid classification of specular and diffuse reflection from image velocities

K. Doerschner^{a,b,*}, D. Kersten^c, P.R. Schrater^{c,d}

^a National Research Center for Magnetic Resonance (UMRAM), Bilkent Cyberpark, C-Blok Kat 2, 06800 Ankara, Turkey

^b Department of Psychology, Bilkent University, 06800 Ankara, Turkey

^c Department of Psychology, University of Minnesota, 75 East River Road, Minneapolis, MN 55455, USA

^d Department of Computer Science and Engineering, University of Minnesota, 200 Union Street SE, Minneapolis, MN 55455, USA

ARTICLE INFO

Available online 15 September 2010

Keywords:

Specular flow
Rapid surface reflectance classification
Velocity histogram
Material perception
Spatio-temporal filtering

ABSTRACT

We propose a method for rapidly classifying surface reflectance directly from the output of spatio-temporal filters applied to an image sequence of rotating objects. Using image data from only a single frame, we compute histograms of image velocities and classify these as being generated by a specular or a diffusely reflecting object. Exploiting characteristics of material-specific image velocities we show that our classification approach can predict the reflectance of novel 3D objects, as well as human perception.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

Identifying the surface reflectance of an object is a fundamental problem in vision. Reflectance provides important information about the object's material and identity and, given known reflectance, algorithms for shape reconstruction exist for both, diffuse [1–4] and specular surfaces [5]. Knowing surface reflectance is even more important for interpreting image motion, because of the strong differences in the motion fields generated by specular and diffuse surfaces.

Previous work on diffuse vs. specular reflectance classification has relied on specific assumptions and conditions, including specialized lighting, assumptions about the spectral BRDF, or knowledge of camera motion. The goal of this paper is to show how reflectance can be rapidly classified based on the statistical differences between the image motion generated by moving diffuse and specular surfaces¹ without these restrictive assumptions.

Broadly, the past research can be divided into two categories, one has treated specularities as an undesirable image artifact and thus focused on their removal, while the other has exploited specular reflection as an additional source of information to 3D shape.

Highlight removal: In order to extract 3D shape from an image most machine vision algorithms use the intensity distribution across an object. While this approach works well for matte

surfaces, specular highlights pose a serious problem for these methods, as the region around a highlight entails abrupt and large changes in image intensity—as opposed to the smoothly varying intensity profile (shading) of a matte or Lambertian object. Therefore, material classification into diffuse and specular reflectances has been a by-product those approaches that aim to remove specular highlights.

Within this approach one group of research has employed the Dichromatic Reflection Model developed by Shafer [6] which approximates the light reflected by a surface point as a linear combination of diffuse and specular components, where both components are assumed to have different spectral distributions. Using this model Klinker [7] analyzes the spectral histogram of single images of colored objects and combines it with a sensor model to separate pixels belonging to highlights from those belonging to diffuse object color—each of which form separate spectral lines in the dichromatic plane. Bajcsy et al. [8] extended this work by also segmenting highlights arising from interreflections between objects. Both approaches assume the diffuse color across an object to be uniform, thus would not produce correct results for textured surfaces. Tan et al. [9] proposes a pixel based techniques which overcomes this limitation by comparing the chromaticities of only two neighboring pixels to detect color discontinuities. However, their iterative approach takes a substantial amount of time. A more efficient algorithm has been developed by Chung et al. [10]. They proposed an integrative feature based technique that classifies boundary pixels as either belonging to a highlight or not, without relying on the color signature of the diffuse and specular reflectance components.

An alternative to the approaches relying on the Dichromatic Reflection Model approach [6] has been to identify specular highlights by taking advantage of the fact that the light reflected by specular regions is highly polarized while that reflected by

* Corresponding author at: National Research Center for Magnetic Resonance (UMRAM), Bilkent Cyberpark, C-Blok Kat 2, 06800 Ankara, Turkey.

E-mail addresses: katja@bilkent.edu.tr (K. Doerschner), kersten@umn.edu (D. Kersten), schrater@umn.edu (P.R. Schrater).

¹ The analysis is limited to purely specular and purely diffusely reflecting surfaces.

diffuse body color is not. The polarization-based algorithm proposed by Wolff [11], uses the Fresnel reflectance model to predict the magnitudes of diffuse and specular polarization components of reflected light. The separation algorithms by Nayar et al. [12] combines color-, and polarization profiles, successfully segmenting highlights with underlying varying diffuse components, and highlights across regions of different reflectance properties (e.g. smooth-rough).

Exploiting specular cues to 3D shape: Several machine vision algorithms, instead of treating specular reflections as a source of noise, take advantage of the cues provided by dense specular flow in the computation of 3D shape. While there has been a substantial amount of work for shape from specularity per se (see [5] for an overview), only a few algorithms involve the segmentation of diffuse and specular reflection. Nayar et al. [13], for example, use a special illumination setup and photometric sampling to determine the Lambertian and specular components of surface reflection. Image intensities corresponding to different light source directions are sampled, and their extraction algorithm determines both, the shape of the object shape, and the proportion of specular and Lambertian reflectance of the surface. Oren and Nayar [14] introduce a caustics-based framework which allows to distinguish between real (surface point or texture element) and virtual features (reflection of a real features by a specular surface). Their algorithm involves tracking of surface features during known camera motion. Real and virtual features are classified according to cluster compactness of the corresponding caustics.² Roth et al. [15] model dense optic flow arising from a surface due to known small camera motion as a probabilistic mixture of diffuse and specular reflection components. Assuming distant illumination they show parametrically how specular flow can be related to 3D geometry.

Material classification: DelPozo and Savareze [16] developed an algorithm which identifies regions of static specular flow, and uses these to classify part of a surface as specular or non-specular. Their approach requires only minimal assumptions about the scene, a single frame and no knowledge about 3D shape. Their approach consists of three steps: first, regions of anisotropic patterns are identified, subsequently complex texon image descriptor which characterizes specular regions and non-specular regions is being build, and lastly a classifier distinguish between those two types of regions. To the best of our knowledge this is the only work in addition to ours aiming to identify surface material per se.

Specular flow and human perception: The above discussed works each rely on specific—and often multiple—assumptions (except [16]), be it a specific reflection model, a specific illumination or sensing setup or known camera trajectory. Evidence from human vision, however, suggests that monocular image motion across a few frames provides sufficient information to classify a surface as diffuse or specular, e.g. [17] showed that static objects with ambiguous apparent reflectance could be unambiguously classified as shiny³ or matte when in motion. Additionally, [18] demonstrated that it is also possible to generate reflectance illusions from motion: under certain conditions, rotating specular objects look matte. (For a demonstration of this effect see http://bilkent.edu.tr/~katja/pr_mov.html Movie 1.) Roth et al. [19] simulated specular flow on a sphere (consisting of random dot elements), however, their simulations were not perceived as shiny in appearance. The question arises: What aspects of specular flow explain both, the rapid material classification and the perceptual errors?

Although specular flow patterns can be quite complex, we will show in this paper that simple statistical measures on *image velocities* can be used to classify moving objects as specular or diffusely reflecting. In contrast to existing work on highlight removal and shape-from-specularity our approach does not require any additional assumptions or conditions. Unlike the algorithm by DelPozo and Savareze [16] our approach does not rely on the computation of complex image features or require the presence of oriented features in the image, hence it is more reliable. Moreover, we can link our proposed simple statistical measurements directly to a theory of specular flow and 3D shape.

The paper is structured as follows: In Section 2 we will explain how surface curvature variability and specular flow are related and make predictions for the velocity distributions of moving diffuse and specular objects. In Section 3 we will describe the implementation of our classification methods, and report experimental results in Section 4. Finally, in Section 5 we will demonstrate that our classifiers can predict human perception.

This work constitutes a novel approach to material classification, relying on simple measures of image velocities only. Our research provides new insights into how 3D shape and surface material are related. Rapid methods for reflectance classification, such as the one proposed here, constitute an important step towards a fully automated vision system.

2. Specular flow

The relative displacement of a specular feature or highlight due to camera or observer motion (or, conversely due to object motion relative to a stationary camera/observer), is negatively related to the magnitude of surface curvature [3,20], i.e. specular features “rush” across low curvature regions and “stick” to points of high curvature. In contrast, all points on a moving diffusely reflective surfaces stick. This suggests that the distribution of *image velocities*,⁴ across a moving object may contain important information about the object’s material, because all specular surfaces with sufficient curvature variation undergoing a generic motion will have both low velocity “sticky” points and high velocity points, while diffusely reflective surfaces will have only “sticky” points. Moreover, except for rotations around the viewing axis, the flow generated by a rigid body motion will have a principle direction of motion.

For example, for an in-depth rotating specular object (Fig. 1A) the distribution of image velocities generated by the specular flow across the object will have regions of relatively high and low magnitude, whose specific range is directly related to the magnitude and range of surface curvatures. As an extreme case, a rotating cube, (0 curvature across sides and positive curvature at the corners) will produce two kinds of image velocities: high ones, opposite to the direction of object rotation (along the sides) and those congruent with object rotation speed and direction (“sticking” to corners). As an object increases in surface curvature homogeneity the resulting range of image velocities will decrease, the extreme end being a rotating specular sphere: it will produce image velocities of magnitude and range 0. This *velocity variability* can be exploited for reflectance classification: high image velocity variability, which can be easily identified from the image velocity histogram, appears to be crucial to induce the spatio-temporal characteristics associated with perceived shininess [18]. Conversely, specular objects with low curvature variability will, when rotated, generate low variability image velocity distributions which are, not surprisingly, not distinct

² Envelope defined by the family of reflection rays produced by the motion of a specular feature [14].

³ Shininess is a perceptual quality of BRDFs with a specular component.

⁴ We define image velocity as the distance traveled per frame by a flow point (specular or diffuse) along the dominant direction of motion. See Sections 2.1 and 3.2.

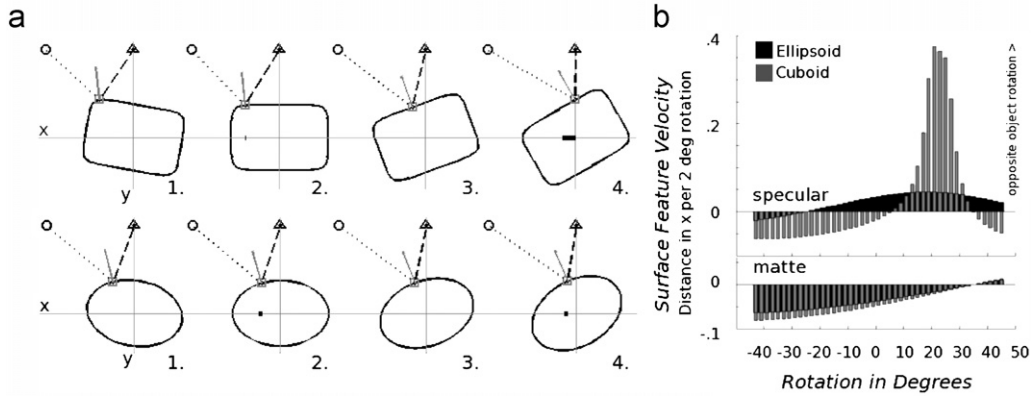


Fig. 1. Specular velocity and curvature variability. (A) Cross-sections through 3D scenes. The position of the 2D camera (triangle) and a point light source (circle) are fixed. We find the surface normal at the point on the object where the specular feature (square) will be visible to the camera. “Specular velocity” is measured as the distance traveled by the specular feature in x (indicated by fat black line) as the object rotates 10° counterclockwise around its origin. Consider the cuboidal cross-section: 1. The specular feature (sf) appears on a high curvature point and “sticks” to this region as the object rotates. 2. The sf moves some distance in the direction of object rotation. 3. The sf appears on a low curvature point. After a 10° rotation the distance that it has traveled, now in opposite the direction of object rotation, has nearly doubled. Compare this to the sf on the ellipsoid. (B) Sf velocities for specular (upper plot) and surface feature velocities for diffusely reflecting (lower plot) objects per 2° rotation. See text for details.

from those generated by diffusely reflecting objects (Fig. 1B). This last observation may account for results by [19]. The simulated specular flow did not look shiny, since a sphere lacks the surface curvature variability crucial to induce the spatio-temporal characteristics associated with perceived shininess.

2.1. Statistical information in flow

Using the statistical properties of flow to classify material properties requires the statistical relationships to be consistent. What conditions are required for optic flow statistics to be reliably different for diffuse and specular surfaces? To answer this question we derive the relationship between object properties and the statistics of optic flow. In particular, we show how the motion of specular points depend critically on the principal curvature of the surface at each point, so that the distribution of curvatures produces the statistics of specular flow. We make this relationship explicit below, and use the resulting equations to discuss the conditions needed for classification accuracy.

Image motion induced by a moving object can be decomposed into two components: image motion due to the change in direction of surface normals (specular flow) and image motion due to the displacement and rotation of surface points (optic flow). Explicit equations for specular flow can be derived assuming orthographic viewing and illumination parametrized by directions on a sphere. The object surface $F(x, y) = (x, y, f(x, y))$ is represented as a function of image coordinates x, y , with $\vec{n}(x, y) = N(\theta, \phi)$ be the surface normal at the surface point $F(x, y)$ with direction (θ, ϕ) and N represents the mapping between spherical and cartesian coordinates. Because the viewing direction is $\vec{v} = (0, 0, 1)$, the mirror direction $\vec{r} = N(2\theta, \phi)$ produces the image point at (x, y) .

When the surface undergoes a rigid body motion $TF(x, y) = RF(x, y) + t$, both surface points and surface normals are transformed, and both induce image motion. For a sufficiently textured surface, optic flow is given by the projection of the motion field:

$$\begin{pmatrix} \frac{dx}{dt} \\ \frac{dy}{dt} \end{pmatrix} = I_{2 \times 3}(-R\dot{R}^T F(x, y) + v_{obj}) \quad (1)$$

where $I_{2 \times 3}$ is the orthographic projection matrix, and $\dot{R}(t) = [\omega \times]_x$ is the cross-product matrix formed from the rotation axis ω . Because translations simply translate the flow under the viewing and

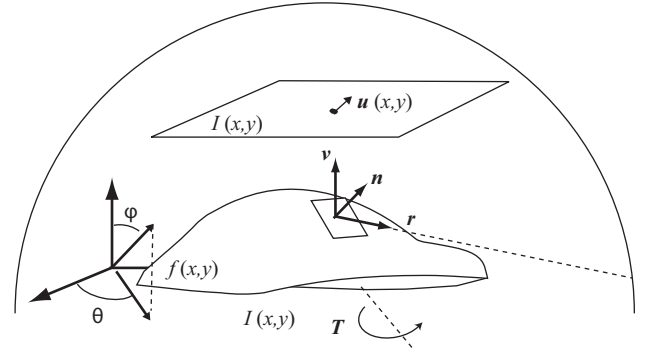


Fig. 2. Analysis of specular flow. A surface $f(x, y)$, reflecting a far-field illumination environment viewed orthographically to produce an image $I(x, y)$, undergoing a rigid body transformation T . The figure has been adapted from [21].

illumination assumptions, we focus on rotations. The relationship shows that the motion field depends on the distribution of depths $F(x, y)$. The other component of image motion is specular flow, caused by the motion of surface normals.

For a specular surface, the change in surface normals caused by rotation around an axis ω induces a specular flow field. The equations simplify if we consider the vector field on a sphere induced by a rotation around expressed in spherical coordinates $\omega = \gamma(\cos\theta_0\sin\phi_0, \sin\theta_0\sin\phi_0, \cos\phi_0)$, where γ is the magnitude of the rotation and θ_0, ϕ_0 encode the direction. Let α, β represent the angular coordinates of the reflection vectors corresponding to each normal. In this representation, $d\alpha/dt$ and $d\beta/dt$ are differential changes to the direction of the reflection vectors induced by the motion (Fig. 2).

A rigid rotation induces a vector field of changes to the reflection vectors:

$$\begin{pmatrix} \frac{\partial\alpha}{\partial t} \\ \frac{\partial\beta}{\partial t} \end{pmatrix} = \gamma \begin{pmatrix} \cos\beta_0 - \sin\beta_0 \cos(\alpha - \alpha_0) \cot\beta \\ \sin\beta_0 \sin(\alpha_0 - \alpha) \end{pmatrix} \quad (2)$$

Specular flow due the change in reflection vectors was derived in [22]. Here we rewrite it in terms of surface curvature using the shape operator, the matrix whose eigenvectors and eigenvalues form the principle curvatures. The shape operator stems from

approximating a surface locally around a point $p=(x, y, z)$ as a quadratic patch:

$$(x, y, f(x, y)) \approx p + x\vec{e}_x + y\vec{e}_y + (x, y)S \begin{pmatrix} x \\ y \end{pmatrix} \vec{e}_z \quad (3)$$

where the matrix S is the shape operator:

$$S = \frac{1}{(1 + |\nabla f|^2)^{3/2}} \begin{pmatrix} 1 + f_y^2 & -f_x f_y \\ -f_x f_y & 1 + f_x^2 \end{pmatrix} \begin{pmatrix} f_{xx} & f_{xy} \\ f_{xy} & f_{yy} \end{pmatrix} \quad (4)$$

Then the specular flow is given by

$$\begin{pmatrix} \frac{dx}{dt} \\ \frac{dy}{dt} \end{pmatrix} = -S^{-1} \begin{pmatrix} f_x & -f_y \\ f_y & f_x \end{pmatrix} \begin{pmatrix} (1 + |\nabla f|^2) & 0 \\ 0 & |\nabla f|^2 \end{pmatrix} \begin{pmatrix} \frac{d\alpha}{dt} \\ \frac{d\beta}{dt} \end{pmatrix} \quad (5)$$

While specular flow $(dx/dt, dy/dt)$ is not directly measurable, it generates measurable image motion in terms of an optic flow field whenever the environment map has sufficient contrast and texture variability. We will assume this is true when discussing the relationship between curvature and image motion.

The three matrices in Eq. (5) have the following interpretation, from right to left. The first two express the effect of the orientation of the tangent plane—large gradients means more specular flow. The last matrix is the inverse of the shape operator. We can use this fact to forge a relationship between specular flow and principal curvatures, which are the eigenvalues of the shape operator. Specifically, $S = VDV^{-1}$, where D is a diagonal matrix with the principal curvatures k_1 and k_2 as the entries, and V contains the directions of principal curvature. Because $S^{-1} = V^{-1}D^{-1}V$, which means specular flow is proportional to inverse principal curvatures. For example a large curvature yields a large eigenvalue, and hence produces no specular flow—the image motion at those points is only due to optic flow. In contrast, small curvatures produce exceptionally fast specular flow. The direction of the flow is determined by the projection of the direction of motion onto the direction of principle curvature and by the sign of the curvature—convex produces motion away from the

surface rotation, concave towards. Parabolic points are especially simple because they have only one non-zero eigenvalue. The matrix S is singular for these points and the specular flow is parallel to the principle curvature. From this analysis we see that the distribution of principal curvatures has a direct effect on the characteristics of specular flow. Fig. 3 panel 2 shows the principle curvature field and simulated flow (panel 3) for a surface (panel 1) with a simulated rotation around the y -axis in the surface plane. Note the relationship between the curvature field and the simulated flow. Locally, the norm of a given flow vector can be a good indicator for the image velocity of the corresponding image point. Large flow will lead to high image velocities, conversely small flow will cause relatively low image velocities. Since we are interested in the global distribution of flow, i.e. across the object, we first need to estimate the dominant direction of image motion and then project each flow vector onto this axis, before computing the corresponding image velocities. Fig. 3 panel 4 shows a 2D density estimate of simulated velocity measurements, for which the bimodal distribution is clearly apparent.

It should be noted that a rotating *planar* specular surface (flat mirror) constitutes a singularity and could not be classified by the proposed method. The specular flow, in this case, would be identical at every image point, hence the distribution of image velocities would contain only a single value. Therefore, a minimal requirement for our algorithm to work is that sufficient surface curvature modulation is present across the object. Exactly what *sufficient* entails in terms of curvature classes (hyperbolic, parabolic, or elliptic) and magnitudes present is the subject of ongoing research in our lab.

3. Implementation

3.1. Algorithm description

To rapidly classify reflectance properties from image velocities our strategy was to (1) estimate velocities from rotating *specular*

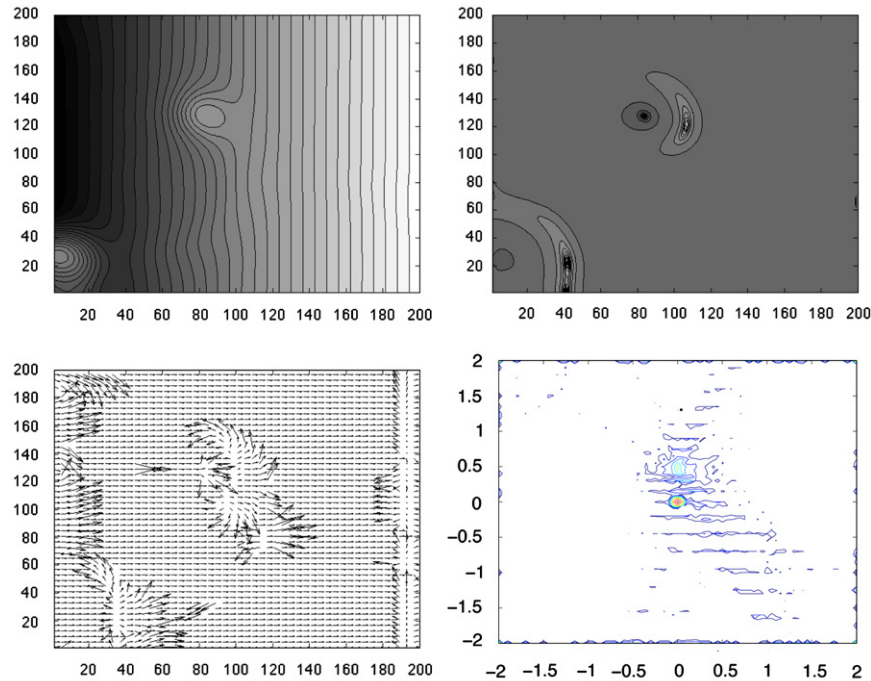


Fig. 3. 3D curvature and specular flow. From left to right starting with the upper row: Panel 1 shows a surface as a contour map. Panel 2 depicts the corresponding magnitude of principle curvatures, darker gray values correspond to concavities, lighter gray values to convexities. Panel 3 shows the theoretical specular flow for a left to right motion, and Panel 4 a 2D density estimate of simulated velocity measurement. See text for details.

objects using spatio-temporal filters, (2) find the principal direction of motion, and (3) classify the velocity histogram in that principal direction using three different approaches: parametric, and non-parametric density estimation, as well as non-negative matrix factorization. We chose to classify movies on the basis of histogram velocities because we expected the velocity signature of specular or matte (appearing) reflectances to be largely object (identity) invariant (but see Section 2 for the special role of 3D curvature). Furthermore, by focusing on the principal direction of motion we achieve object motion invariance.

3.2. Spatio-temporal filtering

We filtered image sequences by directionally selective filters G_2 (second derivative of a 3D Gaussian) and H_2 (and its Hilbert transform) at orientations $(\alpha, \beta, \gamma)_i$ [23]:

$$f^\Omega(x, y, z) = G(r)Q_N(x') \quad (6)$$

are the even and odd filters formed by a n th order polynomial $Q_N(x')$ ⁵ times a separable windowing function $G(r)$ (e.g. a Gaussian-like function), both of which are assumed to be rotationally symmetric. \mathbf{R} is the transformation that these functions are rotated by such that their axis of symmetry points along the direction of cosines α, β and γ . We estimated velocity vectors⁶ from the filter coefficients using the max-steering method of Simoncelli [24]. Subsequent analysis of these velocities was restricted to include samples only from *within object boundaries* in order to avoid contamination with boundary motion. Velocity vectors were sampled from a grid indicated by the colored dots in Fig. 5C.

3.3. Dominant direction of motion

We performed principle components analysis on image velocities to estimate the dominant direction of motion for a given movie frame. Image velocities were projected onto this direction vector.

3.4. Parametric and non-parametric density estimation

To develop statistical classifiers for reflectivity we estimated the conditional probabilities of the projected velocities for both diffuse and specular objects. To verify our results did not depend on the details of a specific density estimation learning procedure, we used three different density learning approaches. The three classification algorithms are described below.

3.4.1. Cross-entropy density estimation

1. Compute histogram estimates of the conditional densities of velocity given shiny.
2. Compute histogram estimates of the conditional densities of velocity given matte.
3. Compute likelihood of a sample given the shiny density.
4. Compute likelihood of a sample given the matte density.
5. Take likelihood ratio and compare against a threshold.

Histogram densities were estimated with a generalized cross-entropy density estimator [25] that uses a gaussian kernel and data-driven bandwidth selection. To classify a given movie frame into shiny or matte we used histogram estimates of the conditional densities of velocity ξ given shiny S , $P(\xi|S)$, and

matte M , $P(\xi|M)$, from image sequences judged shiny and matte in [18] (also see Section 5). A sample velocity ξ' from a test image sequence was classified by comparing the likelihood ratio $P(\xi'|S)/P(\xi'|M)$ against a threshold k .⁷ Note, that we also used the value of the likelihood ratio as a graded material measure for the data set. Graded measures are particularly useful for comparisons to human perception, as discussed below.

3.4.2. Mixture of Gaussians

1. Fit a mixture of Gaussians with two components [26] to a sample frame of each movie.
2. Compute index of bimodality (*velocity contrast*) for each sample frame velocity contrast of each sample.
3. If index > 1 ; sample=specular; else sample=matte.

To confirm that the shape of a given histogram was indeed driven by “diagnostic” (high and low curvature) regions we fitted a mixture of Gaussians with two components [26] to frames of each movie. Given the analysis in Section 2.1 we reasoned that a two-component model would best capture the bimodal nature of specular reflectance velocity distributions. If the estimated Gaussian distributions would significantly overlap it would indicate the absence of high and low velocity regions, hence be indicative for diffuse reflectance. From the two estimated Gaussian means (μ_1, μ_2) we compute the velocity contrast of the sample

$$C_b = \frac{|\mu_1 - \mu_2|}{\max(\sigma_1, \sigma_2) * 2} \quad (7)$$

which is derived from the common index of bimodality, i.e. normal distribution means need to be separated by at least twice the common standard deviation. Instead of the common-, we used the larger standard deviation, which leads to a more stringent criterium of bi-modality and we multiplied this value by 2 for cosmetic purposes such that our cutoff value would be 1. All 36 movies were analyzed, frames were chosen such that the orientation of the superellipsoids were approximately the same. If $C_b > 1$, i.e. if the distribution of image passes the criterion of bimodality, the sample is classified as specular, else as matte. The value of C_b also forms a graded surface material measure.

We further computed the posterior probability of each pixel given either Gaussian distribution. Pixel classifications are illustrated by mapping color coded velocity samples back onto the frame they were taken from.

3.4.3. Mixture of histograms using non-negative matrix factorization

1. Factorize velocity distributions using non-negative matrix factorization.
2. Compute shininess criterion by taking a weighted ratio of specular and matte components.
3. If index > 1 ; sample=specular; else sample=matte.

To smooth the likelihoods and form a low-dimensional representation for the densities, we factorized the velocity distributions using convolutive non-negative matrix factorization (NNMF) [27]. We preserved three components based on an initial estimate that three components account for as much as 97% of the approximation error (see fig.). Because the histogram of a test sequence can be represented as a weighted combination of the three components, these weights can be used to represent the velocity

⁵ $x' = \alpha x + \beta y + \gamma z$.

⁶ These were indicating both, direction and magnitude of the sample.

⁷ k was obtained by a bootstrapping procedure used to constrain the false alarm rate to 5%.

distributions of novel objects. To estimate the weights for a novel sequence, we maximized the likelihood of the total sample evaluated on the components with respect to the weights. The best fitting weight values were used to classify a sample as shiny or matte. A very simple shininess criterion can be computed by taking the ratio of the weights of the two “specular components” and the weight of the “matte component”, e.g. $C_w = 1/2(w_{f1} + w_{f3})/w_{f2}$, with values larger than 1 being classified as specular (see Fig. 7B).

3.5. Movies

The test set consisted of 36 movies (6 shapes \times 6 light probes) of rotating specular superellipsoids (http://bilkent.edu.tr/~katja/g_run.html). Objects were constructed according to

$$1 = \left[\left| \frac{x}{r_x} \right|^{2/n_2} + \left| \frac{y}{r_y} \right|^{2/n_2} \right]^{n_2/n_1} + \left| \frac{z}{r_z} \right|^{2/n_1} \quad (8)$$

We set $r_x=1$ and $r_y=r_z=0.64$. Surface curvature was determined by setting n_1, n_2 to: 0.3, 0.5, 0.7, 0.8, 0.9 or 1.0 (Fig. 4). Each object rotated in depth. Its angular speed was adjusted (0.1, 0.35, 0.61, 0.74, 0.87, $1.0^\circ/\text{frame}$) such that the resulting image velocities were in the

range that our filters were sensitive to. Superellipsoids were rendered under six different light probes: two natural (L1 (“grace”), L4 (“uffizi”) from <http://gl.ict.usc.edu/Data/HighResProbes/>), two partially (L2, L5), two fully phase-scrambled (L3, L6) versions of L1 and L3, respectively. For each movie 40 512×512 images were rendered with *Radiance* [28], using projective projection.

4. Experimental results

4.1. Histograms

Fig. 5 illustrates the characteristic changes that the velocity histogram undergoes as the object decreases in surface curvature variability (left to right). Table 1 shows normalized log-likelihood ratios (LLR) for all histograms testing H_0 that a given histogram has been generated by a matte object.

4.2. Mixture of Gaussians pixel classification

Fig. 6 shows that the simple velocity distribution measure was successful in roughly identifying image regions of high (blue

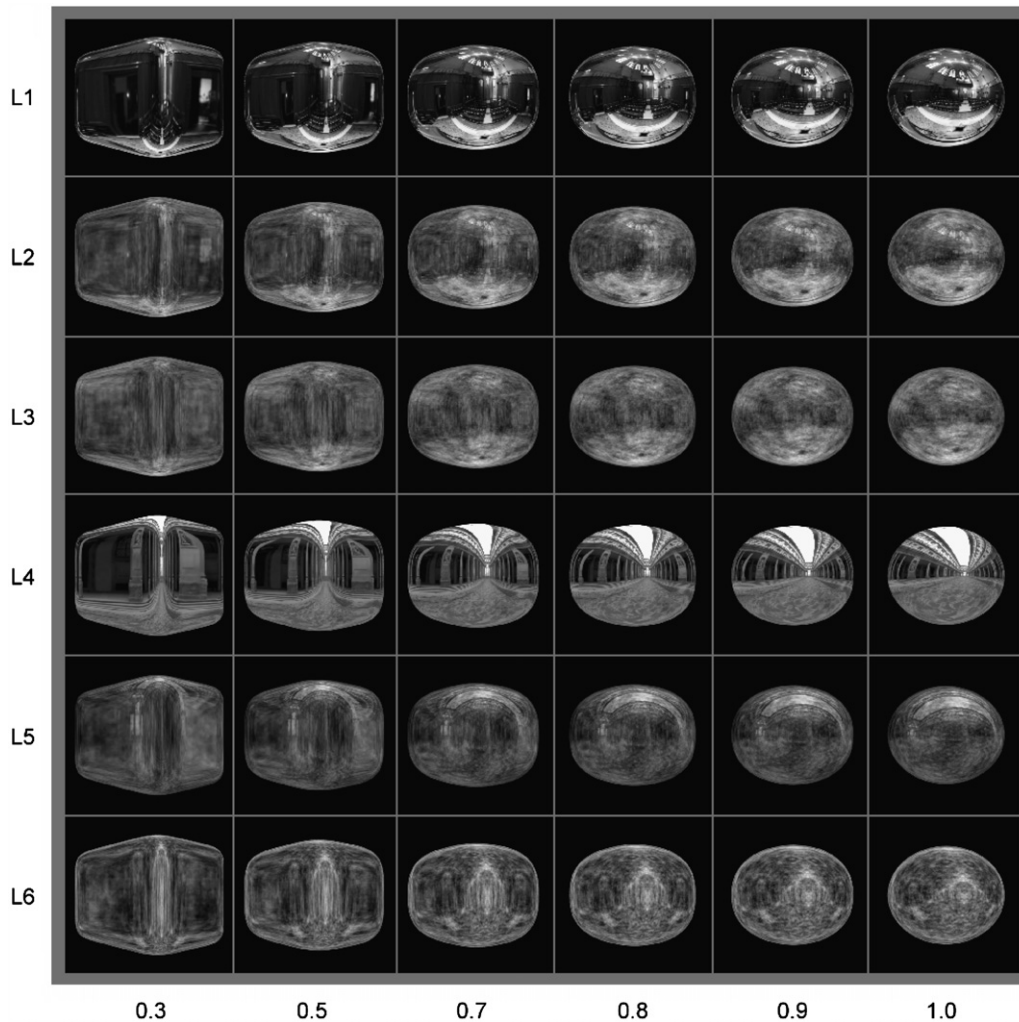


Fig. 4. Renderings. Sample frames from the test set. Panels are labeled according to the object shape and the light probe. Numbers on the x-axis indicate values for n_1, n_2 , as described in Eq. (8). Labels on the y-axis denote the light probe that the object was rendered under (see Section 3.5 for details.) Movies can be viewed at http://bilkent.edu.tr/~katja/g_run.html. As the shape of the object exhibits less surface curvature variability (left to right) we expect the corresponding image velocity distribution to be increasingly more homogeneous (see Section 2).

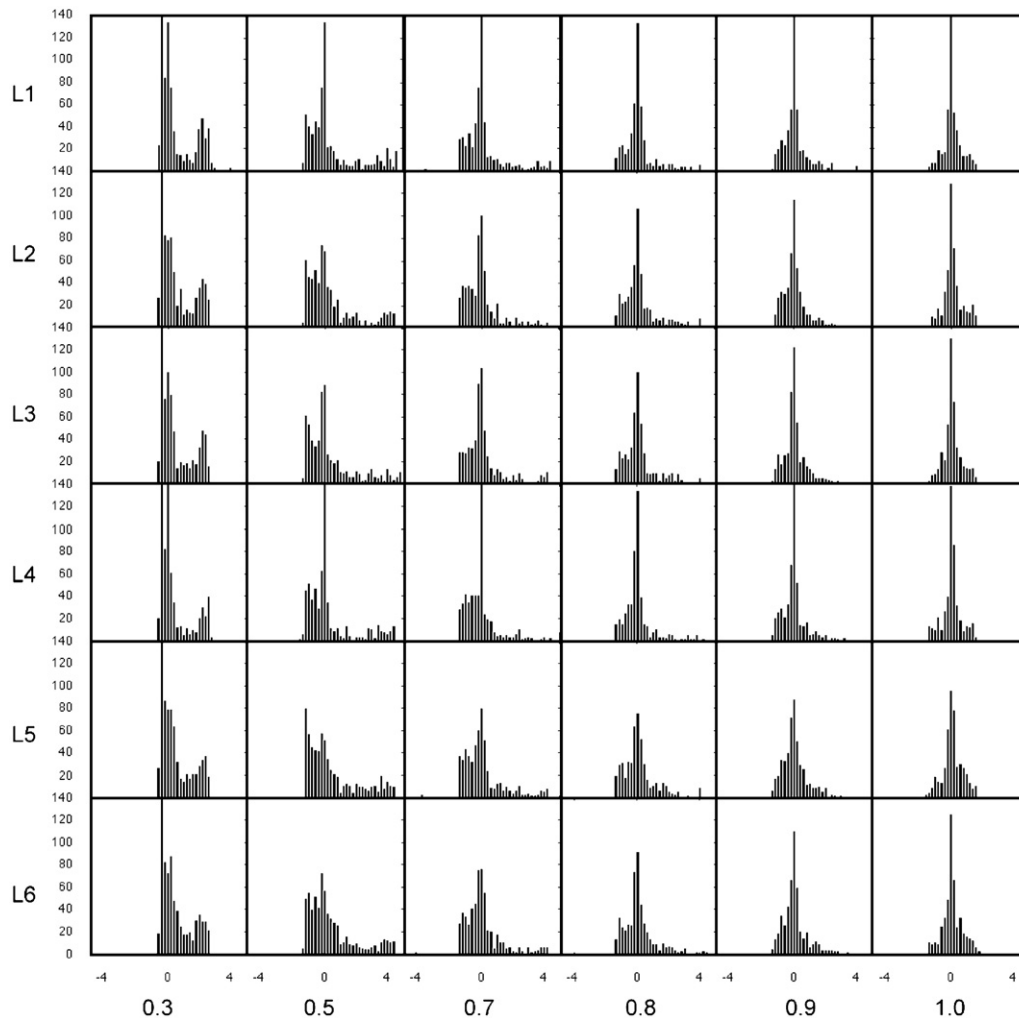


Fig. 5. Histograms. Velocity histograms for all 36 movies. Labels are as in Fig. 4. The bimodal shape of the histograms is evident for specular appearing objects (e.g. $n_1, n_2=0.3$) but not for matte appearing ones (e.g. $n_1, n_2=1.0$).

Table 1
Normalized log-likelihood ratios.

Light probe	Superellipsoid shape coefficient n_1, n_2					
	0.3	0.5	0.7	0.8	0.9	1.0
L1	1.000^T	0.362	0.145	0.153	0.114	0 ^T
L2	0.961	0.362	0.184	0.215	0.139	0.031
L3	0.877	0.365	0.184	0.270	0.103	0.011
L4	0.749	0.267	0.178	0.114	0.114	0.003
L5	0.766	0.476	0.223	0.187	0.142	0.014
L6	0.805	0.368	0.159	0.187	0.148	0.003
Average	0.860	0.367	0.179	0.188	0.127	0.010

Values larger than k ($k=0.16$) (in **bold**) were classified as shiny with a predicted error rate of less than 5%. Training data are indicated by ^T.

pixels) and low (orange pixels) velocities. Purplish colors indicate that the sample could come from either Gaussian distribution. Note, that the distinctiveness of the high and low velocity regions decreases as the amount of the surface curvature variability decreases: in the corresponding two-Gaussian model fit, the two components approach a uni-modal mixture. The measure C_b exploits the bi-modality of specular velocity distributions to classify the material of test sequences (see Table 2).

4.3. Non-negative matrix factorization

The distribution of estimated weights across the stimulus set is shown in Fig. 7A. Ellipsoidal objects' velocity histograms (multiples of 6) tended to have high weights on component 2 (solid triangle) whereas most cube-like objects tended have high weights on components 1(circle) and/or 3(square).

4.4. Objective classification of material of novel 3D objects

To verify that the velocity distribution can be sufficient for objectively classifying material we tested an object with more complex shape variation. We generated 40 frames of a rotating version of the Utah "Teapot". This object was rendered with a diffuse [29] and with a specular reflectance (see Fig. 8). We evaluated the sequence using histograms, mixture of Gaussians, and NMF approaches. Teapots were correctly classified as shiny and matte for all three methods. Histograms: LLR specular and diffusely reflecting teapot were 0.26 (classified as shiny) and 0.008 (classified as matte). Note that the classifier has been trained on specular movies only (superellipsoids), yet the matte object has been classified correctly. Mixture of Gaussians: C_b s for specular and diffusely reflecting teapot were 1.16 (classified as shiny), and 0.87 (classified as matte). NMF: The specular teapot

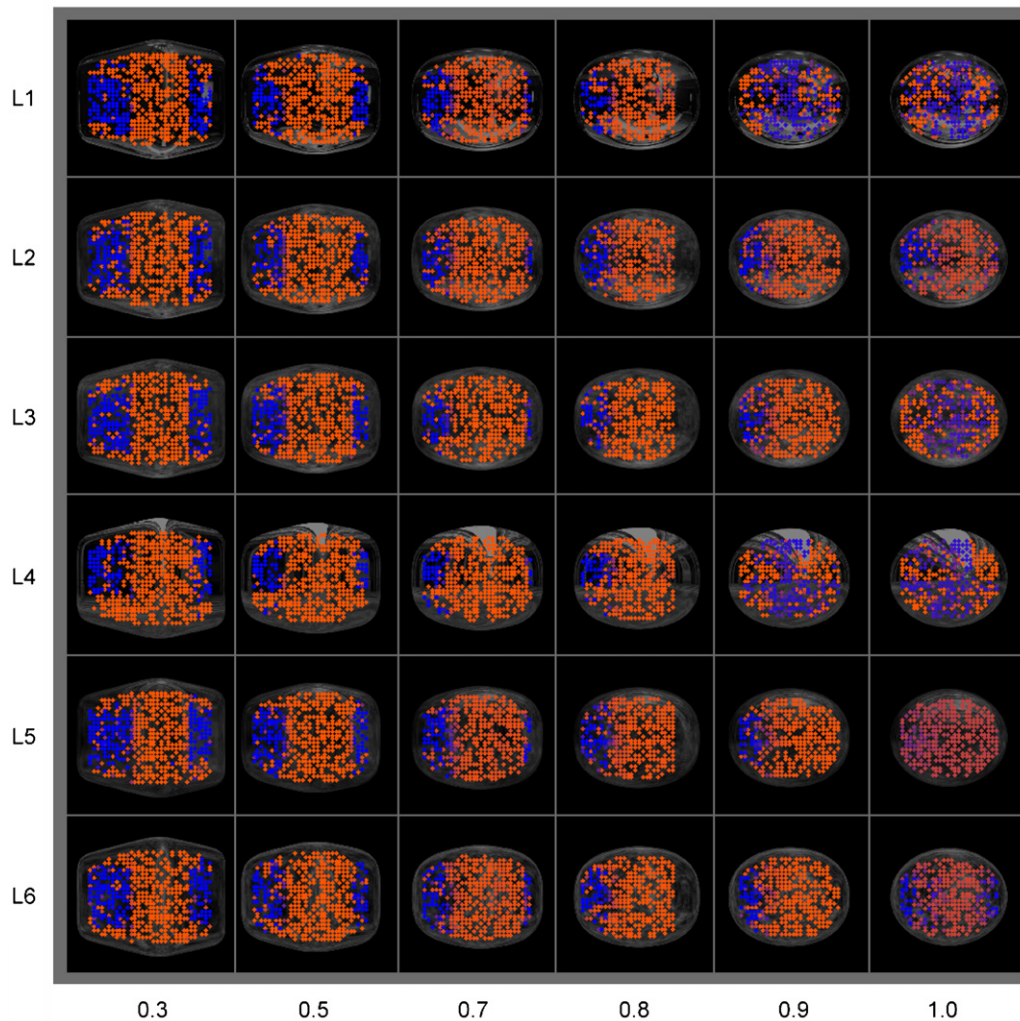


Fig. 6. Pixel classification using mixture of two Gaussians. Each panel shows a movie frame with a subset of classified (fast vs. slow) velocity samples mapped back onto the frame. Labels are as in Fig. 4. Color indicates whether a given sample belongs to a high velocity region (blue) or a low velocity region (orange). This approach works well for bimodal velocity histograms (see corresponding panels in Fig. 5). Purplish color values indicate that a given sample is equally likely to have generated from either Gaussian distribution—which would occur for unimodal histograms.

Table 2
Average C_b .

Light probe	Superellipsoid shape coefficient n_1, n_2					
	0.3	0.5	0.7	0.8	0.9	1.0
Average C_b	1.658	1.4143	0.6824	0.7247	0.4778	0.1341

The average was computed across light probes for superellipsoids with shape coefficients $n_1=n_2$ from 0.3 (cuboidal) to 1 (ellipsoidal). Values > 1 (in bold) indicate that the velocity histogram was classified as bimodal, which could be a rough predictor of material shininess. Compare the relative magnitudes of values to average observer ratings in Table 3.

classified as shiny $C_w=33.2$, and the diffusely reflecting teapot was classified as matte $C_w=0.7954$.

5. Predicting human perception

5.1. Behavioral procedure and data

We used the same set of movies (Section 3.5) in a behavioral experiment with the following modifications: (1) the angular

speed was adjusted to $2.95^\circ/\text{frame}$, (2) a given superellipsoid rotated in depth from -45° to 135° , 0° being the direction to the camera. Frames were assembled into a movie using Quicktime Pro, and set to loop back and forth. The size of the rotating objects at their maximum visible extend was 8.9° visual angle. All movies consisted of 61 frames and were played at 50 frames/s on a G5 workstation Sony GDMC520 (1024×1280) Refresh rate 75 Hz, NVIDIA GeForce 6800 UltraDLL.

Observers were seated in a dark room with their heads stabilized through a chin rest. The viewing distance to the screen was 70 cm. On a given trial observers saw a clip of a rotating superellipsoid either rotating from left to right or right to left (this was achieved by simply playing movies backwards), clips could be re-viewed if desired and the order of presentation of individual clips was randomized. Four observers (three naive, one author KD) indicated via keyboard press on a scale from 1 (matte)–7 (mirror reflection) how shiny a given superellipsoid appeared. Prior to the experiments observers were familiarized with the concepts of shininess.⁸

⁸ Note, that in a separate experiment we measured perceived rigidity for the same set of stimuli and found that the two percepts were significantly correlated

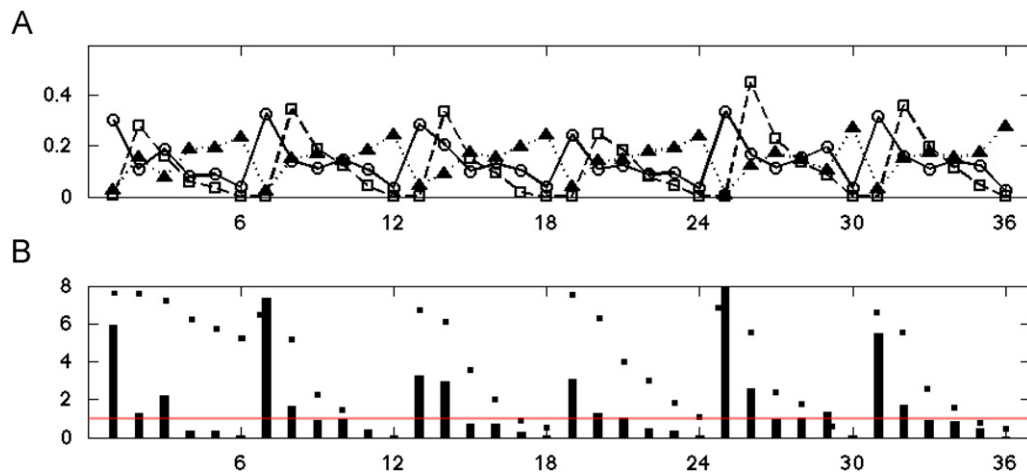


Fig. 7. NMF of velocity histograms. (A) Estimated weights for our test set. (B) Average values of the shininess criterion C_w are 5.4, 1.8, 1.0, 0.7, 0.5, 0.06 for shapes n_1 , $n_2=0.3, 0.5, 0.7, 0.8, 0.9, 1.0$, respectively. The red line represents the cut off value for shininess: $C_w=1$. The black square on top or next to each bar indicates average observer data for the same movie (note, observer values are plotted on a different scale). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

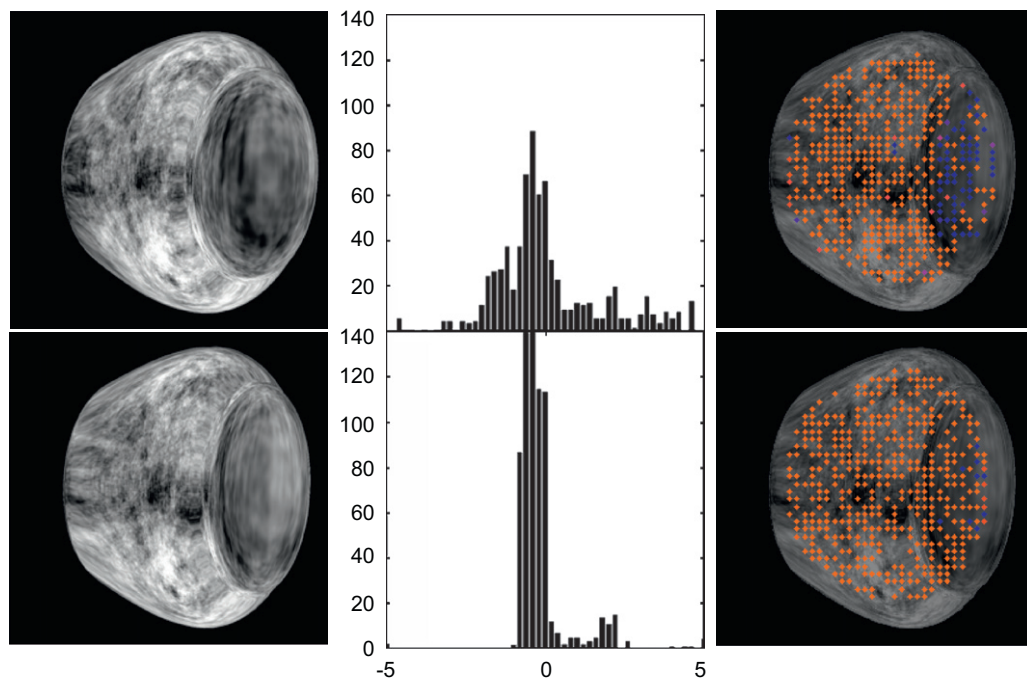


Fig. 8. Classification of novel shape and surface material. Histograms and pixel classification are shown for the specular (upper) and diffusely (lower) reflecting Utah Teapot. Note that our classifiers have been trained on specular movies only, yet the matte object has been classified correctly. This supports our notion that physically matte and apparent matte moving objects share the same flow characteristics.

We transformed all rating data to fall within the interval $[0, 1]$ by $(X_i - 1)/6$, where X_i is the individual rating on a given trial i , and analyzed the data with respect to effects of surface curvature variability.⁹ Fig. 9 shows mean shininess ratings for all shapes and

light probes. Results demonstrate that the more surface curvature variability a rotating object has the shinier it is perceived $F(5,1860)=674.29$ $p < 0.0001$. A subset of average shininess ratings are reported in Table 3.

5.2. Regression results

Regressing normalized LLRs (Table 1) onto normalized observer data (Fig. 9) yielded $R^2=0.45$, $p < 0.00001$. Repeating the analysis with only the most shiny and matte data points yielded $R^2=0.75$, $p=0.0003$. Training data was excluded from the regression.

(footnote continued)

$r=0.69$, $p < 0.00001$, i.e. an object perceived as shiny tended to also be perceived as more rigid. In a separate work we show how objects can be classified according to both, rigidity and reflectivity using optic flow information only [30].

⁹ The second variable in this experiment was the degree of phase scrambling ("realism") of the light probe. However, we will not discuss those results at this point as the primary concern of this study is surface curvature variability. For additional information contact the authors.

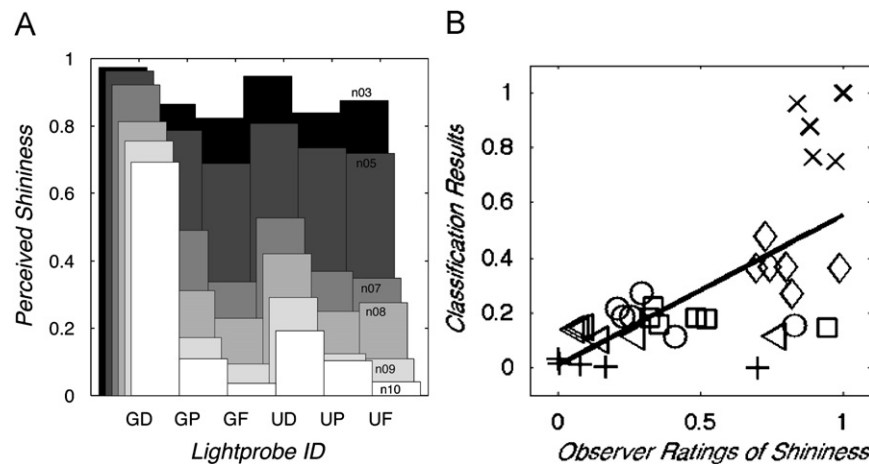


Fig. 9. Behavioral data. (A) Mean shininess ratings for all shapes and light probes. Shape IDs are coded by gray values as indicated. (B) Regression of histogram classifications (LLRs) onto observer data. See text for details.

Table 3
Human shininess ratings.

Light probe	Perceived shininess of shape n_1, n_2					
	0.3	0.5	0.7	0.8	0.9	1.0
L1	0.9740	0.9635	0.9219	0.8125	0.7552	0.6927
L3	0.8229	0.6875	0.3385	0.2292	0.0938	0.0365
Average	0.8872	0.7830	0.4991	0.3837	0.2578	0.1962

Shown are ratings for two light probes (those eliciting on average highest and lowest shininess ratings) as well as the average data (across all light probes and observers). Differences in relative apparent shininess for different light probes is consistent with previous research [31]. In the experiment observers rated apparent shininess of all 36 light probe–shape combinations.

6. Conclusion

We provided a first account of how to rapidly classify surface reflectance from a single frame of object motion, without any assumptions. We show that moving diffusely reflecting, and specular objects with sufficient curvature variability, generate distinct image velocity distributions whose respective characteristics can be captured by simple, invariant statistical measures. Our results account for the misperception of material in [18,19], demonstrating that diffusely reflecting and apparently matte objects, i.e. those that are specular but with insufficient surface curvature variability, share the same velocity histogram characteristics. Thus, we were able to correctly classify a diffusely reflecting object on the basis of a classifier that was trained on a matte-appearing (but physically specular) object.

In future work we will extend our analysis to a velocity region-based approach, as the extent and spatial relationships between high and low velocity regions is likely to be another important diagnostic feature in classifying surface reflectance.

Acknowledgement

KD has been supported by an EC FP7 Marie Curie IRG-239494. Further funding was provided by an NIH RO1 EY015261.

References

- [1] B. Horn, Shape from shading: a method for obtaining the shape of a smooth opaque object from one view, 1970.
- [2] B. Horn, Shape from Shading Information, McGraw-Hill, 1975.
- [3] J. Koenderink, A. Van Doorn, Photometric invariants related to solid shape, *Optica Acta* 27 (1980) 981–996.
- [4] A. Pentland, Shape information from shading: a theory about human perception, *Spatial Vision* 4 (1989) 165–182.
- [5] I. Ihrke, K. Kutulakos, M. Magnor, W. Heidrich, EUROGRAPHICS 2008 STAR—State of The Art Report State of the Art in Transparent and Specular Object Reconstruction, 2008.
- [6] S. Shafer, Using color to separate reflection components, *Color* 10 (1985) 210–218.
- [7] G. Klinker, S. Shafer, T. Kanade, A physical approach to color image understanding, *International Journal of Computer Vision* 4 (1990) 7–38.
- [8] R. Bajcsy, S. Lee, A. Leonardis, Detection of diffuse and specular interface reflections and inter-reflections by color image segmentation, *International Journal of Computer Vision* 17 (1996) 241–272.
- [9] R. Tan, K. Ikeuchi, Separating reflection components of textured surfaces using a single image, *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2005) 178–193.
- [10] Y. Chung, S. Chang, S. Cherng, S. Chen, Dichromatic reflection separation from a single image, *Lecture Notes in Computer Science* 4679 (2007) 225.
- [11] L. Wolff, T. Boulton, Constraining object features using a polarization reflectance model, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13 (1991) 635–657.
- [12] S. Nayar, X. Fang, T. Boulton, Removal of specularities using color and polarization, in: 1993 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Proceedings CVPR'93, 1993, pp. 583–590.
- [13] S. Nayar, K. Ikeuchi, T. Kanade, Determining shape and reflectance of Lambertian, specular, and hybrid surfaces using extended sources, in: *International Workshop on Industrial Applications of Machine Intelligence and Vision*, IEEE, 1989, pp. 169–175.
- [14] M. Oren, S. Nayar, A theory of specular surface geometry, *International Journal of Computer Vision* 24 (1997) 105–124.
- [15] S. Roth, M. Black, Specular Flow and the Recovery of Surface Structure, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, 2006, pp. 1869–1876.
- [16] A. DelPozo, S. Savarese, Detecting specular surfaces on natural images, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition CVPR '07*, 2007, pp. 1–8.
- [17] B. Hartung, D. Kersten, Distinguishing shiny from matte, *Journal of Vision* 2 (2002) 551.
- [18] K. Doerschner, D. Kersten, Perceived rigidity of rotating specular super-ellipsoids under natural and not-so-natural illuminations, *Journal of Vision* 7 (2007) 838.
- [19] S. Roth, F. Domini, M. Black, Specular flow and the perception of surface reflectance, *Journal of Vision* 3 (2003) 413.
- [20] A. Blake, Specular stereo, in: *Proceedings of the International Joint Conference on Artificial Intelligence*, 1985, pp. 973–976.
- [21] Y. Adato, Y. Vasilyev, O. Ben Shihar, T. Zickler, Toward a theory of shape from specular flow, in: *International Conference on Computer Vision (ICCV'07)*, 2007, pp. 1–8.
- [22] Y. Vasilyev, Y. Adato, T. Zickler, O. Ben-Shihar, Dense specular shape from multiple specular flows, in: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2008*, 2008, pp. 1–8.
- [23] K. Derpanis, J. Gryn, Three-dimensional n th derivative of Gaussian separable steerable filters, in: *IEEE International Conference on Image Processing*, 1993.
- [24] E. Context Simoncelli, Distributed analysis and representation of visual motion, Ph.D. Thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, MA, 1993.
- [25] Z. Botev, Z. Botev, A Novel Nonparametric Density Estimator, The University of Queensland, 2006.

- [26] I. Nabney, NETLAB: Algorithms for Pattern Recognition, Springer, 2002.
- [27] P.D. O'Grady, B.A. Pearlmutter, Convolutional non-negative matrix factorisation with a sparseness constraint, in: Proceedings of the IEEE International Workshop on Machine Learning for Signal Processing (MLSP 2006), Maynooth, Ireland, 2006, pp. 427–432.
- [28] G. Larsen, R. Shakespeare, Rendering with Radiance: The Art and Science of Lighting Visualisation, 1998.
- [29] R. Fleming, Rendering sticky reflections with radiance, personal communication, 2007.
- [30] D. Zang, K. Doerschner, P. Schrater, Rapid inference of object rigidity and reflectance using optic flow, in: Proceedings of the 13th International Conference on Computer Analysis of Images and Patterns, Springer-Verlag, 2009, p. 888.
- [31] R. Fleming, R. Dror, E. Adelson, Real-world illumination and the perception of surface reflectance properties, Journal of Vision 3 (2003) 347–368.

Katja Doerschner is Assistant Professor of Psychology at Bilkent University. She received her Ph.D. in Experimental Psychology from New York University. Her research interests include computational and biological vision, and neuroimaging.

Daniel Kersten is Professor of Psychology at the University of Minnesota. He received his B.S. in Mathematics from M.I.T. and Ph.D. in Psychology from the University of Minnesota. His research interests include visual object perception, brain mechanisms of vision, and theories of optimal visual performance.

Paul Schrater is Associate Professor of Psychology and Computer Science and Engineering at the University of Minnesota. He received his Ph.D. from the University of Pennsylvania. His research involves probabilistic approaches to studying vision, motor control, learning and decision making in humans and artificial agents.