Rapid Inference of Object Rigidity and Reflectance using Optic Flow

Di Zang[1][*], Katja Doerschner[2], and Paul R. Schrater[1]

[1] Dept. of Computer Science & Engineering, University of Minnesota, USA
{zangx019,schrater}@umn.edu
[2] National Research Center for Magnetic Resonance (UMRAM) & Dept. of
Psychology, Bilkent University, Turkey
katja@bilkent.edu.tr

Abstract. Rigidity and reflectance are key object properties, important in their own rights, and they are key properties that stratify motion reconstruction algorithms. However, the inference of rigidity and reflectance are both difficult without additional information about the object's shape, the environment, or lighting. For humans, relative motions of object and observer provides rich information about object shape, rigidity, and reflectivity. We show that it is possible to detect rigid object motion for both specular and diffuse reflective surfaces using only optic flow, and that flow can distinguish specular and diffuse motion for rigid objects. Unlike nonrigid objects, optic flow fields for rigid moving surfaces are constrained by a global transformation, which can be detected using an optic flow matching procedure across time. In addition, using a Procrustes analysis of structure from motion reconstructed 3D points, we show how to classify specular from diffuse surfaces.

Key words: Optic flow, rigidity detection, specular motion, reflectance classification

1 Introduction

For some computer vision applications like shape analysis from motion, it is typically required to know the material and rigidity of the objects. For instance, there would exist some difficulties to track highly reflective objects like cars without knowing if the object appearance remains constant across frames. Hence, most algorithms usually have strong assumptions about both the reflectivity and rigidity. For example, structure from motion algorithms assume rigidity and it is difficult to extract the point motion information needed without diffusely reflective and patterned objects [1]. Although there are methods to handle both nonrigid structure from motion and shape from specular flow, these methods are derived under the assumption that the rigidity and reflective properties of the object are known [2–5].

---

[*] Corresponding author

Detecting that an object is shiny and rigid would allow a tracking system to rely more on appropriate measurements and improve performance. Methods for rapidly classifying the reflectivity and rigidity of an object would provide the basis for automated recovery. Further, to be most useful, such methods should have minimal information demands. Ideally, we would like an assumption free, fast, image-based method for material and rigidity classification. In this paper we show how optic flow information from a single camera can be used to classify both rigidity of moving objects, and the reflectivity of rigid objects.

Previous methods for classifying material have largely relied on the ability to control the lighting in the scene, using multiple lights, structured lights, color, stereo, or combinations of these. For examples, see [6–11]. Oren and Nayar [12] develop a classification strategy to distinguish image points whose motions affected by specular reflectance from points behaving like diffuse reflectors based on caustic curves. To our knowledge, we are the first to suggest that rigidity can be classified for both diffuse and specular surfaces from optic flow information alone.

In this paper we develop an approach to classify the rigidity and reflectivity of a moving body using only optic flow information. Our approach consists of two parts. We show that rigidity produces characteristic transformations in optic flow that holds for objects with both diffuse and specular reflectance. We exploit this information to develop an optic flow matching algorithm for rigidity classification. We also show how an analysis of the consistency of structure from motion reconstruction can be used to identify diffuse rigid objects.

2   Rigidity from Optic Flow

To detect the rigidity of a specular or diffusely reflecting object from optic flow, we show a simple relationship exists between the optic flow fields at two time points for far-field environmental illumination and orthographic (or paraperspective) viewing. In particular, the flow fields generated by a rigid body motion that differ by a global transformation is derived below.

In order to derive a relationship between optic flow and rigid object motion, we assume that both the viewer and the environment are far from the object, approximated by orthographic viewing and illumination parameterized by direction on a sphere. These assumptions are not overly restrictive as [2] has shown that paraperspective is an exceedingly good approximation for most scenes. As shown in Fig. 1, the object surface $F(x, y) = (x, y, f(x, y))$ is represented as a function of image coordinates x, y, $\boldsymbol{n}(x, y) = S(\theta, \phi)$ indicates the surface normal at the surface point $F(x, y)$ with direction (θ, ϕ), S represents the mapping between spherical and cartesian coordinates, $\boldsymbol{u}(x, y)$ is the optic flow results from the rigid body transformation T. Because the viewing direction is $\boldsymbol{v} = (0, 0, 1)$, the mirror direction $\boldsymbol{r} = S(\theta, 2\phi)$ produces the image point at (x, y).

Rigid body transformation T can be applied to the surface F as $T\left[F(x, y)\right] = R\left[F(x, y)\right] + \boldsymbol{t}$, with R and $\boldsymbol{t}$ refer to the rotation matrix and the translation
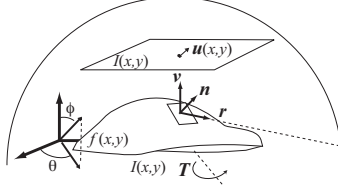
Fig. 1. Assumptions for our treatment of the rigidity from optic flow problem, adapted from [4]. A surface $f(x,y)$, reflecting a far-field illumination environment viewed orthographically to produce an image $I(x,y)$, undergoes a rigid body transformation T.

vector. This induces a motion field in spatial coordinates:

$$\begin{pmatrix} \frac{dx}{dt} \\ \frac{dy}{dt} \end{pmatrix} = I\left(-R\dot{R}^T F(x,y) + \boldsymbol{t} \right) , \tag{1}$$

where $I = \begin{pmatrix} 1\ 0\ 0 \\ 0\ 1\ 0 \end{pmatrix}$ is the orthographic projection matrix, and $\dot{R}^T$ is the transpose of the cross product matrix $\dot{R}$ formed from the rotation axis ω, where $\dot{R}$ takes the following form:

$$\dot{R} = [\omega_\times] = \begin{pmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{pmatrix} . \tag{2}$$

For a fixed rotation axis, $\dot{R}^T F(x,y)$ is a constant flow. Thus, the optic flow pattern generated by a rigid-body transformation is an added translation and a global transformation that is the projection of the rotation onto the cartesian plane $-IR$: the flow is being rotated across time. This means that a global transformation of the motion field across time provides critical information about rigidity. For textured diffusely reflective objects, this motion field result translates directly into optic flow. After removing a global translation, we expect a rigid body motion to produce optic flow patterns that are projected rotations of an initial flow pattern.

We next show a similar result for specular surfaces, which reveals that the global transformations of optic flow patterns is a key piece of information about object rigidity. Because translations simply translate the flow under the viewing and illumination assumptions, we focus on rotations. For a specular surface, if the surface normals are rotated by a rotation R around an axis ω, then the transformation as a function of time is given by $R(t)\,[\omega_\times]$. In cartesian coordinates, $\frac{d\boldsymbol{n}}{dt} = R(t)\,[\omega_\times]\,\boldsymbol{n}$. This transformation of the normal field induces a specular flow field. Adapting the results in [5] to the case of object motion (rather than environment motion), an explicit relationship between the reflection direction and the first order derivatives of the surface can be used to relate differential changes in surface normals to optic flow, when the surface normals are expressed in spherical coordinates:

$$\begin{pmatrix} \frac{d\phi}{dt} \\ \frac{d\theta}{dt} \end{pmatrix} = \begin{pmatrix} \frac{1}{2|\nabla f|(1+|\nabla f|^2)} & 0 \\ 0 & \frac{1}{2|\nabla f|^2} \end{pmatrix} \begin{pmatrix} f_x & f_y \\ -f_y & f_x \end{pmatrix} \begin{pmatrix} f_{xx} & f_{xy} \\ f_{xy} & f_{yy} \end{pmatrix} \begin{pmatrix} \frac{dx}{dt} \\ \frac{dy}{dt} \end{pmatrix} . \tag{3}$$

To convert the normal flow between spherical coordinates and cartesian coordinates, we use the jacobian $\mathbf{J}$ of the cartesian to spherical coordinates mapping: $\left(\frac{d\phi}{dt}, \frac{d\theta}{dt}\right)^T = \mathbf{J}\frac{d\boldsymbol{n}}{dt}$. Chaining these relationships, the difference between a flow at an initial time $t = 0$ and a later time t is a rotation of the flow. This shows that specular flow patterns will differ by global transformations for rigid body motions.

Consequently, by matching optic flow patterns for motion sequences across time, classification can be made based on the measure of average angular error (AAE) [13]. The magnitude of AAE can be used to classify surface points as rigid, with small AAE indicating rigid and large AAE indicating nonrigid.

3   Distinguishing Specular and Diffuse Rigid Bodies

To distinguish rigid motions from diffusely reflective and specular objects, we use structure from motion [14] to reconstruct a candidate shape, and then assess the variation of the shape across time using Procrustes analysis [15]. For diffuse reflective and rigid objects, we would expect the variation in the reconstructed shape to be low and much higher for specular and nonrigid surfaces. Structure from motion is applied to a set of points that are tracked using normalized correlation [16]. To assess shape variation, we used a Procrustes analysis that removed the means of the set of tracked points within each time frame and aligned the points by finding a global rotation that minimized the least-squares difference between corresponding points. But unlike the normal Procrustes analysis, the scale is not removed. The average of the Euclidean distances between corresponding aligned points provides a measure of shape change that can be computed across time lags. Large values of this average shape change (ASC) measure indicate the surface is not both rigid and diffuse reflective. Combined with the optic flow matching measure, these optic flow based measures can distinguish rigid from nonrigid objects, and diffuse rigid from specular rigid motions.

4   Optic Flow Computation

We use a combined global local differential method (CLG) for optic flow computation based on Bruhn et al. [17]. CLG yields accurate, dense flow fields that are robust against noise. The method estimates the flow field by minimizing an energy function:

$$E(\boldsymbol{u}) = \int_{\Omega} (\psi_1(\boldsymbol{u}^T J_\rho(\nabla_3 f)\boldsymbol{u}) + \alpha\psi_2(1 - |\nabla \boldsymbol{u}|^2))dxdy , \tag{4}$$

where Ω denotes the image domain, α serves as regularization parameter, $\boldsymbol{u} = [u, v, 1]^T$ is the flow field, ∇ refers to the spatial gradient, and ∇_3 is the spatio-temporal gradient. The function J_ρ takes the form $J_\rho(\nabla_3 f) = K_\rho * (\nabla_3 f \nabla_3 f^T)$, where K_ρ means a Gaussian kernel with standard deviation ρ. Two nonquadratic penalisers $\psi_1(\cdot)$ and $\psi_2(\cdot)$ are computed as

$$\psi_i(z) = 2\beta_i^2 \sqrt{1 + \frac{z}{\beta_i^2}} \quad i \in \{1, 2\} , \tag{5}$$

with β_1 and β_2 as scaling parameters to handle outliers. For all the parameters, we take suggested values from [17].

5 Experimental Results

Test set. Our test set was comprised of novel 3D objects, generated by sinusoidally modulated spheres, which were organized into 4 categories according to their reflectivity (specular vs. diffuse) and rigidity (rigid vs. nonrigid). Nonrigid deformations were achieved by animating a phase shift of one sinusoidal modulator, in addition to scaling the object either in width (specular) or width and height (diffuse). For each measure (ASC, AAE) we generated 4 (1 per object category) 100-frame test sequences, some example frames are shown in Fig. 2. For ASC experiments, objects underwent a 90° rotation around the viewing direction and an xy-translation, whereas for AAE experiments, objects underwent a 90° rotation only.

Average shape change (ASC). We track object features across the duration of a sequence (see Fig. 3), and compute the ASC by comparing shape changes between the first and second 50-frame block. As shown in Fig. 4A, the ASC measure stabilizes when more than 100 feature points are tracked. Small ASC values reliably indicate the diffusely reflective, rigid object.

Average Angular Error (AAE). Fig. 5 shows sample optic flow fields for each object category. As expected, the flow fields generated by the specular rigid object are very similar between frames - up to a rotation (this is also true for the diffuse, rigid object - but not shown here). However, flow fields for nonrigid objects of either reflectance can vary in non-systematic ways. The AAE was computed by comparing the initial flow field (computed between frames 1 and 2) and the inverse rotated subsequent two-frame flow fields, incrementing frame counts by 10. As results in Fig. 4B illustrate, the AAEs for specular rigid and diffuse rigid objects are relatively stable and small compared to nonrigid objects of either reflectance. Thus it provides a reliable measure of the rigidity of an object.

Table 1 summarizes qualitatively results of each step (1.ASC, 2. AAE) in our approach.

6 Conclusions

We have shown that it is possible to distinguish the rigidity and reflectance of moving objects on the basis of the optic flow fields they generate. Rigidity
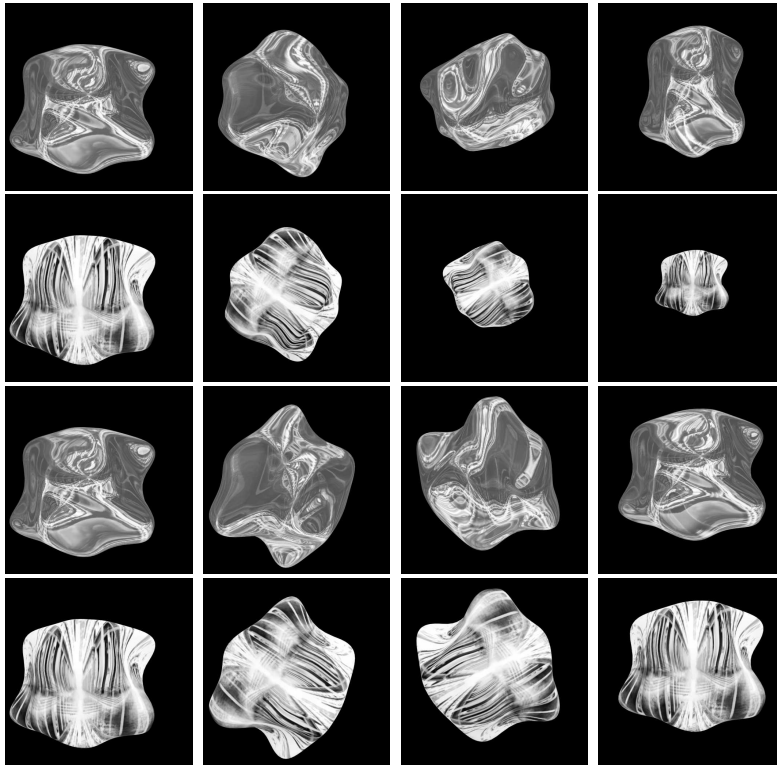
Fig. 2. Example frames (left to right: 1, 34, 67, 100) from our test for each of the 4 objects categories (top to bottom): specular nonrigid, diffuse nonrigid object, specular rigid and diffuse rigid. See text for details.
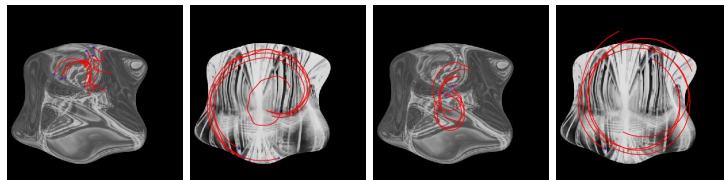

Fig. 3. Selected feature points tracked through 100 frames shown for all 4 object categories (left to right): specular nonrigid, diffuse nonrigid, specular rigid, and diffuse rigid.
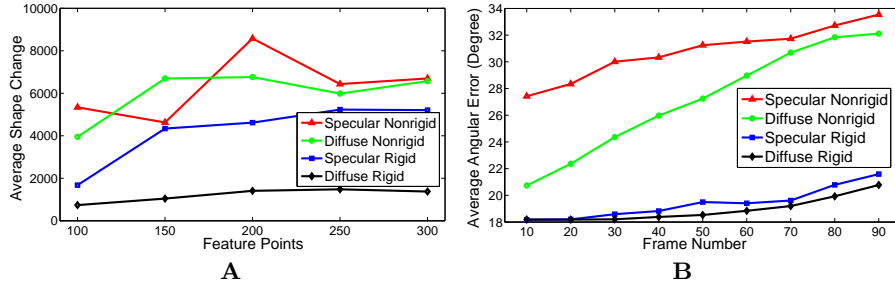
Fig. 4. A. ASC for all 4 object categories as a function of number of tracked feature points. **B.** Average angular errors between an initial flow field based on frames 1-2 and subsequent fields as a function of frame number. AAEs become larger with increasing lag, and are reliably small and stable for rigid objects of either reflectance.

Table 1. Our method allows for a sequential classification approach: In step 1 diffuse rigid objects are successfully classified. In step 2, the AAE reliably distinguishes between rigid and non-rigid objects.

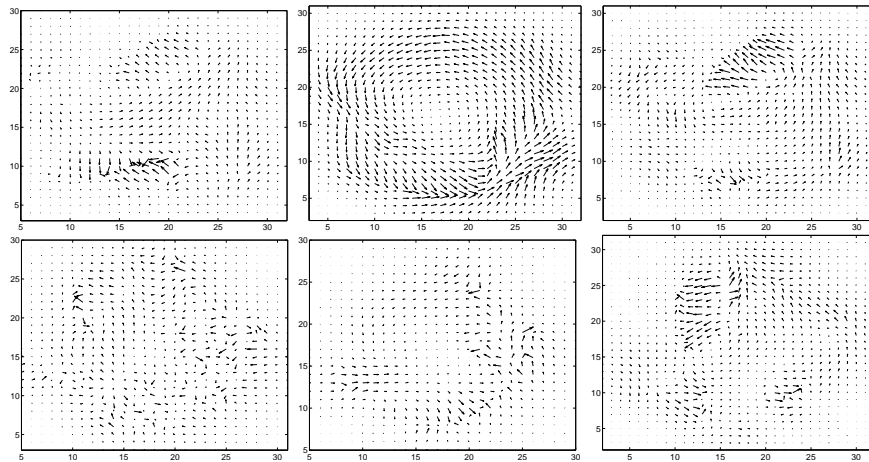| Step in Analysis	Object Class			
	specular		diffuse	
	rigid	nonrigid	rigid	nonrigid
1. ASC	large	large	small	large
2. AAE	small, stable	large, > diffuse	small, stable	large, <specular

Fig. 5. The top row shows initial flow fields (see text) for specular nonrigid, diffuse nonrigid and specular rigid objects, respectively (see Fig 2 Column 1 for corresponding sequence frames).The bottom row shows optic flow fields between frames 51 and 52.

for both specular and diffuse surfaces constrains the optic flow to follow a projected transformation across time. Using a structure from motion reconstruction criterion, it is possible to distinguish specular from diffuse reflectance of rigid motions. In future work it will be possible to formulate a statistical optic-flow based rigidity and reflectivity classifier and quantify the error rates.

References

1. Hartley, R.I., Zisserman, A.: Multiple View Geometry in Computer Vision. Second edn. Cambridge University Press (2004)
2. Zisserman, A., Giblin, P., Blake, A.: The information available to a moving observer from specularities. IVC **7**(1) 38–42
3. Roth, S., Black, M.J.: Specular flow and the recovery of surface structure. In: CVPR '06: Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. (2006) 1869–1876
4. Adato, Y., Vasilyev, Y., Ben Shahar, O., Zickler, T.: Toward a theory of shape from specular flow. In: ICCV07. (2007) 1–8
5. Vasilyev, Y., Adato, Y., Zickler, T., Ben Shahar, O.: Dense specular shape from multiple specular flows. In: CVPR08. (2008) 1–8
6. Healey, G., Binford, T.: Local shape from specularity. CVGIP **42**(1) (1988) 62–86
7. Bhat, D., Nayar, S.: Binocular stereo in the presence of specular reflection. In: ARPA94. (1994) 1305–1315
8. Saito, M., Kashiwagi, H., Sato, Y., Ikeuchi, K.: Measurement of surface orientations of transparent objects using polarization in highlight. In: Proc. of IEEE Computer Society Conference on Computer Vision and Pattern Recognition. (1999) 1381
9. Lin, S., Li, Y., Kang, S., Tong, X., Shum, H.: Diffuse-specular separation and depth recovery from image sequences. In: In Proc. of European Conference on Computer Vision (ECCV). (2002) 210–224
10. Lellmann, J., Balzer, J., Rieder, A., Beyerer, J.: Shape from specular reflection and optical flow. International Journal of Computer Vision **80**(2) (2008) 226–241
11. Chow, S.K., Chan, K.L.: Removal of specular reflection component using multi-view images and 3d object model. In: PSIVT09. (2009) 999–1009
12. Oren, M., Nayar, S.: A theory of specular surface geometry. International Journal of Computer Vision **24**(2) (1997) 105–124
13. Barron, J.L., Fleet, D.J., Beauchemin, S.S.: Performance of optical flow techniques. International Journal of Computer Vision **12**(1) (1994) 43–77
14. Tomasi, C., Kanade, T.: Shape and motion from image streams under orthography: a factorization method. International Journal of Computer Vision **9**(2) (1992) 137–154
15. Gower, J., Dijksterhuis, G.: Procrustes Problems. Oxford University Press (2004)
16. Gonzalez, R.C., Woods, R.E.: Digital Image Processing. Addison-Wesley (1992)
17. Bruhn, A., Weickert, J., Schnörr, C.: Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. International Journal of Computer Vision **61**(3) (2005) 211–231