Dynamic Shading, Motion Parallax and Qualitative Shape

Stephen Waldon Charles R. Dyer

Computer Sciences Department
University of Wisconsin
Madison, WI 53706

Abstract

We address the problem of qualitative shape recovery from moving surfaces. Our analysis is unique in that we consider specular interreflections and explore the effects of both motion parallax and changes in shading. To study this situation we define an image flow field called the reflection flow field*, which describes the motion of reflection points and the motion of the surface. From a kinematic analysis, we show that the reflection flow is qualitatively different from the motion parallax because it is discontinuous at or near parabolic curves. We also show that when the gradient of the reflected image is strong, gradient-based flow measurement techniques approximate the reflection flow field and not the motion parallax. We conclude from these analyses that reliable qualitative shape information is generally available only at discontinuities in the image flow field.*

1 Introduction

We consider the qualitative shape recovery problem for a moving specular surface surrounded by multiple objects. Our analysis accounts for the effects of reflection and therefore requires consideration of the combined effects of motion parallax and changes in radiance from an object's surface (dynamic shading). Discussion of these combined effects has been limited in previous work on qualitative shape analysis. Prior work has focused either on motion parallax [9, 1, 15, 14, 6, 16, 5] or dynamic shading [10, 11, 4, 17, 3, 2], but not both. Although Verri and Poggio [15] considered how photometric effects influence the measured parallax, their goals were different from ours. Their objective was to determine conditions under which the measured parallax agreed qualitatively with the true projected motion field. They found that photometric effects are a potential source of error and that conditions exist where even qualitative interpretations of motion parallax are not valid. However, their conclusions were that these conditions do not occur often and can be neglected in most situations. As we will show, for the class of problems we consider their conclusions are not valid.

Research focusing on qualitative dynamic shading has also not adequately addressed the combined effects of motion parallax and dynamic shading. In this domain the issues are important because motion parallax creates difficulties in measuring changes in shading. Koenderink and van Doorn [10, 11] present their results exclusively in terms of object surface radiance, thus avoiding the need to discuss the effects of motion parallax on image irradiance. Blake *et al.* [4] presented a technique based on Koenderink and van Doorn's results, which provided a means of locating parabolic curves. However, they avoided the problem of object motion by moving the light source instead of the object. Closely related to our work is that of Blake and Bulthoff [2], who used information from the movement of specular points on an object's surface to recover qualitative aspects of surface curvature. However, their method requires compensating for the effects of motion parallax; a computation that requires knowledge of the relative viewer/object velocity. In our analysis we do not assume this knowledge. In addition, their work deals only with specularities from a single source, and does not consider reflections from other objects.

Although previously not considered, we believe that the class of scenes where reflections of surrounding objects are relevant, is fairly large. Examples include the body of a car on a street, that reflects an image of nearby cars, buildings, and overhead trees; or a coffee

The support of the National Science Foundation under Grant No. IRI-9022608 is gratefully acknowledged. We also thank J. Barron and S. Beauchemin for providing the optical flow computation software. The authors can be reached by email at waldon@cs.wisc.edu or dyer@cs.wisc.edu.

cup on an office desk, which reflects an image of the walls and other objects in the room. Since, in these situations, there will be many points on an object's surface that reflect light radiated from other objects, the surface of a moving shiny object will appear as if an image (the reflection of other objects) is moving on its surface. The movement of this reflected image, including the motion of the object as viewed by a stationary observer, can be characterized by an image flow field that we will call the *reflection flow field.* The velocity at any point in the reflection flow field is the perspective projection of the sum of the velocity of the object's surface, plus the velocity of the reflecting point on the surface. This flow field, therefore, combines both the effects of motion and the dynamics of photometry.

In the rest of this paper we analyze in detail the reflection flow field. In Section 2 we reinterpret and extend existing results on specular motion to show the qualitative properties of the reflection flow field. Section 3 presents results that show how the reflection flow behaves with respect to surface reflectance properties. We then relate this to gradient-based techniques for measuring image flow. Section 4 discusses the qualitative properties of the flow and their relationship to qualitative shape description.

2 Kinematic Analysis

To begin our analysis of the reflection flow field, we assume that the surface reflectance properties of the observed object are such that the reflected image of surrounding objects is clearly visible, and that the texture of the object's surface does not interfere with the observation of the reflected image. By neglecting these aspects of photometry, we are free to consider only the kinematics of the reflection flow field. In Section 3 we will return to photometry in order to determine how to measure the reflection flow, but for now we will concentrate on its description and interpretation.

2.1 Modeling Specular Motion

In this section we discuss the construction of a kinematic model that describes the reflection flow field. Figure 1 shows the relevant geometry: Inside a room is an observer positioned at a distance λ from an object's surface. Surrounding the viewer and object are additional objects (e.g. walls). The walls are textured, and the surface of the object has a "strong" specular component. Because of the shininess of the observed

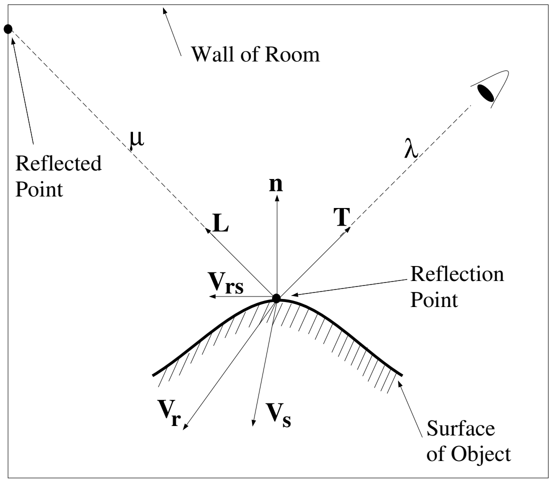

Figure 1: Geometry for moving specular surface.

object's surface, a distorted image of the walls will be reflected from its surface. The distance from the reflection point on the object's surface to any given reflected point in the room is denoted by μ. We assume the viewer is stationary and the object is in motion with velocity V_s (V_s being the instantaneous velocity of the surface at a specific reflection point). Because the object is in motion, the point of reflection for any specific reflected point also moves. The velocity of the reflection point, V_r, can be viewed as the sum of an apparent motion on the objects surface, V_{rs}, and the motion of the surface itself. If the surface is assumed to be smooth, the Weingarten map, W, can be used to describe the rate of change of the surface normal, n_t, as a function of motion in the tangent plane, V_n:

$$n_t = WV_n$$

Under these assumptions, it is possible to derive a mathematical model describing the motion of each reflection point in terms of the motion of the surface. Blake and Bulthoff [2] analyzed the motion of specularities for the case of a moving observer and a stationary object. These results are relevant since for the specular reflectance model they used, motion of a reflection point is the same as that of a specularity (i.e., the brightest point of a specularity is determined by the law of reflection). Therefore, following the analysis of Blake and Bulthoff, we can derive a description of the flow of reflection points for the case of a moving object and stationary viewer (details of this derivation are in Appendix A):

$$V_r{}^p = P(V_{r\,s} + V_s) \qquad (1)$$

where

$$V_{r\,s} = (2(T \cdot n)W - K\Pi_{nT})^{-1}(\frac{\Pi_{nT}}{\lambda} + \frac{\Pi_{nL}}{\mu})V_s$$

$$K = \frac{1}{\lambda} + \frac{1}{\mu}$$

P is the perspective projection map (note: we use superscript p to denote quantities that result from applying this operator). Π_{nT}, Π_{nL} are the composition of Π_n, Π_T and Π_n, Π_L, respectively. Π_n denotes a mapping to the tangent plane and Π_T, Π_L denote mappings to the planes defined by T and L, respectively. We emphasize that in our work we are concerned with the properties of $V_r{}^p$ where as Blake *et al.* were concerned with the properties of $V_{r\,s}{}^p$ as they relate to $V_s{}^p$. They used $sign(V_{r\,s}{}^p \cdot V_s{}^p)$ for a single specular point [17] to determine whether a surface was convex or concave. In addition, they used $V_{r\,s}{}^p$, $V_s{}^p$ and λ (obtained using stereo measurements) to recover a quantitative measure of surface curvature for a single point on the surface. In our work we do not require knowledge of either $V_s{}^p$ or $V_{r\,s}{}^p$.

2.2 Reflection Flow Field Properties

Next, we study the discontinuity structure of the reflection flow field. The main reason for focusing on this particular aspect of the flow is that the location of discontinuities is directly related to qualitative properties of an object's shape (to be shown below). By inspecting the reflection flow equation, Eq. (1), it is evident that the flow contains two types of discontinuities. The first type occurs when the perspective projection map, P, is singular. Physically, this results when the viewing direction is tangent to the object's surface. The perspective projection map discontinuities are present in both the reflection flow field and the motion parallax, and are more commonly known as the occluding contour. These discontinuities were used by Koenderink and van Doorn [9] as a basis for qualitative shape description *vis-a-vis* the aspect graph. A second type of discontinuity occurs when $(2(T \cdot n)W - K\Pi_{nT})$ is singular. As pointed out by Blake and Bulthoff [2], this happens on or near a parabolic curve on the surface. In fact, under orthographic projection, a discontinuity will exist whenever W is singular (i.e., when either of the principle curvatures is zero).

In order to visualize the discontinuities of the reflection flow field, and to better understand the qualitative differences between the reflection flow and the motion parallax, we constructed ray-traced image sequences for the bivariate polynomial shown in Figure 2. The convex portion of this object is shaded in white to reveal the location of the parabolic curve. It is located at the border between the white and shaded regions.

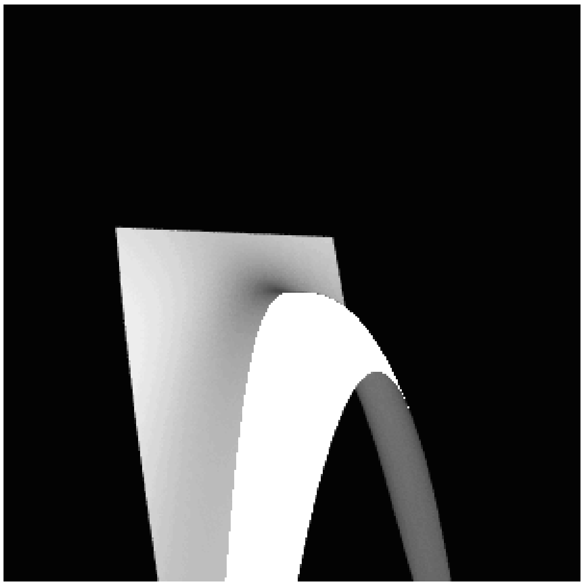

Figure 2: Image of a bivariate polynomial surface. The convex region is marked in white.

In Figure 3 two image sequences and a flow field are shown for the object in Figure 2. The general viewing geometry is shown schematically in Figure 3(a). In the actual ray-traced sequence, the walls were textured with a checkerboard pattern, a distorted image of which can be seen in a reflection from the object's surface. The first image of this sequence is shown in Figure 3(b). The convex portion of the surface is shaded in white. As the object moves, the reflected image moves on the object's surface. The motion of this reflection in the hyperbolic region (unshaded region) is down towards the parabolic curve and the motion of the reflection in the convex region is up towards the parabolic curve. Figure 3(c) shows this flow for the white boxed region in Figure 3(b).

For the sequence just discussed, the viewing conditions were such that the discontinuity was very close to the parabolic curve. In general, the discontinuity can drift from the parabolic curve and will occur where one of the eigenvalues of $(2(T \cdot n)W - K\Pi_{nT})$ is zero. Using a convenient basis for the matrix form of the linear operators of this transform, Zisserman *et al.* [17] showed that the eigenvalues of $(2(T \cdot n)W - K\Pi_{nT})$

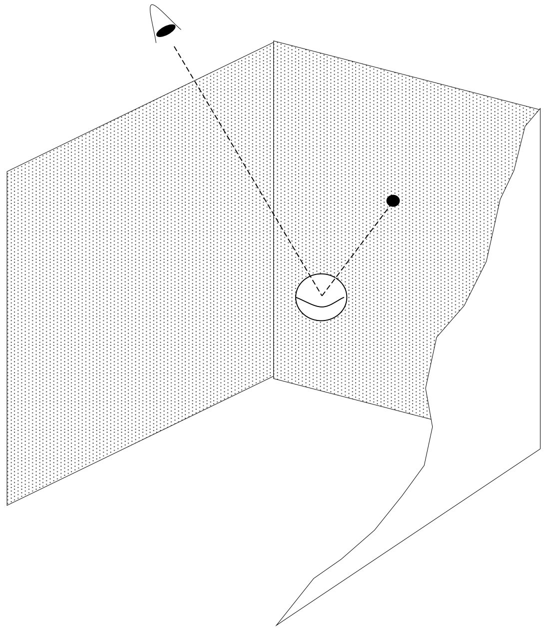

(a)

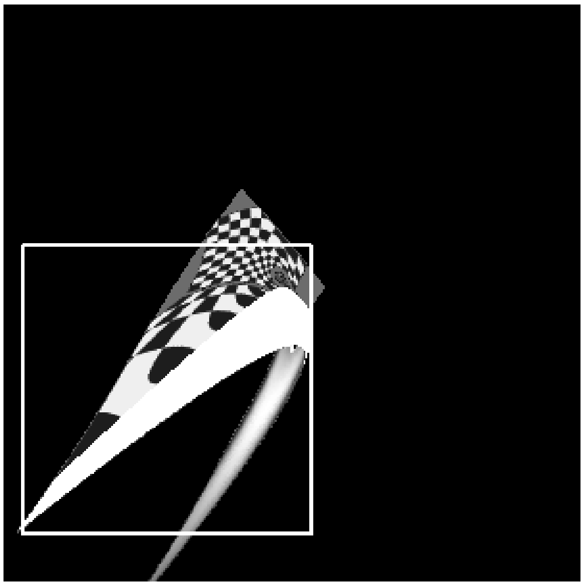

(b)

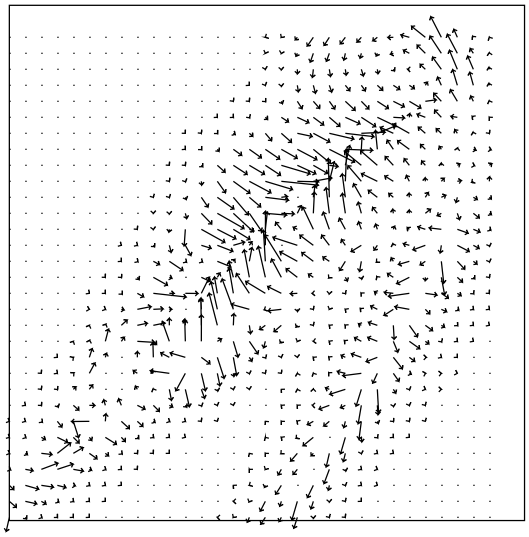

(c)

(d)

Figure 3: (a) Overall viewing geometry for the sequence in (b). (b) First image of sequence for geometry in (a). (c) Reflection flow field for the sequence in (b) (computed using Horn's method [7]). (d) First image of another sequence where the geometry is similar to that in (a) but the object is closer to both the walls and the viewer. Note that the discontinuity has drifted from the parabolic curve.

are given by:

$$\lambda_1 = k_1 - sec(\phi)K \quad (2)$$

$$\lambda_2 = k_2 - sec(\phi)K \quad (3)$$

where ϕ is the angle between the viewing direction and the surface normal, and k_1, k_2 are the principle surface curvatures. Figure 3(d) shows an example of when the viewing conditions are such that the actual discontinuity has drifted off of the parabolic curve. To make this easier to see, the convex region is shaded in dark gray and the region next to it in light grey. The light gray region begins at the parabolic curve and extends into the hyperbolic region. It terminates at the point where the concave curvature meets the condition for an eigenvalue to be zero (i.e., $k_1 = sec(\phi)K$). Therefore, the distance from the white/checkerboard border to the light gray/dark gray border is the amount the discontinuity has drifted from the parabolic curve.

Using the eigenvalues from the previous analysis, we can also determine when the spectral velocity will be large enough to make the discontinuity detectable in a discrete image. In theory, since the reflected point's velocity is infinite at singular points on the object's surface, its magnitude will be much greater than the surface velocity itself. Therefore, a discontinuity should always be observable in the measured field. Because of finite resolution, however, if in the neighborhood of the discontinuity the spectral velocity drops very quickly, the discontinuity itself may appear only as noise. In order for the velocity of the reflection point V_{r_s} to dominate, we require that

$$\|(2(T\cdot n)W - K\Pi_{nT})^{-1}(\frac{\Pi_{nT}}{\lambda} + \frac{\Pi_{nL}}{\mu})V_s\| \gg \|V_s\| \quad (4)$$

Therefore

$$\|(2(T\cdot n)W - K\Pi_{nT})^{-1}\| \gg 1 \quad (5)$$

From linear algebra it is known that

$$\|(2(T\cdot n)W - K\Pi_{nT})^{-1}\| \approx \frac{1}{\lambda_i}$$

where λ_i is the smallest eigenvalue of $(2(T\cdot n)W - K\Pi_{nT})$. From Eq. (2) or (3) we can now see what conditions are required for the relation in Eq. (5) to hold. One of the surface curvatures, say k_1, in the neighborhood of the discontinuity must remain as close to $-sign(k_1)sec(\phi)K$ as possible. This will make the associated eigenvalue as small as possible, and thus Eq. (5) will hold. In addition, for the discontinuity to be close to the parabolic curve, $sec(\phi)K$ should be as small as possible. This will happen when the viewer and reflected point are far away (as in Figure 3b) and when the viewing direction is close to the surface normal direction.

Another significant conclusion that we can draw from this analysis, is that when the inequality in Eq. (5) holds, photometric effects dominate the reflection flow field. This means that the direction and magnitude of the reflection flow is no longer qualitatively similar to the motion parallax and hence qualitative structure from motion parallax techniques will no longer be applicable.

3 Photometric Analysis

In the last section we described qualitative kinematic properties of the reflection flow field. We assumed that the surface reflectance properties were such that reflections of surrounding objects could be detected on the object's surface. In this section we present a more detailed photometric analysis that will shed additional light on the relationship between the motion parallax and reflection flow fields. We begin by developing a reflectance model that accounts for reflections of other objects. Using this model we then show how the strength of the motion parallax and reflection flow signals are dependent on surface reflectance characteristics. We then relate this analysis to existing gradient-based flow measurement techniques, and present results for several image sequences.

3.1 The Reflectance Model

Because the reflection flow field is dependent on the reflection of points from nearby objects, it is necessary to take into account the effects of nonuniform illumination. In this section we develop a photometric model that characterizes the effects relevant to the reflection flow. It is known from photometry that the total irradiance at any point on a given surface, $E_{\bar{s}}$, can be calculated by integrating the incident light distribution over the hemisphere, Ω, surrounding the point on the surface:

$$E_{\bar{s}} = \iint_{\Omega} E_s(\Omega)d\Omega$$

To calculate the radiance, R, in a given direction, T, from any point on the surface, the incident light distribution must be weighted by the surface BDRF:

$$R(T) = \iint_{\Omega} f(\Omega, T)E_s(\Omega)d\Omega$$

To relate the surface radiance to the image irradiance, E_i, we will assume that the image forming system has been calibrated so that

$$E_i(T) = R(T)$$

In order to simplify the model, we will assume that the BRDF is a linear combination of a Lambertian component, f_{lb}, and a specular component, f_{sp}. Therefore

$$E_i(T) = \iint_{\Omega} (f_{lb}(\Omega, T) + f_{sp}(\Omega, T)) E_s(\Omega) d\Omega$$

To handle the nonuniform illuminant for the Lambertian case, we will assume that the relation, $\nabla E_{\bar{s}} = 0$, holds over all visible portions of the scene. This means that the total amount of light falling on the hemisphere surrounding each visible point is constant. Using this assumption, we can replace the nonuniform illumination by an equivalent uniform illuminant, E_{su}, in the direction I_u. Also, by the Lambertian assumption, f_{lb} is constant over Ω, and thus the Lambertian component of E_i is given by:

$$E_{i_{lb}}(T) = k_{lb} N \cdot I_u \qquad (6)$$

where k_{lb} is the Lambertian reflectance coefficient and its magnitude is governed by the magnitude of E_{su} and the Lambertian BDRF.

To simplify the specular component, we use a single lobe specular reflectance model similar in spirit to Phong's model [12].

$$E_{i_{sp}}(T) = \iint_{\Omega} k_{sp} B(\rho, L) E_s(\Omega) d\Omega \qquad (7)$$

Here, k_{sp} is the specular reflection coefficient. The weighting function B is such that its value is high in the neighborhood of the direction of perfect reflection, $L = T - 2(T \cdot N)N$, and drops off quickly for directions away from L. This rate of drop off is a function of the parameter ρ. The larger ρ is, the steeper the fall off is. Physically, ρ controls how "crisp" the reflected image is. As ρ approaches infinity, B approaches a delta function at L, and consequently models a perfect specular reflector. Lower values of ρ are typical of surfaces that reflect "fuzzy" images of surrounding objects. We can simplify Eq. (7) by assuming that $\nabla B(\rho, \mathrm{L}) = 0$, (i.e., the shape of B is the same regardless of the direction of perfect reflection). The simplification results because we can now use linearity to convolve the BDRF with the nonuniform illumination incident on the hemisphere Ω. The result of this convolution is then sampled at the reflected direction (the sampling is accomplished by multiplying the integrand by a delta function at the reflected direction L). Under our assumption about B, the following equation and Eq. (7) should be identical.

$$E_{i_{sp}}(T) = \iint_{\Omega} \delta(L) k_{sp} B(\rho, \Omega) * E_s(\Omega) d\Omega$$

which can be rewritten as

$$E_{i_{sp}}(T) = k_{sp} E'_{\bar{s}}(L) \qquad (8)$$

where

$$E'_{\bar{s}}(L) = B(\rho, \Omega) * E_s(\Omega) \mid_L$$

and $*$ denotes convolution.

Combining the Lambertian and specular components from Eqs. (6) and (8), the complete model becomes

$$E_i(T) = k_{lb} N \cdot I_u + k_{sp} E'_s(L) \qquad (9)$$

To demonstrate the effects of the parameters in Eq. (9), two images are shown in Figure 4. Each image is of the same scene as the image in Figure 3(b). Figure 4(a) uses a Lambertian reflectance coefficient that is stronger than the specular reflectance coefficient. This results in a dim reflected image. However, ρ is large (B is a delta function) and therefore the reflected image also appears "crisp." In Figure 4(b), the specular reflection coefficient is stronger than the Lambertian reflection coefficient, but ρ is relatively small. This causes the reflected image to appear "fuzzy" although brighter than in Figure 4(a).

3.2 Gradient-Based Flow Measurement

Up to this point, the example reflection flow fields were simulated using a modified ray-tracing system. In this section, we will discuss the measurement of reflection flow from image sequences. Specifically, we will concentrate on gradient-based techniques developed to measure motion parallax. For the sake of simplicity, we will use the first-order motion constraint equation [7] as a basis for the rest of our discussion on flow measurement. To develop this equation, we begin with

$$\frac{dE}{dt} = \frac{\partial E}{\partial t} + \|\nabla E\| V_{\bar{s}}{}^{p}$$

where $V_{\bar{s}}{}^{p}$ is the component of the motion parallax in the direction of the image intensity gradient and is sometimes referred to as the normal flow (note: in the above equation we use bar on the velocity subscript to denote the component normal to the intensity gradient). If the radiance from an object's surface is assumed to be stationary over time (i.e., $dE/dt = 0$), the

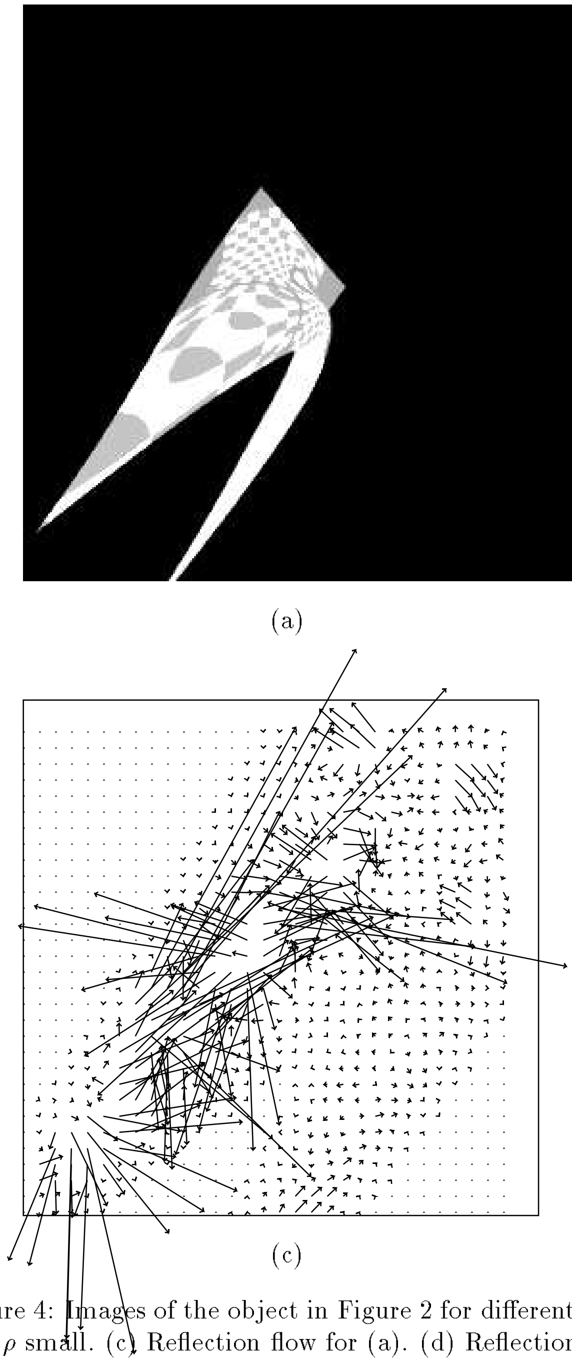

(a) (b)

(c) (d)

Figure 4: Images of the object in Figure 2 for different values of $k_l, k_{sp} and \rho$. (a)$k_l > k_{sp}$ and ρ large. (b) $k_l < k_{sp}$ and ρ small. (c) Reflection flow for (a). (d) Reflection flow for (b).

normal flow can be computed from an image sequence using

$$V_{\bar{s}}{}^{p} = \frac{\partial E/\partial t}{\|\nabla E\|}\frac{\nabla E}{\|\nabla E\|} \tag{10}$$

The assumption that the radiance from an object's surface is stationary over time means that all changes in the image intensity gradient are due to relative movement of object and viewer. That is, the dynamic image intensity signal contains only a motion parallax component. This assumption has been shown to hold generally only under uniform illumination, Lambertian reflectance, and planar motion with no rotation [15].

To see the effect of nonuniform illumination and specular reflectance on the motion constraint equation, we can use the model formulated in the last section to derive an expression for the image intensity gradient. Differentiating Eq. (9) we get

$$\nabla E_i = \nabla E_{i_{lb}} + \nabla E_{i_{sp}} \tag{11}$$

where

$$\nabla E_{i_{lb}} = k_{lb}\nabla(N \cdot I_u)$$
$$\nabla E_{i_{sp}} = k_{sp} J_{T,L} \nabla_L E_s'(L)$$

and $J_{T,L}$ is the Jacobian, which relates a change in viewing direction to a change in the reflected direction. ∇ is the gradient operator with respect to viewing direction, and ∇_L is the gradient operator with respect to the illuminating hemisphere at the reflecting point. Using Eq. (11) it is possible to relax the stationarity assumption made in the original formulation of the motion constraint equation. By assuming that changes in image intensity are caused by a shift in a stationary Lambertian gradient plus a shift in a stationary specular gradient, the following relation is obtained:

$$\nabla E_i \cdot V_{mc} = \nabla E_{i_{lb}} \cdot V_{\bar{s}}{}^{p} + \nabla E_{i_{sp}} \cdot V_{\bar{r}}{}^{p}$$

where V_{mc} is the velocity measured by the motion constraint equation for the moving specular surface, and $V_{\bar{r}}{}^{p}$ is the component of the reflection flow field in the direction of the specular image gradient. Substituting this into the motion constraint equation, Eq. (10), we get

$$V_{mc} = \frac{\nabla E_{i_{lb}} \cdot V_{\bar{s}}{}^{p} + \nabla E_{i_{sp}} \cdot V_{\bar{r}}{}^{p}}{\|\nabla E_i\|}\frac{\nabla E_i}{\|\nabla E_i\|}$$

From this expression it is evident that if the stationary assumptions for Eq. (12) hold and, in addition, if

$$\|\nabla E_{i_{sp}}\| \gg \|\nabla E_{i_{lb}}\| \tag{12}$$

then the reflection flow field will dominate the change in image intensity and V_{mc} will approximate $V_{\bar{r}}{}^{p}$ and not $V_{\bar{s}}{}^{p}$. To recover the actual specular motion field, $V_r{}^{p}$, second-order gradient-based techniques can be used [13], and an analysis similar to that performed above will apply.

To demonstrate how the measured flow is affected by different photometric parameters of our model, we used Horn's method to recover the reflection flow field for the sequence in Figure 3(b) and two additional image sequences. The results of the two additional sequences are shown in Figures 6(c) and 6(d). Figure 3(b) represents the ideal conditions for reflection flow recovery. The specular coefficient is high compared to the Lambertian coefficient, and the value of ρ is also large. Both of these conditions serve to increase $\|\nabla E_{i_{sp}}\|$ and hence increase the strength of the reflection flow field signal. Figures 6(c) and 6(d) show flow fields for dim and fuzzy reflected image sequences (see Figures 6(a) and 6(b) for the first images of these two sequences). Each case results in a decrease of $\|\nabla E_{i_{sp}}\|$. For the dim image sequence, this results mostly in a noisy recovered field, but for the fuzzy image sequence the flow appears to be even better than for the ideal image sequence. This is because the additional smoothing provided by an increase in ρ aids the process of numerical differentiation. However, too large values of ρ will degrade the reflected image to the point where the gradient will be too small to measure. This again will result in a very noisy recovered field.

4 Discussion and Conclusions

In this paper we described the qualitative properties of an image flow field that combines information from dynamic shading and motion parallax. The dynamic shading component of the flow arises from the motion of reflection points on an object's surface. The motion parallax component arises from relative viewer/object motion. From a kinematic analysis of the flow, we showed that the dynamic shading effects dominate the flow in the vicinity of a parabolic curve, and give rise to a discontinuity in the flow on or near the parabolic curve. From a photometric analysis, we showed that when the gradient of the reflected image is strong, gradient-based flow measurement techniques approximate the reflection flow field and not the motion parallax. Unfortunately, without detailed knowledge of object velocity, viewing geometry, surface reflectance properties and illumination, it is not possible to determine whether a region of the image flow is dominated by changes in shading or by motion parallax. In general, this means it is not possible to know when to apply qualitative motion parallax analysis or qualitative dynamic shading analysis (not to mention the inadequacy of using quantitative methods).

Fortunately, some invariants do exist. We have shown that discontinuities of the flow occur either on or near parabolic curves or at occluding boundaries. Knowledge of the location and connectivity of these curves has been shown to determine significant topological characteristics of an object's shape [8]. In addition, since the flow velocity in the neighborhood of the parabolic curve is large compared to the flow on other parts of the surface, it should be possible to determine the underlying cause of the discontinuity. This would make it possible to determine whether the discontinuity corresponds to a parabolic curve or an occluding boundary. We are currently investigating methods for locating and identifying discontinuities in the flow, and we are investigating suitable representation schemes for this information.

References

[1] F. Aicardi. On the existence of limit cycles in motion field. *Biological Cybernetics*, 62:99–106, 1989.

[2] A. Blake and H. Bulthoff. Shape from specularities: Computation and psychophysics. *Phil. Trans. R. Soc. London B*, 331:237–252, 1991.

[3] A. Blake, R. Cipolla, and A. Zisserman. Towards qualitative vision: motion parallax. In *Proc. AAAI-90 Workshop on Qualitative Vision*, pages 6–10, 1990.

[4] A. Blake, A. Zisserman, and G. Knowles. Surface descriptions from stereo and shading. *Image and Vision Computing*, 3:183–191, 1985.

[5] N. da Vitoria Lobo and J. K. Tsotos. Extracting qualitative shape from image motion: applications to stereo-pairs. In *Proc. AAAI-90 Workshop on Qualitative Vision*, pages 36–40, 1990.

[6] J. Herve and J. Aloimonos. Catastrophe theory and computational vision. In *Proc. AAAI-90 Workshop on Qualitative Vision*, pages 96–100, 1990.

[7] B.K.P. Horn and B.G. Shunk. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.

[8] J. J. Koenderink. *Solid Shape*. MIT Press, Cambridge, MA, 1990.

[9] J. J. Koenderink and A. J. van Doorn. The internal representation of solid shape with respect to vision. *Biol. Cybernetics*, 32:211–216, 1979.

[10] J. J. Koenderink and A. J. van Doorn. Photometric invariants related to solid shape. *Optica Acta*, 27:981–996, 1980.

[11] J. J. Koenderink and A. J. van Doorn. Perception of solid shape and spatial layout through photometric invariants. In R. Trappl, editor, *Cybernetics and Systems Research*, pages 943–948. North Holland, Amsterdam, 1982.

[12] B. T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.

[13] S. Uras, F. Girosi, A. Verri, and V. Torre. A compuational approach to motion percption. *Biological Cybernetics*, 60:79–87, 1988.

[14] A. Verri, F. Gerosi, and V. Torre. Mathematical properties of the 2d motion field: From singular points to motion parameters. In *Proc. Workshop on Visual Motion*, pages 190–200, 1989.

[15] A. Verri and T. Poggio. Motion field and optical flow: Qualitative properties. *IEEE Trans. Pattern Anal. and Mach. Intell.*, 11:490–498, 1989.

[16] D. Weinshall. Shortcuts in the computation of qualitative 3-d shape and motion invariants. In *Proc. AAAI-90 Workshop on Qualitative Vision*, pages 96–100, 1990.

[17] A. Zisserman, P. Giblin, and A. Blake. The information available to a moving observer from specularities. *Image and Vision Computing*, 7:38–41, 1989.

A Derivation of the Kinematic Model

This section presents the derivation of a model of the reflection flow kinematics. It follows very closely that of Blake and Bulthoff [2]. However, in our case, the viewer is stationary and the object is moving. Referring to Figure 1, the velocity of a reflected point, V_r, can be considered as the sum of the velocity of the surface, V_s, and the velocity of the reflecting point on the surface, V_{rs}:

$$V_r = V_s + V_{rs}$$

Since the instantaneous change in the unit vectors T and L, which define the viewing and reflection directions, respectively, are restricted to their associated orthogonal planes, the velocity of the reflected point is related as follows:

$$\lambda T_t = -\Pi_T V_r$$

$$\mu L_t = -\Pi_L V_r$$

The geometry of reflection is given by

$$n = \frac{T+L}{|T+L|}$$

By differentiating this expression, a relation between the instantaneous change in the unit normal and the viewing and reflecting directions is obtained:

$$n_t = \frac{\Pi_n(T_t + L_t)}{2(T \cdot n)}$$

By definition, the velocity of the reflected point on the surface is restricted to the tangent plane. Also, by assumption, the surface is smooth. The Weingarten map can therefore be used to relate $V_{r\,s}$ to the rate of change of the unit normal:

$$n_t = W V_{r\,s}$$

Using the above relations, the following expression results:

$$W V_{r\,s} = \frac{1}{2T \cdot n} \Pi_n(-\Pi_T(\frac{V_s + V_{r\,s}}{\lambda}) + -\Pi_L(\frac{V_s + V_{r\,s}}{\mu}))$$

A little more algebra yields

$$(2(T \cdot n)W - (\frac{\Pi_{nT}}{\lambda} + \frac{\Pi_{nL}}{\mu}))V_{r\,s} = (\frac{\Pi_{nT}}{\lambda} + \frac{\Pi_{nL}}{\mu})V_s$$

Noting that $V_{r\,s}$ is restricted to the tangent plane and using the symmetry of the reflection geometry

$$\Pi_{nT} V_{r\,s} = \Pi_{nL} V_{r\,s}$$

hence,

$$(2(T \cdot n)W - K\Pi_{nT})V_{r\,s} = (\frac{\Pi_{nT}}{\lambda} + \frac{\Pi_{nL}}{\mu})V_s$$

where

$$K = \frac{1}{\lambda} + \frac{1}{\mu}$$

Finally, projecting into the image plane,

$$V_r{}^p = P((2(T \cdot n)W - K\Pi_{nT})^{-1}(\frac{\Pi_{nT}}{\lambda} + \frac{\Pi_{nL}}{\mu})V_s + V_s)$$

where P is the perspective projection operator and $V_r{}^p$ is the velocity in the image plane.