Rapid Classification of Surface Reflectance from Image Velocities

Katja Doerschner^{*1}, Dan Kersten², and Paul Schrater^{2,3}

 ¹ Dept. of Psychology, Bilkent University
² Dept. of Psychology, University of Minnesota
³ Dept. of Comp. Science & Eng., University of Minnesota katja@bilkent.edu.tr, {kersten,schrater}@umn.edu

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.

Key words: specular flow, rapid surface reflectance classification, velocity histogram, material perception, spatio-temporal filtering

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 and specular surfaces [1]. However, because of the strong differences in the image motion generated by specular and diffuse surfaces, unknown reflectance is a serious problem for these methods. Previous work on diffuse vs. specular reflectance classification has relied on specific assumptions and conditions, such as the tracking tracking of surface features during known camera motion [2], known surface shape [3], the use of structured lights [4], color [5], or a specific reflectance model [6].

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. [7] showed that static objects with ambiguous apparent reflectance could be unambiguously classified as shiny or matte when in motion. Additionally, [8] demonstrated that it is also possible to generate reflectance illusions from motion: under certain conditions, rotating specular objects look matte (also see [9]). What aspects of specular motion explain both the rapid material classification and the perceptual errors? Although specular motion patterns can

^{*} Corresponding author

be quite complex, we will show that simple statistical measures on image velocities can be used to classify moving objects as specular or diffusely reflecting, without any additional assumptions or conditions. We will demonstrate that these classifiers can predict human perception, as well as the material of novel objects. 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 [10, 11], 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 velocities 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 [8]. Conversely, specular objects with low curvature variability will, when rotated, generate low variability image velocity distributions which are, not surprisingly, not distinct from those generated by diffusely reflecting objects (Fig. 1B).

3 Implementation

General Strategy. To rapidly classify reflectance properties from image velocities our strategy was to 1) estimate velocities from rotating *specular* objects using spatio-temporal filters, 2) find the principal direction of motion, and 3) classify the velocity histogram in that principal direction using 3 different approaches:



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.

parametric and non-parametric density estimation, as well as non-negative matrix factorization. We choose 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.

Spatio-temporal Filtering. We filter image sequences by directionally selective filters G_2 (second derivative of a 3D Gaussian) and H_2 (and its Hilbert Transform) at orientations $(\alpha, \beta, \gamma)_i$) [12].

$$f^{\Omega}(x, y, z) = G(r)Q_N(x') \tag{1}$$

are the even and odd filters formed by a *n*th order polynomial $Q_N(x')^4$ times a separable windowing function G(r) (e.g. a Gaussian-like function), both of which are assumed to be rotationally symmetric. **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 estimate velocities from the filter coefficients using the max-steering method of Simoncelli [13]. Subsequent analysis of these velocities was restricted to include velocity samples only from *within object boundaries* in order to avoid contamination with boundary motion. Velocities were sampled from a grid indicated by the colored dots in Fig. 2C.

⁴ $x' = \alpha x + \beta y + \gamma z$

4 Doerschner et. al

Parametric and Non-parametric Density Estimation. 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. To develop a statistical classifier for reflectivity we estimated the conditional probabilities of 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.

Histograms. Histogram densities were estimated with a generalized crossentropy density estimator [14] 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 [8]. A sample velocity ξ' from a test image sequence was classified as shiny or matte by comparing the likelihood ratio $P(\xi'|S)/P(\xi'|M)$ against a threshold k^5 . 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 discussed below.

Mixture of Gaussians. 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 [15], and computed the posterior probability of each pixel given either Gaussian distribution. Pixel classifications are illustrated by mapping the samples back onto the frame they were taken from. From the two estimated Gaussian means (μ_1, μ_2) we compute the velocity contrast of the sample

$$C_b = \frac{|\mu_1 - \mu_2|}{\max(\sigma_1, \sigma_2)} \ . \tag{2}$$

If $C_b > 1$ the sample is classified as specular, else as matte, and the value of C_b forms a graded material measure.

Mixture of Histograms Using Non-negative Matrix Factorization. To smooth the likelihoods and form a low-dimensional representation for the densities, we factorized the velocity histograms using convolutive non-negative matrix factorization (NNMF) [16]. We preserved 3 components based on an initial estimate that 3 components account for as much as 97% of the approximation error. Because the histogram of a test sequence can be represented as a weighted combination of the 3 components, these weights can be used to represent the velocity 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.

 $^{^5~}k$ was obtained by a bootstrapping procedure used to constrain the false alarm rate to 5%.

Movies. The test set consisted of 36 movies (6 shapes x 6 light probes) of rotating specular superellipoids. Objects were constructed according to

$$1 = \left[\left| \frac{x}{r_x} \right|^{\frac{2}{n_2}} + \left| \frac{y}{r_y} \right|^{\frac{2}{n_2}} \right]^{\frac{n_2}{n_1}} + \left| \frac{z}{r_z} \right|^{\frac{2}{n_1}} .$$
(3)

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. 2A). Each object rotated in depth and its angular speed was adjusted $(0.1, 0.35, 0.61, 0.74, 0.87, 1.0^{\circ}/frame)$ such that resulting image velocities were in the range that our filters were sensitive to.

4 Experimental Results

Histograms. Figure 2B 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 matter object.



Fig. 2. Renderings, Histograms, and Pixel Classification. A. Sample frames for superellipsoids (SE) and for the specular and diffusely reflecting Utah Teapot. Numbers indicate values for n_1, n_2 , in (3). SEs were rendered under 6 different light probes: 2 natural (L1 ("grace"), L3 ("uffizi") from http://gl.ict.usc.edu/Data/HighResProbes/), 2 partially- (L2, L4), 2 fully phase-scrambled (L3, L6) versions of L1 and L3, respectively. For each movie 40 512x512 images were rendered with *Radiance* [17]. B. Corresponding velocity histograms. C. Corresponding pixel classification results. See text for details.

Mixture of Gaussians Pixel Classification. Figure 2C shows that the simple velocity distribution measure was successful in roughly identifying image regions of high (blue pixels) and low (orange pixels) velocities. Purplish colors indicate that the sample could come from either Gaussian distribution. Note,

6 Doerschner et. al

Table 1. Normalized Log-Likelihood Ratios. 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.

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	

Table 2. Average C_b . 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.

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	

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).

Non-negative matrix factorization. The distribution of estimated weights across the stimulus set is shown in Fig. 3A. 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). A very simple shininess criterion can be computed by taking the ratio of the weights of the 2 "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. 3B).

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 [18] and with a specular reflectance (see Fig. 2A (right)). We evaluated the sequence using histograms, mixture of Gaussians, and NNMF 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). Mixture of Gaussians: C_b s for specular



Fig. 3. NNMF of velocity histograms. A. Estimated weights for our test set. B. Average values of C_w : 5.4, 1.8, 1.0, 0.7, 0.5, 0.06. 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). C. Regression of histogram classifications onto observer data. See text for details.

Table 3. Human Shininess Ratings. Shown are ratings for 2 light probes (those eliciting highest and lowest shininess ratings) as well the average data (across all light probes and observers). Differences in relative apparent shininess for different light probes is consistent with previous research [19]. In the experiment observers rated apparent shininess of all 36 light probe – shape combinations.

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

and diffusely reflecting teaport were 1.16 (classified as shiny), and 0.87 (classified as matte). NNMF: The specular teaport classified as shiny $C_w = 33.2$, and the diffusely reflecting teaport was classified as matte $C_w = 0.7954$.

Predicting Human Perception. In the experiment 4 observers indicated via keyboard press on a scale from 1 (matte) - 7 (mirror reflection) how shiny a given superellipsoid appeared. A subset of results are reported in Table 3. Additional experimental details can be obtained from [8]. Regressing normalized LLRs (Table 1) onto normalized observer data (Fig. 3) 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.

5 Discussion

We provide 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 [8,9], demonstrating that diffusely reflecting

8 Doerschner et. al

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.

References

- 1. Ihrke, I., Kutulakos, K., Lensch, H., Magnor, M., W, H.: State of the Art in Transparent and Specular Object Reconstruction (2008)
- Oren, M., Nayar, S.: A Theory of Specular Surface Geometry. International Journal of Computer Vision 24(2) (1997) 105–124
- Roth, S., Black, M.: Specular Flow and the Recovery of Surface Structure. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Volume 2. (2006) 1869–1876
- 4. Healey, G., Binford, T.: Local shape for specularity. Jones and Bartlett Publishers, Inc., USA (1992)
- Nayar, S., Fang, X., Boult, T.: Removal of specularities using color and polarization. In: 1993 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 1993. Proceedings CVPR'93. (1993) 583–590
- Chung, Y., Chang, S., Cherng, S., Chen, S.: Dichromatic Reflection Separation from a Single Image. Lecture Notes in Computer Science 4679 (2007) 225
- 7. Hartung, B., Kersten, D.: Distinguishing shiny from matte. J. Vis. $\mathbf{2}(7)$ (11 2002) $551{-}551$
- Doerschner, K., Kersten, D.: Perceived rigidity of rotating specular superellipsoids under natural and not-so-natural illuminations. J. Vis. 7(9) (6 2007) 838–838
- Roth, S., Domini, F., Black, M.: Specular Flow and the Perception of Surface Reflectance. J. Vis. 3(9) (10 2003) 413–413
- Koenderink, J., Van Doorn, A.: Photometric invariants related to solid shape. Journal of Modern Optics 27(7) (1980) 981–996
- 11. Blake, A.: Specular stereo. In: Proc. Int. J. Conf. on Artificial Intell. (1985) 973-976
- 12. Derpanis, K., Gryn, J.: Three-dimensional nth derivative of Gaussian separable steerable filters. In: IEEE International Conference on Image Processing. (2005)
- Context Simoncelli, E.: Distributed analysis and representation of visual motion. PhD thesis, PhD Thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, MA (1993)
- 14. Botev, Z., Botev, Z.: A Novel Nonparametric Density Estimator. The University of Queensland (2006)
- 15. Nabney, I.: NETLAB: algorithms for pattern recognition. Springer (2002)
- O'Grady, P.D., Pearlmutter, B.A.: Convolutive 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 (September 2006) 427–432
- 17. Larsen, G., Shakespeare, R.: Rendering with Radiance: The Art and Science of Lighting Visualisation (1998)
- Fleming, R.: Rendering Sticky Reflections with Radiance. Personal Communication (2007)
- Fleming, R., Dror, R., Adelson, E.: Real-world illumination and the perception of surface reflectance properties. Journal of Vision 3(5) (2003) 347–368