Image Invariants for Smooth Reflective Surfaces

Aswin C. Sankaranarayanan[1], Ashok Veeraraghavan[2], Oncel Tuzel[2], and Amit Agrawal[2]

[1] Rice University, Houston, TX 77005, USA
[2] Mitsubishi Electric Research Labs, Cambridge, MA 02139, USA

Abstract. Image invariants are those properties of the images of an object that remain unchanged with changes in camera parameters, illumination etc. In this paper, we derive an image invariant for smooth surfaces with mirror-like reflectance. Since, such surfaces do not have an appearance of their own but rather distort the appearance of the surrounding environment, the applicability of geometric invariants is limited. We show that for such smooth mirror-like surfaces, the image gradients exhibit degeneracy at the surface points that are parabolic. We leverage this result in order to derive a photometric invariant that is associated with parabolic curvature points. Further, we show that these invariant curves can be effectively extracted from just a few images of the object in uncontrolled, uncalibrated environments without the need for any a priori information about the surface shape. Since these parabolic curves are a geometric property of the surface, they can then be used as features for a variety of machine vision tasks. This is especially powerful, since there are very few vision algorithms that can handle such mirror-like surfaces. We show the potential of the proposed invariant using experiments on two related applications - object recognition and pose estimation for smooth mirror surfaces.

1 Introduction

Image invariants are those properties of the images of an object that remain unchanged with changes in camera parameters, illumination etc. Any geometric invariant (eg., cross ratio) is true for surfaces with any reflectance characteristics including diffuse, specular and transparent surfaces. But, in order to actually use these geometric invariants from observed images of an object, one needs to identify point correspondences across these images. Establishing such point correspondences from images of diffuse objects is a meaningful task since these objects have photometric features of their own. But surfaces with mirror reflectance do not have an appearance of their own, but rather present a distorted view of the surrounding environment. Therefore, establishing physical point correspondences using image feature descriptors (such as SIFT) is not meaningful. Such descriptors find correspondences between environment reflections, and therefore are not physically at the same point on the surface. Thus, there is a need to find photometric properties of specular surfaces that are invariant to the surrounding environment. In this paper, we study and present such an invariant for the images of smooth mirrors.

The main results of this paper arise by studying the photometric properties of the images of mirror surfaces around points that exhibit parabolic curvature. Parabolic curvature points are fundamental to perception of shape both for diffuse [15, 17] and for

specular surfaces [23]. In this paper, we first derive a photometric invariant that is associated with parabolic curvature points of the mirror surface. We show that a smooth mirror imaged by an orthographic camera, reflecting an environment feature at infinity, exhibits degenerate gradients at parabolic curvature points. This degeneracy is characterized by the image gradients being orthogonal to the direction of zero curvature at the parabolic point. Although the invariant holds exactly for the aforementioned imaging setup, we empirically show that for a range of practical imaging conditions (with perspective camera and finite scene), the invariance still holds to a high degree of fidelity.

The set of parabolic points is a geometric property of a surface and each surface has its own distinct set of parabolic curves. The photometric invariant that we propose allows us to detect these parabolic curves from just images of the specular object without any a priori knowledge about its 3D shape or the surrounding environment. Since these parabolic points are a geometric property of the surface, they can then be used for a variety of machine vision tasks such as object recognition, pose estimation and shape regularization. In this paper, we demonstrate a few such applications.

Contributions: The specific technical contributions of this paper are:

- We present a theoretical study of the properties of images of mirrors. We show that under a certain imaging setup, the image derivatives at the points of parabolic curvature exhibit degeneracies independent of the surrounding environment.
- We show that this degeneracy can be measured quantitatively using just a few images of the object under arbitrary illumination, thereby allowing us to recover the parabolic curvature points associated with the mirror.
- We show new applications of this invariant to challenging machine vision problems such as object recognition and pose estimation for mirror objects.

2 Prior Work

In this paper, we are interested in identifying invariants for images of mirrors. Additional assumptions are needed for obtaining something meaningful/non-trivial. A planar mirror viewed by a perspective camera is optically the same as a perspective camera, and hence, can produce arbitrary images.

The qualitative properties of images of specular/mirror objects have been well studied (see [14] for a survey). Zisserman et al. [25] show that local surface properties such as concave/convexity can be determined under motion of the observer without knowledge of the lighting. Blake [4] analyze stereoscopic images of specular highlights and show that disparity of highlights observed on the mirror is related to the qualitative properties of the shape such as its convexity/concavity. Blake and Brelstaff [5] quantify local surface ambiguities given stereo images of highlights. Fleming et al. [11] discuss human perception of shape from images of specular objects even when the environment is unknown and show that humans are capable of accurately determining the shape of the mirror; potentially from image compression cues. In another study of human perception of specular surfaces, Savarese et al. [20] report poor perception when the surrounding scene comprises of unknown but structured patterns.

It is worth noting that points/curves of parabolic curvature have been studied in terms of their photometric properties. Our search is motivated in part from classical results in photometric stereo and more recent work in the area of specular flow. Koenderink and van Doorn [15] demonstrate that the structure of isophotes is completely determined by parabolic curves of the surface. Further, they also show that the a local extremum of the field of isophotes occur on parabolic curves, and move along these curves under motion of the light source. Isophotes as a construct are useful for diffuse objects with constant albedo and mirrors under simple lighting (such as a point light source), and do not extend well to scenes/object with rich textures.

Much of prior work using properties of parabolic points revolve around the idea of consistency of highlights at parabolic curvature points across small changes in view or illumination. Miyazaki et al [18] use parabolic curves for registration of transparent objects across views.

Recent literature has focused on estimating the shape of the mirror from the specular flow [1, 21, 7] induced under motion. Specular flow is defined as the movement of environmental features on the image of a mirror due to motion of the mirror/scene. It has been shown that parabolic curvature points exhibit infinite flow under infinitesimal motion. The infinite flow is a result of appearance of new scene features and disappearance of existing ones, an observation made earlier by Longuet-Higgins [16] and Walden and Dyer [22] as well. Waldon and Dyer [22] suggest that, for mirrors, reliable qualitative shape information is available only at the parabolic curves in the forms of discontinuous image flow fields. Studies in perception [17] hint at the ability of humans to detect and use parabolic curvature curves to perform local shape analysis. Some existing approaches in perception [23] and detection [9] of mirrors remark on the anisotropy of gradients in the images of smooth mirrors. However, these papers do not identify the existence of the invariance, the assumptions required for the invariance to hold, the geometric interpretation behind its occurrence and the stability of the invariance for practical imaging scenarios. More importantly, in addition to exploring these properties, we also show that parabolic curvature points of the mirror (a surface descriptor) can be recovered from a few images of the mirror.

3 Deriving the Invariant

This section describes the main technical contributions of the paper. We begin with a brief overview of parabolic curvature points. Then, we discuss the image formation model for mirror objects and show that the observed image gradients exhibit a degeneracy at the points of parabolic curvature, irrespective of the environment. This leads us to define an invariant for the surfaces of smooth mirrors.

3.1 Parabolic Curvature Points

Let us model the shape of the (smooth) mirror in its Monge form $(x, y, f(x, y)) = (\mathbf{x}, f(\mathbf{x}))$ in a camera coordinate system where the function f is twice continuously differentiable. At a given point on the surface, the curvature along a curve is defined as the reciprocal of the radius of the osculating circle. The principal curvatures are

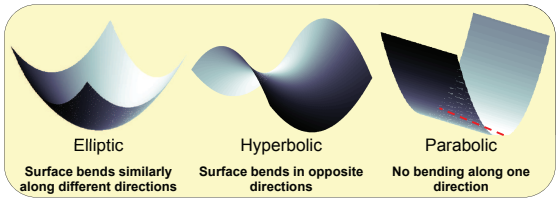

Fig. 1. Local properties of a surface can be classified into 4 types: elliptic, hyperbolic, parabolic and flat umbilic (planar). This classification deals with the bending of the local surface in various directions. The parabolic curve is shown in red.

defined as the minimum and maximum values of the curvature measured along various directions at a given point. The product of the principal curvatures is defined as the Gaussian curvature. It can be shown [15] that the Gaussian curvature is given by

$$\frac{f_{xx}f_{yy} - f_{xy}^2}{1 + f_x^2 + f_y^2}. \tag{1}$$

Points at which one of the principal curvatures is zero are termed parabolic curvature points or simply parabolic points. Defining the Hessian at point $\mathbf{x}$ of the surface as

$$H(\mathbf{x}) = \frac{1}{2}\begin{bmatrix} f_{xx} & f_{xy} \\ f_{xy} & f_{yy} \end{bmatrix}_{(\mathbf{x})}, \tag{2}$$

parabolic curvature points are defined by points where $\mathrm{rank}[H(\mathbf{x})] = 1$. When both principal curvatures at a point are zero, the point is referred to as *flat umbilic*. Planes are examples of surfaces which are flat umbilic everywhere. Shown in Figure 1 are characterization of local properties of a surface.

3.2 Image Formation For Mirror Objects

Mirrors do not have an appearance of their own, and image of mirror are warps of the surrounding environment. Modeling the shape of the mirror as $(\mathbf{x}, f(\mathbf{x}))$, image formation can be described by identifying the camera and the environment. We model the camera as orthographic. Under an orthographic camera model, all the rays entering the camera are parallel to its principal direction.

The surface gradient at pixel location $\mathbf{x}$ is given as $\nabla f = (f_x, f_y)^T$, and the surface normal is given as

$$\mathbf{n}(\mathbf{x}) = \frac{1}{\sqrt{1 + \|\nabla f\|^2}} \begin{pmatrix} -\nabla f \\ 1 \end{pmatrix}. \tag{3}$$

The camera viewing direction $\mathbf{v}$ at each pixel is the same, $\mathbf{v} = (0, 0, 1)^T$. Under perfect mirror reflectance, we can compute the direction of the ray that is reflected onto the camera as $\mathbf{s} = 2(\mathbf{n}^T\mathbf{v})\mathbf{n} - \mathbf{v}$ The corresponding Euler angles $\Theta(\mathbf{x}) = (\theta, \phi)$ are given as,

$$\tan\phi(\mathbf{x}) = \frac{f_y}{f_x}, \qquad \tan\theta(\mathbf{x}) = \frac{2\|\nabla f\|}{1 - \|\nabla f\|^2}. \tag{4}$$

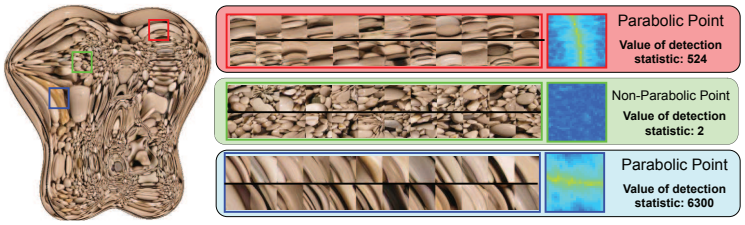

Fig. 2. Degeneracy of image gradients at parabolic points. Shown are image patches at three locations on a mirror from *multiple* images rendered under a rotating environment. The image patches corresponding to parabolic curves do not have have gradients along the parabolic curve. This is in contrast to a non-parabolic point which can have arbitrary appearance.

The scene/environment that is seen at pixel $\mathbf{x}$ is hence defined by the intersection of the scene and the ray in the direction of $\mathbf{s}(\mathbf{x})$ from the location of the mirror element $(\mathbf{x}, f(\mathbf{x}))^T$. In the special case of environment at infinity, the dependence on the location of the mirror is completely suppressed, and the environment feature observed depends only on the surface gradient ∇f.

Under the assumption of environment at infinity, we can define the environment map over a sphere. Let $E : \mathbb{S}^2 \mapsto \mathbb{R}$ be the environment map defined on the sphere $\mathbb{S}^2$ under the Euler angle parametrization. Under *no inter-reflectance* within the object, the forward imaging equation for the intensity $I(\mathbf{x})$ observed at pixel $\mathbf{x}$ is given as

$$I(\mathbf{x}) = E(\Theta(\mathbf{x})) \tag{5}$$

where $\Theta(\mathbf{x})$ is the Euler angle of the observed ray as given in (4). Differentiating (5) with respect to $\mathbf{x}$, the image gradients are given by

$$\nabla_{\mathbf{x}} I = 2H(\mathbf{x}) \left[\frac{\partial \Theta}{\partial \nabla f}\right]^T \nabla_{\Theta} E \tag{6}$$

where the Hessian $H(\mathbf{x})$ is defined in (2). The full derivation is in the supplemental material (and similar to that of [1]).

For parabolic curvature points, $H(\mathbf{x})$ is singular. As a consequence, $\nabla_{\mathbf{x}} I$ takes values that are proportional to the non-zero eigenvalue of $H(\mathbf{x})$, immaterial of what the environment gradient $\nabla_{\Theta} E$ is. Figure 2 shows the local appearance of parabolic and non-parabolic points under various environment maps.

3.3 Invariant

Given a smooth mirror $(\mathbf{x}, f(\mathbf{x}))$, where f is $\mathbb{C}^2$ continuous, placed with the surrounding environment at infinity and viewed by an orthographic camera, the proposed invariant is a statement on the observed image gradient at parabolic curvature points of the mirror. Under this setting, the image gradients at parabolic curvature points are *degenerate* and take values along a single direction that is defined by the local shape of the surface. This property is independent of the scene in which the mirror is placed.

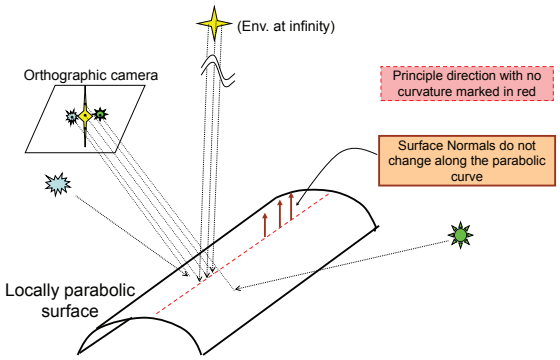

Fig. 3. The proposed invariant can be geometrically described using the ray diagram above. At a parabolic point, by definition, the normal (and curvature) do not change along one direction. Under orthographic imaging and scene at infinity, the same feature is imaged onto the camera as we move along the parabolic curve and hence, the image gradient disappears along this direction.

The invariant arises directly due to the principal direction of zero curvature at parabolic point (see Figure 3). By definition, an infinitesimal movement on the surface along this direction does not change the surface normal. Under our imaging model, the environment feature imaged at a point depends only on the surface normal. Hence, an infinitesimal displacement on the image plane along the projection of this direction does not change the environment feature imaged. As a consequence, the image gradient along this direction is zero. This geometric understanding is related to recent work on ray-space specular surface analysis [24, 10], where the authors study local mirror patches as general linear cameras and associate different camera models to different local surface properties. Interesting connections to our work can potentially be derived from these papers and remain an important direction for future research.

Mathematically, the invariant can be expressed in various forms. From (6) and under the assumed imaging conditions, a parabolic curvature point at $\mathbf{x}_0$ satisfies

$$\nabla_{\mathbf{x}} I(\mathbf{x}_0) = \|\nabla_{\mathbf{x}} I(\mathbf{x}_0)\| \mathbf{v} \tag{7}$$

where $\mathbf{v}$ is the eigenvector of $H(\mathbf{x}_0)$ with non-zero eigenvalue. An alternate interpretation that does not involve $H(\mathbf{x})$ uses the matrix $M(\mathbf{x})$ defined as:

$$M(\mathbf{x}) = \sum_{E} \left((\nabla_{\mathbf{x}} I(\mathbf{x}; E))(\nabla_{\mathbf{x}} I(\mathbf{x}; E))^T \right) \tag{8}$$

where $I(x; E)$ is the intensity observed at pixel $\mathbf{x}$ under environment defined in $E(\Theta)$. Note that the summation in (8) is over all possible environment maps. At points of parabolic curvature,

$$\text{rank}[M(\mathbf{x}_0)] = 1 \tag{9}$$

In contrast, for elliptic and hyperbolic points, the matrix M is full rank. For flat umbilical points, $H(\mathbf{x})$ is the zero matrix and the image gradients are zero as well. Therefore, the matrix defined in (8) will be zero rank.

4 Detecting Parabolic Curvature Points

We derive a practical algorithm for detecting points of parabolic curvature from multiple images of a mirror. The algorithm exploits the consistency (or degeneracy) of the image gradients as given in (9). Under motion of the camera-mirror pair (or equivalently, rotation of the environment), the environment feature associated with each point changes arbitrarily. However, parabolic points are *tied* to the surface of the mirror, and hence, the direction of image gradients associated with them do not change. This motivates an acquisition setup wherein the environment is changed arbitrary and consistency of image gradient at a pixel indicates whether or not it has parabolic curvature. Since movement of the camera-object pair simultaneously is the equivalent to that of rotation of the environment, we use environment rotation to denote both. In practise, moving the camera-object pair is easier to accomplish.

Given a set of images $\{I_j\}$, compute the matrix

$$M(\mathbf{x}) = \sum_j \left(\nabla_{\mathbf{x}} I_j(\mathbf{x})\right) \left(\nabla_{\mathbf{x}} I_j(\mathbf{x})\right)^T \tag{10}$$

using image gradients computed at each frame. We use the ratio of the eigenvalues of $M(\mathbf{x})$ as the statistic to decide whether or not a pixel $\mathbf{x}$ observes a parabolic curvature point. Figure 4 shows estimates of parabolic points of different surfaces. Images for this experiment were rendered using PovRay. Each image was taken under a arbitrary rotation of the environment. As the number of images increase, the detection accuracy increases significantly as $M(\mathbf{x})$ at non-parabolic points become well-conditioned.

We believe that our approach is unique in the sense that it recovers a 'dense' characterization of parabolic curvature points from uncalibrated images of a mirror. Much of the existing literature on using the photometric properties of parabolic points rely on the stability of highlights at parabolic points under changes in views. However, such a property is opportunistic at best, and does not help in identifying all the parabolic curvature points associated with the visible surface of the mirror. In this sense, the ability to recover a dense set of parabolic curvature points opens the possibility of a range of applications. We discuss these in Section 5.

Theoretically, the invariance is guaranteed only for an orthographic camera and an environment at infinity. However, in practice, the invariance holds with sufficient fidelity when these assumptions are relaxed. We explore the efficacy of the proposed invariant for a range of practical operating conditions in Section 6.

Mis-detection: The proposed invariant does not take inter-reflections into account. Inter-reflections alter the physics of the imaging process locally, and violates the relations made earlier in physical models. Imaging resolution also affects the detection process. For low resolution images, the curvature of the surface observed in a single pixel might deviate significantly from parabolic. Such a scenario can potentially annul the invariance at the parabolic point due to corruption from the surrounding regions.

False Alarm: It is noteworthy that the invariant describe image gradients at parabolic points. However, degenerate gradients do not necessarily imply the presence of parabolic points. Clearly, for small number of images, it is possible that a surface pixel/patch do

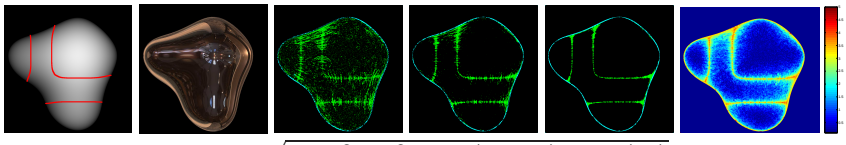

(a) $z = \sqrt{4 - x^2 - y^2 - cos(2x - 2) - sin(2y)}$

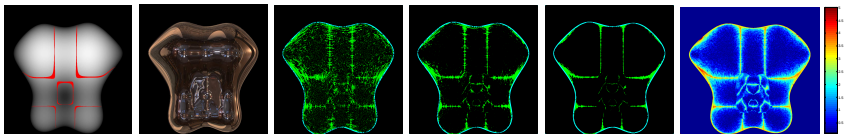

(b) $z = \sqrt{4 - x^2 - y^2 - cos(3x - 6) - 2sin(2y)}$

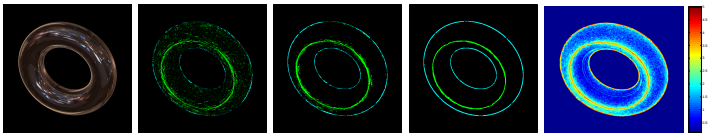

(c) Rotated Torus

Fig. 4. Detecting parabolic curvature points from rendered images for various surfaces. (From left to right) the depth map of the mirror with the parabolic points highlighted in red; a rendered image of the surface using the *Grace cathedral* environment map; detected parabolic points from 2 images; from 5 images; from 25 images; Log of decision statistic estimated from 25 images. The occluding contour is shown in cyan and the parabolic points in green.

not observe environment features that are sufficiently rich. Similarly, discontinuities in the surface such as occluding contours can lead to consistent degeneracy in the observed image gradients.

Note that, the detection of statistics does not require the environment texture to be rich. Using increasing number of images (camera-mirror pair rotations), the degeneracies due to environment become incoherent and can be filtered out easily. Our experiments include textures such as the Grace cathedral which exhibit large regions with little or no textures and the method succeeds to capture the statistics regardless.

5 Applications

In this section, we describe three applications of the presented theory; (1) pose estimation, (2) recognition of mirror-like objects; and (3) a possible extension for surface reconstruction. The equivalent algorithms designed for diffuse/textured surfaces require establishing correspondences between image observations and a model of the object [13]. For specular objects, the highlights on the objects serve as an informative cue for object detection and pose estimation [8]. Similarly, Gremban and Ikeuchi [12] use specular highlights for object recognition, and plan novel views that are discriminative between objects with similar highlights. However, these methods do not generalize to objects with mirror reflectance. In a calibrated setup (camera and environment), it is possible to infer about surface normals through image and environment

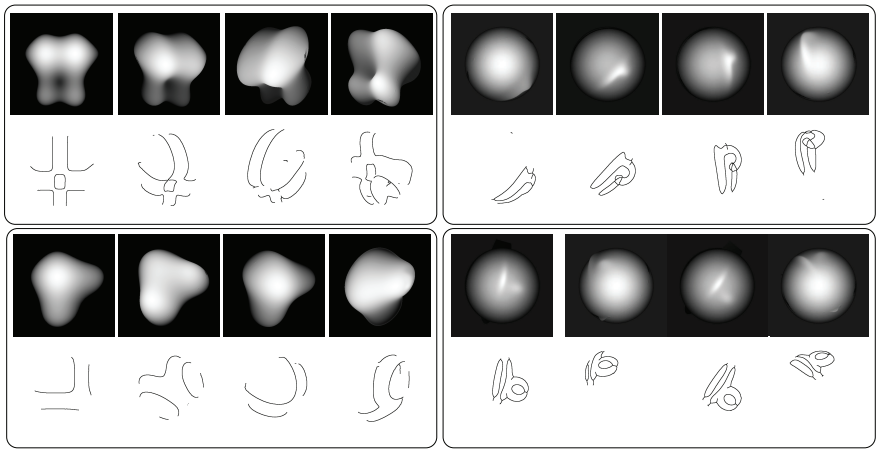

Fig. 5. Parabolic curves provide a unique signature for object pose and identity. Shown are the parabolic curves of four different objects in differnet poses. In each instance, the depth map and the analytically computed parabolic curves are shown.

correspondences which can be further utilized for estimation algorithms [6, 19]. However the required calibration process is a tedious task. Pose estimation and recognition in an uncalibrated setup remains to be a challenge and we show that the proposed image invariants provide necessary information for such tasks.

Pose Estimation: The pose estimation algorithm recovers the 3D rotation and 2D translation parameters with respect to a nominal pose of the object. Since the camera model is assumed orthographic, the object pose can only be recovered up to depth ambiguity. We assume that either the parametric form or the 3D model of the object is given in advance. Based on the representation, the 3D positions of the parabolic points at object coordinates are recovered either analytically (using parametric form) or numerically (using the 3D model). In an offline process, we generate a database of curve templates by rotating the parabolic curvature points with respect to a set of sampled 3D rotations and projecting visible points to the image plane. Since rotation of the object along the principal axes (θ_z) of the camera results in an in-plane rotation of the parabolic points on the image plane, it suffices to include only out-of-plane rotations (θ_x and θ_y) to the database which is performed by uniform sampling of the angles on the 2-sphere. A few samples included into the database is given in Figure 5.

The initial pose of the object is recovered by searching for the database template together with its optimal 2D Euclidean transformation parameters $\mathbf{s} = (\theta_z,\, t_x,\, t_y)$, which aligns the parabolic points of the template to the image parabolic curvature points. We use a variant of chamfer matching technique [2] which measures the similarity of two contours. The precision of the initial pose estimation is limited by the discrete set of out-of-plane rotations included into the database. We refine the estimation using a combination of iterative closest point (ICP) [3] and Gauss-Newton optimization algorithms.

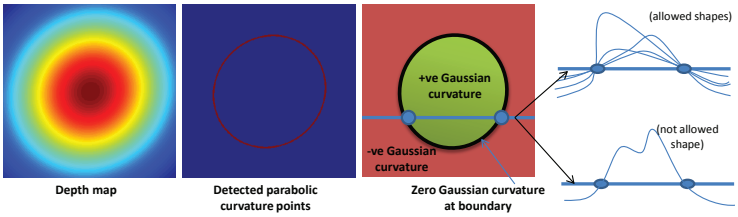

Fig. 6. A stylistic example showing how the parabolic points provide global shape priors. Such priors are extremely useful in restricting the solution space in a surface recovery algorithm as well as in identifying regions where simple parametric models describe the surface accurately.

Recognition: The parabolic curvature points provide unique signatures to recognize many objects in variable poses. The object recognition algorithm is a simple extension of the presented pose estimation approach. For each object class we repeat the pose estimation process and recover the best pose parameters. The object class is then given by the minimum of the chamfer cost function [2] over all classes.

Shape Priors: Knowledge of the parabolic curvature points gives a strong prior on the shape of the mirror. It is well known that curves of parabolic curvature separate regions of elliptic and parabolic curvature. Toward this end, we can constrain the range of possible shapes (Figure 6). Further, in each region we can use simple non-parametric surface models such as splines and regularize their parameters to satisfy the curvature properties. This forms a compelling direction for future research.

6 Experiments

We use both real and synthetically generated images for our experiments. For synthetic experiments, we use publicly available ray-tracing software POV-Ray for photo realistic rendering which provides high quality simulations of real world environments including inter-reflections. Real data was collected with a Canon SLR camera using a 300mm lens, and placing the mirror approximately 150cm from the camera. Both camera and mirror were rigidly mounted to a platform, which was moved around to change the environment features seen on the mirror.

6.1 Detecting Parabolic Curves

In Figure 4, we present results for detection of parabolic curvature points from synthetically rendered images under the ideal imaging condition of orthographic camera and scene at infinity. We show the performance of the detection when these assumptions are violated. Figure 7 shows the detection of parabolic points when the scene is at a finite distance from the mirror. In particular, the detection of parabolic curvature points is reliable even when the minimum distance of the mirror to the object is the same as that of the variations in the depth of the object itself. This shows the stability of the detection statistic to finite scenes. Figure 8 shows stable detection results when the camera is

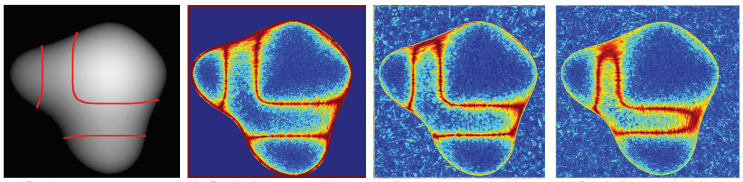

Fig. 7. Detection of parabolic points when the environment is at a finite distance from the mirror.

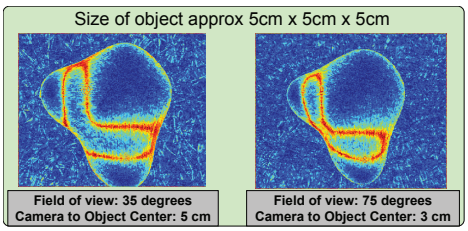

Fig. 8. Detection of parabolic points using a perspective camera under medium to large deviation from the orthographic case. The parabolic curvature points remain stable in both cases. Note that as the camera approaches the object and the field of view of the camera is increased, the relative locations of the (projection of the) parabolic points on the image plane changes. This, in part, explains the drift of the parabolic curvature points.

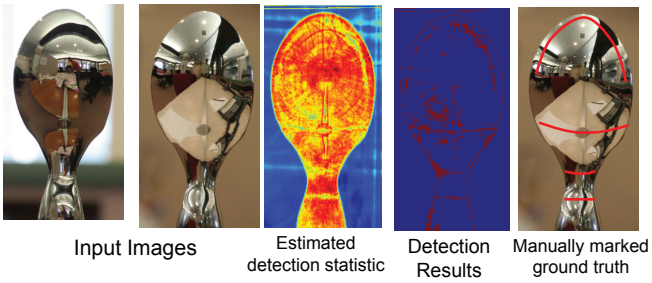

Fig. 9. Estimation of parabolic points of a real object using multiple images. Results were estimated using 17 images.

heavily perspective. These figures reinforce the detection of parabolic curves based on the invariant for practical imaging scenarios. In Figures 9 and 10, we show detection of parabolic curvature points using real images for two highly reflective objects.

6.2 Pose Estimation and Recognition

In the synthetic experiments, we randomly sample six parameters of the 3D object pose and render the object under several environment rotations. The parabolic curves on the

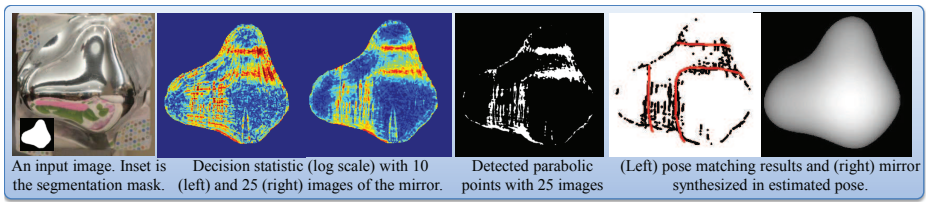

Fig. 10. Estimation of parabolic points of a real surface using variable number of images. The parametric form of the surface was given in Figure 4 and it is manufactured using a CNC machine. As the number of images increase, the degeneracies due to environment become incoherent and detection becomes more reliable.

image plane is detected using the 25 rendered images which are then utilized to recover the object pose via the algorithm described earlier.

We provide results for four different surfaces which are shown in Figure 5. In Figure 11a, we present several pose estimation results. The simulation is repeated 30 times for each surface using a different pose and mean absolute estimation errors for five parameters of the 3D pose is given in Figure 11b. We note that, since the camera model is orthographic, the object pose can be recovered only up to a depth ambiguity. In all our trials the pose estimation algorithm converged to the true pose. As shown, the parabolic curves provide extremely robust features for pose estimation, and average rotation error is less than 2 degrees and 5 pixels. In Figure 10, we show pose detection results on real images of a mirror. The estimated pose was $(-4.24, -1.66, 1.9)$ for a ground truth of $(0, 0, 0)$.

For recognition, we place four objects simultaneously to the environment and recognize identities of these surfaces. This is a challenging scenario due to heavy interreflections of the surfaces. The same rendering scenario of the pose estimation experiment is repeated. The object identities are given via the minimum of the cost function after pose estimation. The average recognition rate over 10×4 trials is 92.5% and typical recognition examples are shown in Figure 12. We note that, two of the surfaces have exactly the same occluding contour, therefore in this scenario this statistic is expected to fail whereas parabolic curvature points provide unique signatures.

7 Conclusions

In this paper, we propose a photometric invariant for images of smooth mirror. We show that images of mirror exhibit degenerate image gradients at parabolic curvature points when the camera is orthographic and the scene is at infinity. We demonstrate the practical effectiveness of the invariant even under deviations from this imaging setup. In particular, the invariant allows for a dense recovery of the point of parabolic curvature from multiple images of the mirror under motion of the environment. This allows us to recover a geometric property of the mirror. We show that recovery of the parabolic curvature points opens up a range of novel applications for mirrors.

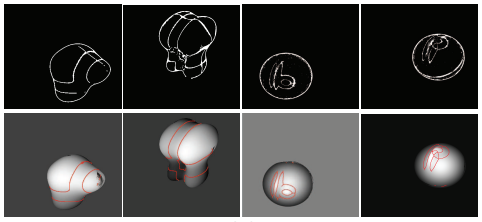

(a)

Mean Abs. Err.	t_x	t_y	θ_x	θ_y	θ_z
Surface-1	1.94	1.60	0.55	0.67	0.46
Surface-2	1.86	2.87	0.96	0.59	0.28
Surface-3	3.54	8.55	2.74	1.04	0.93
Surface-4	4.63	6.69	1.95	1.39	0.48

(b)

Fig. 11. (a) Visualization of the pose estimation results. For each test object, we show the pose estimate at one of the 30 random poses used. (Top) Parabolic curvature points detected from 25 images of the mirror under a rotating environment. (bottom) Estimated pose of the mirror with the true parabolic curvature points overlaid. (b) Mean pose estimation errors. Translational error is in pixels and rotational error is in degrees. The results are averaged over 30 trials.

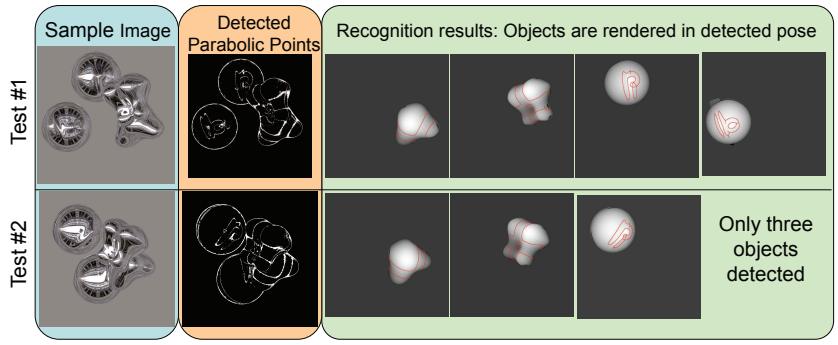

Fig. 12. Recognition experiment on synthetic images. Our test setup consisted of arbitrarily placing all four test objects in a virtual scene and rendering multiple images. The recovered parabolic curvature points were used to recognize the object and estimate its pose.

Acknowledgments. We thank the anonymous reviewers for their feedback in improving the paper. We also thank Jay Thornton, Keisuke Kojima, John Barnwell, and Haruhisa Okuda, Mitsubishi Electric, Japan, for their help and support. Aswin Sankaranarayanan thanks Prof. Richard Baraniuk at Rice University for his encouragement and support.

References

1. Adato, Y., Vasilyev, Y., Ben-Shahar, O., Zickler, T.: Toward a Theory of Shape from Specular Flow. In: ICCV (October 2007)
2. Barrow, H.G., Tenenbaum, J.M., Bolles, R.C., Wolf, H.C.: Parametric correspondence and chamfer matching: two new techniques for image matching. In: Joint Conf. on Artificial Intelligence. pp. 659–663 (1977)
3. Besl, P., McKay, H.: A method for registration of 3-D shapes. TPAMI 14(2), 239–256 (1992)
4. Blake, A.: Specular stereo. In: Int. Joint Conf. on Artificial Intelligence. pp. 973–976 (1985)
5. Blake, A., Brelstaff, G.: Geometry from specularities. In: ICCV. pp. 394–403 (1988)

6. Bonfort, T., Sturm, P.: Voxel carving for specular surfaces. In: ICCV (October 2003)
7. Canas, G.D., Vasilyev, Y., Adato, Y., Zickler, T., Gortler, S., Ben-Shahar, O.: A Linear Formulation of Shape from Specular Flow. In: ICCV (September 2009)
8. Chang, J., Raskar, R., Agrawal, A.: 3D Pose Estimation and Segmentation using Specular Cues. In: CVPR (June 2009)
9. DelPozo, A., Savarese, S.: Detecting specular surfaces on natural images. In: CVPR (June 2007)
10. Ding, Y., Yu., J., Sturm, P.: Recovering specular surfaces using curved line images. In: CVPR (June 2009)
11. Fleming, R.W., Torralba, A., Adelson, E.H.: Specular reflections and the perception of shape. Journal of Vision 4(9), 798–820 (2004)
12. Gremban, K., Ikeuchi, K.: Planning multiple observations for object recognition. IJCV 12(2), 137–172 (1994)
13. Haralick, R., Joo, H., Lee, C., Zhuang, X., Vaidya, V., Kim, M.: Pose estimation from corresponding point data. IEEE Trans. on Systems, Man and Cybernetics 19(6), 1426–1446 (1989)
14. Ihrke, I., Kutulakos, K.N., Lensch, H.P.A., Magnor, M., Heidrich, W., Ihrke, I., Stich, T., Gottschlich, H., Magnor, M., Seidel, H.P.: State of the Art in Transparent and Specular Object Reconstruction. In: IEEE Int. Conf. on Image Analysis and Processing. vol. 12, pp. 188–193 (2005)
15. Koenderink, J., Van Doorn, A.: Photometric invariants related to solid shape. Journal of Modern Optics 27(7), 981–996 (1980)
16. Longuet-Higgins, M.S.: Reflection and refraction at a random moving surface. I. Pattern and paths of specular points. Journal of the Optical Society of America 50(9), 838 (1960)
17. Mamassian, P., Kersten, D., Knill, D.: Categorical local-shape perception. Perception 25, 95–108 (1996)
18. Miyazaki, D., Kagesawa, M., Ikeuchi, K.: Transparent surface modeling from a pair of polarization images. TPAMI 26(1), 73–82 (2004)
19. Savarese, S., Chen, M., Perona, P.: Local shape from mirror reflections. IJCV 64(1), 31–67 (2005)
20. Savarese, S., Fei-Fei, L., Perona, P.: What do reflections tell us about the shape of a mirror? In: Applied perception in graphics and visualization (August 2004)
21. Vasilyev, Y., Adato, Y., Zickler, T., Ben-Shahar, O.: Dense specular shape from multiple specular flows. In: CVPR (June 2008)
22. Waldon, S., Dyer, C.: Dynamic shading, motion parallax and qualitative shape. In: IEEE Workshop on Qualitative Vision. pp. 61–70 (1993)
23. Weidenbacher, U., Bayerl, P., Neumann, H., Fleming, R.: Sketching shiny surfaces: 3D shape extraction and depiction of specular surfaces. ACM Transactions on Applied Perception 3(3), 285 (2006)
24. Yu, J., McMillan, L.: Modelling Reflections via Multiperspective Imaging. In: CVPR (June 2005)
25. Zisserman, A., Giblin, P., Blake, A.: The information available to a moving observer from specularities. Image and Vision Computing 7(1), 38–42 (1989)