ORIGINAL PAPER

Detection and localization of specular surfaces using image motion cues

Ozgur Yilmaz · Katja Doerschner

Abstract Successful identification of specularities in an image can be crucial for an artificial vision system when extracting the semantic content of an image or while interacting with the environment. We developed an algorithm that relies on scale and rotation invariant feature extraction techniques and uses motion cues to detect and localize specular surfaces. Appearance change in feature vectors is used to quantify the appearance distortion on specular surfaces, which has previously been shown to be a powerful indicator for specularity (Doerschner et al. in Curr Biol, 2011). The algorithm combines epipolar deviations (Swaminathan et al. in Lect Notes Comput Sci 2350:508–523, 2002) and appearance distortion, and succeeds in localizing specular objects in computer-rendered and real scenes, across a wide range of camera motions and speeds, object sizes and shapes, and performs well under image noise and blur conditions.

Keywords Specularity detection · Image motion · Surface reflectance estimation

K. Doerschner
National Magnetic Resonance Research Center (UMRAM),
Bilkent Cyberpark, C-Blok Kat 2, 06800 Ankara, Turkey
e-mail: katja@bilkent.edu.tr

K. Doerschner
Department of Psychology, Bilkent University,
06800 Ankara, Turkey

O. Yilmaz (✉)
Department of Computer Engineering, Turgut Ozal University,
06010 Ankara, Turkey
e-mail: ozyilmaz@turgutozal.edu.tr

1 Introduction

Surface reflectance estimation is a fundamental problem in computer vision. Not only is known surface reflectance a prerequisite for the successful recovery of 3D shape [3–8], but it also provides crucial information about the semantic identity of objects (Fig. 1). For an artificial agent, visual extraction of reflectance properties may be a crucial prerequisite for properly planning interactions with the environment, for example in the biomedical context [9], or in real-time, real-world 3D reconstruction scenarios [10,11].

Visual estimation of surface reflectance properties is mathematically under-constrained since the object's reflectance properties, its 3D shape and the illumination have to be estimated simultaneously from the 2D patterns of light arriving at the sensor. Thus, the majority of previous work on reflectance classification has relied on specific assumptions about the spectral BRDF [13–15], knowledge of camera motion [16,17] or has been made under specific conditions, such as specialized sensing or lighting [18].[1] Previously, we showed that reflectance can be rapidly classified based on the statistical differences between the image motion generated by moving diffuse and specular surfaces without any restrictive assumptions [20]. Our approach was inspired by the human visual system which is able to extract surface material properties from single [21–36] and multiple images [1,37–39] with ease, despite the simultaneous shape–material–illumination estimation challenge. More recent evidence suggests that one particular powerful image motion cue that the human visual system seems to be sensitive to when estimating surface reflectance is the distortion of appearance that moving specular surfaces give rise to [1] (Figs. 2, 5).

---

[1] But see [19] who use only minimal assumptions about the scene, motion and 3D shape.

Fig. 1 Surface reflectance, appearance and identity. The shape in these three photographs is the same but the identity of the object changes as a function of its surface reflectance characteristics. Semantic labeling would be impossible on the basis of shape alone. From *left* to *right* ping-pong ball, chrome ball bearing, *black* plastic sphere. To optimally interact with these objects, e.g., to pick them up without breaking or dropping them, visual estimation of surface reflectance is crucial. This figure has been re-printed with permission from [12], copyright John Wiley and Sons, 2013


Fig. 2 Appearance distortion on specular objects. **a** As the object rotates on the platform, the pattern reflected from its surface varies markedly across frames. Appearance distortions tend to be more pronounced at low curvature region [5,40] where the pattern of distortion varies with the sign of the curvature. Examples of appearance distortions on this object are *highlighted* in *red squares* in **a** and extracted and magnified in **b**. Appearance distortion is a powerful cue that the human visual system uses to distinguish specular from diffusely reflecting surfaces [1]. In this paper, we utilize this motion feature to develop an algorithm that detects and localizes specular surfaces in real-world scenes. Photographs by Maarten Wijntjes

Appearance distortion is an attractive motion feature to be extracted by a specularity detecting algorithm because it does not require any assumptions about the object, the illumination or the camera trajectory, it can be computed from just two images of a motion sequence, and it is a robust feature: it consistently occurs on specular surfaces, regardless of object shape, motion, or the reflected environment. These are significant advantages over previous approaches which classified objects as matte or specular solely on the basis of the bi-modality of the image velocity histogram [20]. We show below (Fig. 3) that bi-modality does not predict surface reflectance for complex objects rotating around arbitrary axes, thus its applicability is limited.

We developed an algorithm that relies on scale and rotation invariant feature extraction techniques and uses motion cues to detect and localize specular surfaces in both computer-rendered and real image sequences. Appearance change in feature vectors is used to quantify the appearance distortion on specular surfaces. Our algorithm combines a novel appearance distortion cue and epipolar deviations [2] and succeeds in localizing specular objects across a wide range of camera motions and speeds, object sizes and shapes, and performs well under image noise and blur conditions.

The paper is organized as follows: in Sect. 2 we review related work and introduce the concept of appearance distortion. In Sect. 3 we introduce the algorithm for detection and localization of specularities in image sequences. Section 4 describes the test sets and in Sect. 5 we present the experimental results for various camera motions, camera speeds, object sizes, object shapes, surface reflectance properties in computer-generated scenes, as well as for real indoor and outdoor videos. We end with a brief discussion in Sect. 6.

2 Related work

2.1 Surface reflectance estimation

Surface reflectance is a major factor contributing to an object's appearance (see Fig. 1). To obtain a full description
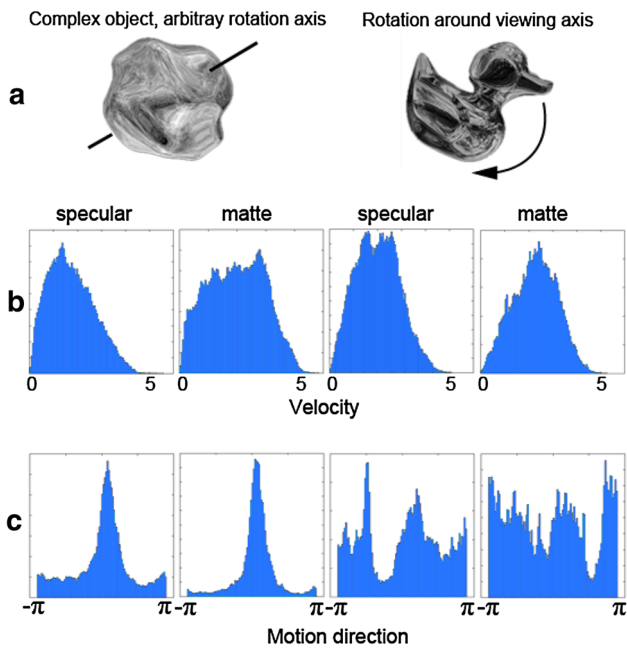
Fig. 3 Velocity and motion histograms. We computed image velocity and motion direction histograms, two potential sources of information which may vary systematically with surface material for two test scenarios: **a** complex shape rotating about an oblique axis and a duck rotating around the viewing axis. Each object was rendered with diffuse (textured) and specular reflectances. **b** There is no apparent difference in the bi-modality of the velocity histograms of shiny and matte surfaces. Velocity corresponds to the magnitude of a flow vector. **c** Also no characteristic differences emerge from for the motion direction histograms

of how an object of a particular material interacts with the illumination one could estimate the bidirectional reflectance distribution function (BRDF) for each point on an object's surface. Apart from systems that require special acquisition devices, there exist image-based acquisition techniques that simultaneously estimate 3D shape and reflectometry from multiple images without the need of special hardware (see [41] for a review). These approaches estimate high-dimensional functions for each point on the object and require very large number of images, hence long computation time. Although these techniques are very useful for model-based rendering and 3D photography, there are robotic vision problems, such as object recognition, or grip force estimation, where a matte–specular binary decision about the object of interest would be sufficient. For such cases, computationally less demanding and more robust techniques would be necessary.

2.1.1 Specular highlight detection

Specular highlights, a special instance of specular reflections, pose problem for methods that use the intensity distribution across an object to recover its 3D shape, since they introduce abrupt and large changes in image intensity around the high-

light region. Thus a large body of research has focused on highlight identification and its segregation from the diffusely reflecting components, which we briefly review.

Assuming different spectral distributions of diffusely and specular reflecting surface material components has proven to be particularly successful in highlight-removal approaches. Shafer [13] introduced the Dichromatic Reflection Model, which approximates the light reflected by a surface point as a linear combination of diffuse and specular components to model this two-spectral component reflection function. The model has been used in several 'flavors' for highlight identification and removal, e.g., by Klinker [42] in combination with a sensor model, by Bajcsy et al. [43] to also segment highlights arising from interreflections between objects, or by Tan et al. [44] and Mallick et al. [45] using a pixel-based techniques to allow highlight segregation for textured surfaces—to name a few.

Light reflected by specular regions is highly polarized while that reflected by diffuse body color is not, thus, an alternative approach to using color information has been to identify specular highlights by looking at the amount of polarization in the reflected light, e.g., [14] and more recently [46]. Nayar et al. [15] combines color and polarization profiles, and Chung et al. [47] proposed an integrative feature-based technique that does not rely on the color signature of diffuse and specular reflectance components.

What this literature has in common is that the specular highlight is treated as a disruption to the 3D shape recovery of a matte object. Specularities, however, can be valuable source of information and constrain, for example, 3D shape recovery if image motion is taken into account[2][48]. Specular feature motion can also be a particularly useful cue for surface reflectance classification, as discussed next.

### 2.2 Image motion and specular surfaces

Image motion has been beneficial in many computer vision problems: structure from motion, image stitching, 3D shape recovery, stereo correspondence, recognition, or pose estimation (see [49] for a review), and has recently received increasing attention in 3D specular shape reconstruction [50–52], specularity detection [2,53], and reflectance classification [1,20].

Bi-modality Previous research suggested that movies of rotating shapes could be classified into specularly-, or diffusely reflecting[3] objects solely on the basis of the shape of the image velocity histogram [20]. This work showed that specular and matte objects give rise to material-specific image velocities. Moreover, for specular objects, the pat-

---

[2] Also see [28] specular shape perception in human observers.

[3] When referring to matte or diffusely reflecting objects, we imply that these objects also have a 2D texture.

Fig. 4 Epipolar deviation vs. appearance distortion: specificity. Moving, diffusely reflecting objects, such as the photograph of the house in the above scene, will distort the optic flow field solely due to camera motion. This will cause epipolar deviations, which a specularity detection algorithm, that is solely based on this cue, may label as specular. **a** An office scene. In addition to camera motion the photograph, located on the desk, moves. **b** *Left* this causes large epipolar deviations (*purple* overlay), thus this region would incorrectly be labeled as specular. *Right* appearance distortion (*purple* overlay) is not affected by this manipulation and correctly labels the specular object. **c** Epipolar deviation and appearance distortion for this sequence

tern of image velocities—in particular the bi-modality of the velocity histogram—strongly depends on the 3D curvature variability of the object. While the classification results correlated well with human performance, the test set used in this study consisted only of simple, cuboidal 3D shapes that rotated from right to left around the vertical axis. Thus, the question arises as to what extent these findings generalize to more generic scenarios. Figure 3 demonstrates that bi-modality of the image velocity distribution does not predict surface reflectance for complex objects rotating around arbitrary axes, thus its applicability may be limited to simple shapes[4].

Epipolar deviations For rigid, diffusely reflecting objects the optic flow due to camera motion obeys epipolar geometry. Swaminathan et al. [2] showed that specular motion violates the epipolar constraint. While these epipolar deviations may signal the existence of specularities [2,53], their usefulness in specular object detection is restricted to linear camera motion and convex shapes. Moreover, problems may arise—due to the global nature of this cue[5]—that diminish its specificity. For example, moving matte, textured objects will distort the optic flow field solely due to camera motion, and would thus cause epipolar deviations (Fig. 4). Moreover, slowly moving, near-planar, specular objects will have negligible epipolar deviations, and may thus not be identified as specular [1]).

2.3 Appearance distortion

As a specular object rotates about its axis (Fig. 5a) specular features, 'rush' towards high curvature points, and appear to become 'absorbed' due to the compression at these locations [28] (Fig. 5b). Additionally, "feature genesis" occurs

---

[4] It may be that bi-modality does not work well as a global parameter; however, at a particular spatial scale it may continue to correctly predict surface reflectance.

[5] At the fundamental matrix estimation stage, motion vectors from the entire image contribute.
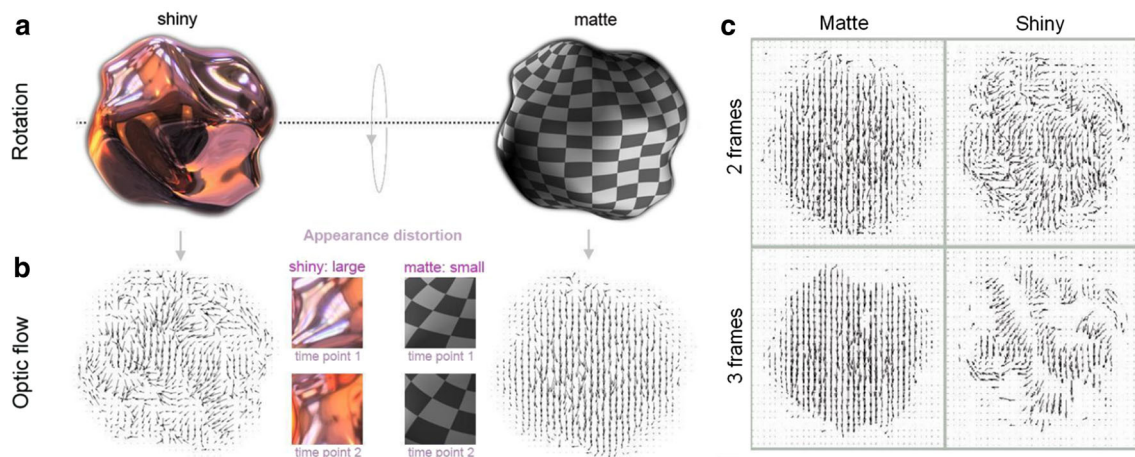
Fig. 5 Appearance distortion and appearance change. **a** A complex shiny (*left*) and a matte, textured object (*right*) are rotating about the *horizontal* axis. **b** This rotation gives rise to distinct flow pattern for each surface material. The shiny object exhibits a marked amount of appearance distortion, i.e., feature absorption and genesis, whereas the appearance of the diffusely reflecting (matte), textured object does not change substantially. **c** In order to assign a motion vector pixels need to be tracked for a period of time. Although increasing integration time reduces noise, it also decreased trackability. Optic flow computed over a 2-frame distance and optic flow computed over a 3-frame distance are shown. Here, trackability diminishes dramatically for the shiny, but not for the matte object. These figures have been reprinted with permission from [1], Copyright 2011 by Cell Press. Because of the current limitations of the applicability of optic algorithms to scenes containing specular flow [51] (see Sect. 5.8), we propose here to measure appearance distortion directly using SIFT features [56]

at local concavities on the object's surface (see also [54] for an analysis on parabolic points, and [55] for specular stereo analyses). Previously, we have shown that the resulting distortion of appearance during object motion impairs the trackability of these features by optic flow mechanisms, i.e., when image features change in appearance too rapidly, they cannot be tracked for a sufficient time interval to estimate their motion. We computed a metric that captures the proportion of image features that are trackable as function of lengthening the frame interval (Fig. 5c), and showed that this metric—which we called coverage change was highly predictive of surface specularity[6].

Although Doerschner et al. [1] found coverage change to be consistently larger for specular objects than matte, textured ones, it is only an indirect measure of appearance distortion and relies on optic flow computations. Adato et al. [51] point out that existing optic flow algorithms are incompatible with fundamental properties of specular flows such as occasional, large specular flow magnitudes which is in direct disagreement with a global smoothness constraint [57]. Changing to a polar representation of optical flow [58] substantially improves the quality of specular flow computation. Alternatively, Adato et al. [48] proposed to estimate specular flow and 3D shape simultaneously, rather than sequentially, thereby improving the estimation of both aspects.

Our aim here is to detect specularities, irrespective of 3D shape. Thus, to avoid the above-mentioned problems in spec-

ular flow estimation, we propose to directly capture appearance distortion using SIFT features [56]. We compute a SIFT feature descriptor that captures local appearance and quantify appearance distortion with the change in matching feature descriptors in consecutive frames. In order to directly compare our SIFT method to an optic flow-based approach, we conducted a control experiment in Sect. 5.8. Our results confirm that an optical flow based-method—although performing sufficiently for classification tasks [1]—is unable to reliably and efficiently detect specular objects.

Appearance distortion vs. epipolar deviation Epipolar deviations are related to localization errors in SIFT features. Other than by specular motion, these are primarily caused by perspective effects, illumination and camera gain changes [59]. Appearance distortion, as it would occur independent of specular motion, primarily arises from errors in the SIFT descriptor [56] which are caused by imaging non-linearities, such as illumination changes, sensor saturation, gamma correction [60,61] as well as perspective effects [62]. While the origins of errors in specularity detection for the two measures overlap, the effects of these variables (illumination, perspective) may alter each cue differently since epipolar deviation is related to the feature detector [56]; whereas, appearance is related to the feature descriptor [56]. Thus, by combining both cues we would significantly boost detection performance (see Fig. 14). In general, however, we find that appearance distortion performs better under conditions that are problematic for epipolar deviations such as independently moving objects (Fig. 4), or image degradations, e.g., added

---

[6] Note that the aim in [1] was to predict human perception. Thus this measure predicts apparent or perceived shininess not physical reflectance.
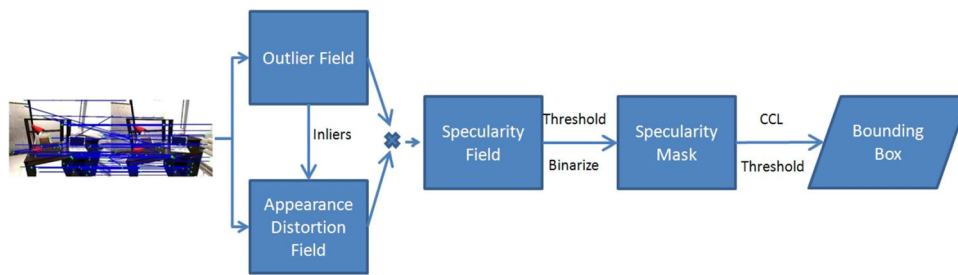
Fig. 6 Flowchart of the algorithm. After applying SIFT, matching outlier and appearance distortion fields are obtained. These two fields are multiplied pixel by pixel to yield the specularity field. A specularity mask is created by thresholding the specularity field. Connected component labeling is used to compute connected regions and these are thresholded according to their sizes. The components that have large regions define bounding *box(es)*

noise and motion blur (Fig. 12a, b). Lastly, unlike previous approaches to specular detection [2,53,63] appearance distortion is not limited to work with linear camera motion and convex surfaces, i.e., would successfully detect also concave specular regions (see inset Fig. 7b), and perform well under non-linear camera trajectories.

3 Algorithm

The specular region detection algorithm that we propose takes a pair of images and outputs bounding boxes of detected specularities. SIFT features [56,64] not only provide scale and rotation invariance in image analysis, but also feature vectors that quantify local appearance. This makes them a perfect tool to directly capture the appearance distortions[7] that occur on specular surfaces, without relying on optical flow.

A flowchart of the algorithm is shown in Fig. 6 and it involves two stages: in stage 1 epipolar deviations and appearance distortion are computed using SIFT features. These two data are combined in the second stage to yield the specularity field. This field is thresholded and specular image regions are extracted using connected component labeling. We evaluate the algorithm's performance by computing precision-recall curves as defined below.

3.1 Stage 1

Epipolar deviations

1. Extract SIFT features for each frame.
2. Eliminate features with low average feature vectors (Appendix A)[8].
3. SIFT nearest neighbor feature matching [66].

4. Apply 2000 RANSAC [67] iterations with 8-point DLT fundamental matrix estimation [68] to matching features.
5. Accept features with Sampson error [69] more than a selected threshold as outliers (Appendix A).
6. Initializing a zero magnitude field (same size as the image).
7. Assign high intensity to outlier pixels and convolve with Gaussian kernel (Appendix A).

This field quantifies the density of epipolar deviations in the image. See Fig. 7a.

Appearance distortion

1. Compute L1 norm of the change in matching feature vectors for consecutive frames (inliers only). L1 norm is chosen over L2 norm, for its relative sensitivity to smaller distances.
2. Initializing a zero magnitude field (same size as the image).
3. Assign appearance distortion values to inlier pixel locations and convolve with the same Gaussian kernel as above.

The resulting field quantifies the appearance distortion in the image; see Fig. 7b. Note that appearance changes for outlier features are not used in order to avoid false high appearance distortion values. The latter arise from errors in feature matching.

3.2 Stage 2

Combination The specularity field is obtained through pixel-by-pixel multiplication of the epipolar deviation and appearance distortion fields (Fig. 7c). Multiplication of two fields acts as an AND operation, but it is more quantitative and flexible since a threshold can be defined on the combined field. This field has high intensity only if both epipolar deviation and appearance distortion exist.

---

[7] SIFT features have also been used for sparse specular surface reconstruction [65].

[8] This is crucial since appearance distortion critically depends on the change in feature vectors.

Fig. 7 Sample Scene. The camera motion is a rotation in azimuth, a single specular object is located on the desk. A sample frame from the image sequences we used in our experiments with epipolar deviation (outlier) field (**a**), appearance distortion field (**b**), and the combined specularity field (**c**) overlaid in *purple*. *Inset* figure in **b** illustrates that appearance distortion, and therefore or algorithm also detects concave specular regions. *Purple squares* indicate specular regions that were detected by the appearance distortion measure. The *white bounding box* in **c** denotes the image region that was labeled as specular by the algorithm

Region extraction A binary specularity mask image is obtained after thresholding, and the bounding box of the specular surface is computed after connected component labeling. Size thresholding is applied on connected components to reject small spurious regions. The bounding box is shown in Fig. 7c, as a white rectangular outline overlaid on the office scene.

3.3 Performance evaluation

We first determined the intersection I of the area delimited by the detected bounding box D with the area delimited by the ground truth bounding box G.

$$I = D \cap G \tag{1}$$

Precision P is the ratio of the intersection area I to the detected area D.

$$P = \frac{I}{D} \tag{2}$$

Recall is the ratio of the intersection area to the ground truth bounding box area.

$$R = \frac{I}{G} \tag{3}$$

Precision is equal to one if the detected bounding box D is completely inside the ground truth bounding box G (Fig. 8b). In this case, a region on the specular object is labeled as specular which is considered sufficient to label the object as specular. Therefore, we will primarily use precision to evaluate the performance of the algorithm for the localization of specular surfaces.

4 Test set

Computer-rendered scenes In order to compare algorithm performance with ground truth we rendered (3ds Max, Copyright 2012 Autodesk, Inc.) 10 frames, 1,600 × 1,200 pixel

Fig. 8 Performance evaluation.
Shown are hypothetical
detection results. **a** Precision P
is the ratio of the intersection
area I to the detected area D
(*white bounding box*), i.e., how
much of the detected region is
actually specular. Recall R is the
ratio of I to the ground truth
bounding box area G, i.e., how
much of the specular object has
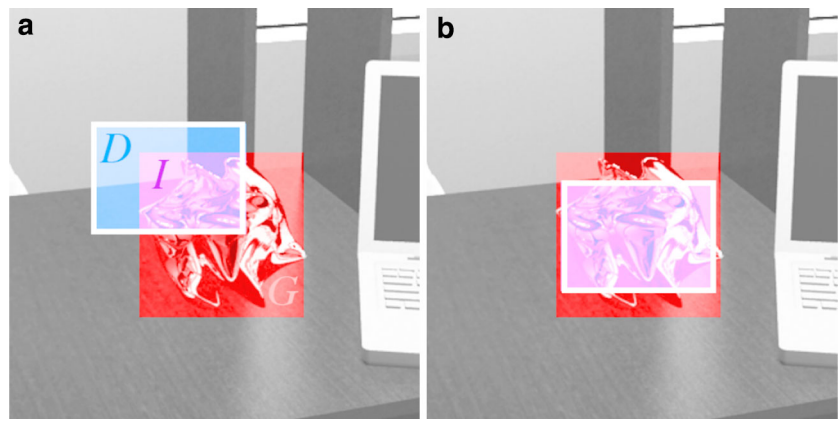been detected. **b** $P = 1$ if D is
completely contained in G

image sequences of a scene set in a study room, that contained
a fully specular (Blinn BRDF [70]), novel object with a mod-
erately undulated 3D surface structure (containing concavi-
ties and convexities, Fig. 7a). We used a computer-rendered
image set in order to have control over all environmental
variables such as the type of camera motion (rotation, trans-
lation, zoom), object size, intensity of the specular reflection
and specular blur, object shape (sphere, ellipsoid, chamfer
box with various degrees of corner roundedness), camera
speed, as well as image noise and blur. The sequences were
evaluated for each subsequent pair of images, and the results
(precision) were averaged for each sequence to give an over-
all performance for different scenarios.
Real-world scenes In order to test the algorithms perfor-
mance in real-world settings, we also captured real image
sequences with a medium-quality camera (resolution $1,600 \times$
$1,200$). We recorded camera trajectories through two indoor
and two outdoor environments that each contained a specu-
lar, undulated sculpture. As a further check of the robustness
of our algorithm in generic non-rigid motion scenarios[9] we
also analyzed the video of a moving person.

All image sequences, rendered and synthetic are available
at our website (http://www.bilkent.edu.tr/~katja/specdec) as
a benchmark set for future studies.

5 Experiments and results

5.1 Camera motion and object size

Previous specular surface detection algorithms [2,53,63]
relied solely on epipolar deviations, which limited their
applicability to linear camera motion. Our SIFT-driven
epipolar deviation feature, however, is expected to give sim-
ilar results across different camera trajectories, since it relies

on the same geometric information. To verify this, we tested
our algorithm for three types of camera motions: rotation,
translation and zoom. In addition, to concurrently examine
the effects of object size on detection quality, we multiplied
the diameter of the largest object by 0.8 and 0.5 to obtain
medium and small-sized objects, respectively. Thus we tested
the algorithm on nine image sequences (three camera motions
$\times$ three object sizes). Precision-recall curves for the three
object sizes, and a rotational camera motion are shown in
Fig. 9. The precision values for all nine conditions are given
in Table 1.

Overall, we see that at higher recall values the precision
of the algorithm drops[10]. The algorithm precisely located
the specular object in the image sequence for the large and
medium-sized object, for each, camera, rotation, translation
and zoom. There was a significant decrease in precision for
the small object, as the number of SIFT features might not
have been sufficient to capture specularity. Consistent with
previous studies [2,53,63], we obtained best results for small-
sized objects with translational camera motion.

For subsequent experiments, we optimized the parame-
ters of the algorithm to give maximal precision for large
and medium objects and all types of camera motions
(Appendix A).

5.2 Surface properties: reflection intensity and blur

We wished to examine the effect of surface reflectance prop-
erties on detection performance, that is, we tested whether the
algorithm is able to detect less-than-perfectly specular sur-
faces. The object size in this experiment was fixed to large and
we used a translational camera motion[11]. Specular intensity
and specular blur were varied by adjusting the reflectance

---

[9] As discussed below: nonrigid and specular motion share similar fea-
tures and may be confused by a classifier; see, for example [1].

[10] Precision-recall curves are obtained by varying a specific threshold
parameter, analogous to ROC curves.

[11] Given the results in experiment 5.1, we did not expect differences in
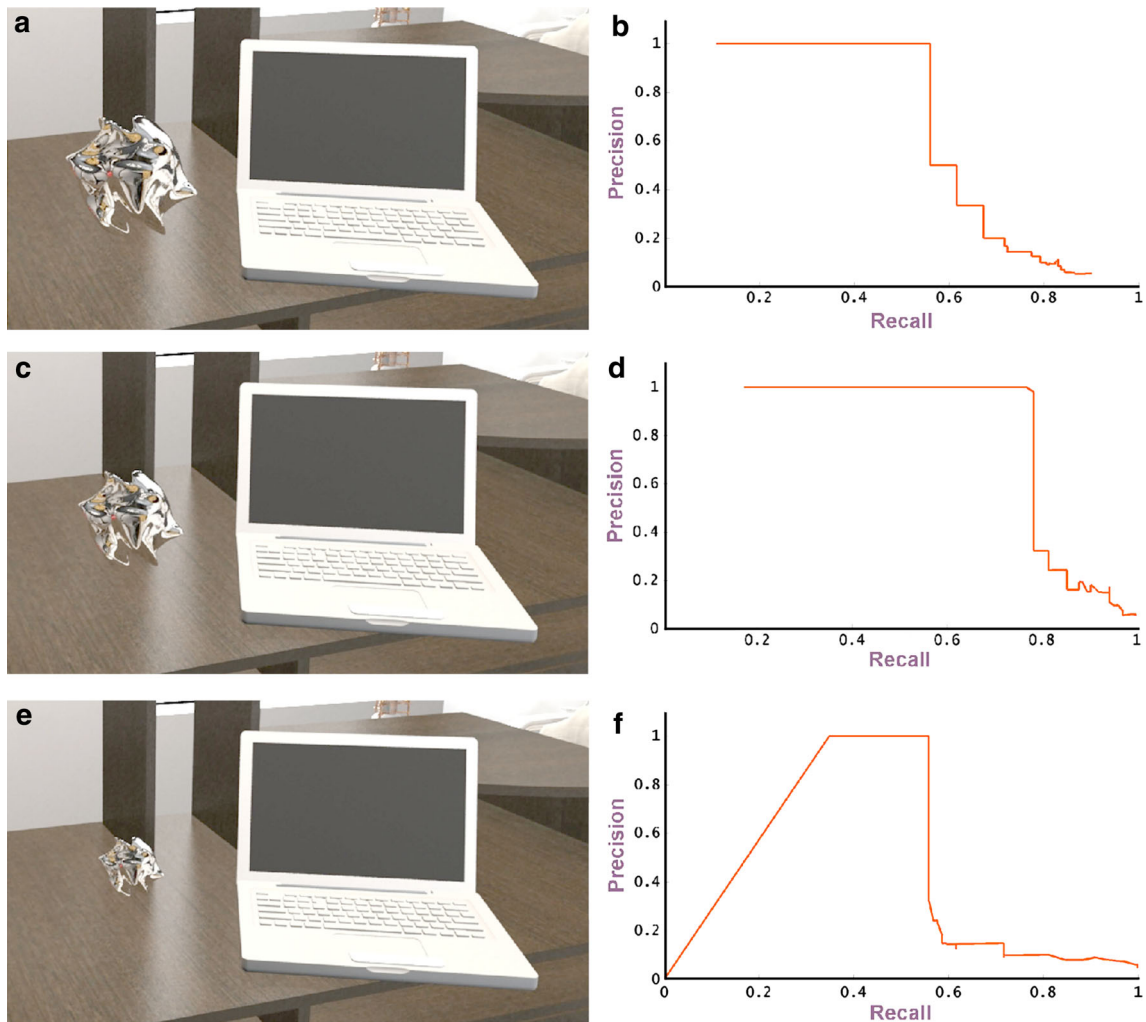performance for rotation and zoom.

Fig. 9 Effect of object size. We tested three different object sizes (**a**, **c**, **e**). Corresponding precision-recall curves are shown on the *left* (**b**, **d**, **f**)

Table 1 Effect of camera trajectory and object size

Size	Camera motion type		
	Rotation	Translation	Zoom
Large	0.94	1	1
Medium	1	1	1
Small	0.22	0.49	0.11

Precision values for three different camera trajectories and three object sizes

[1, 0.25] and glossiness [1, 0.5] parameters of the BLINN model. Increasing specular blur results in a brushed metal appearance, whereas decreasing specular intensity causes the object to appear more glossy, less mirror-like. Renderings and results are shown in Fig. 10. Though our algorithm remains quite success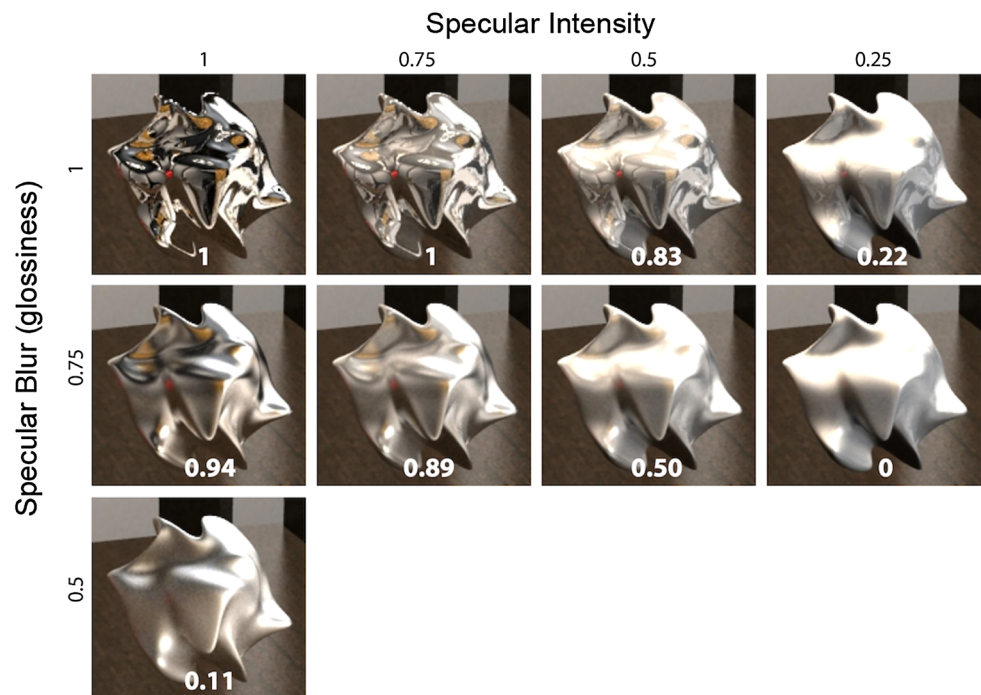ful for small changes in reflectance, we find an overall, gradual decline in precision for decreasing specular intensity and increasing specular blur.

5.3 Object shape

Epipolar deviations in specular surfaces have been shown to peak and then decrease with increasing curvature radius [2]. Appearance distortion is also expected to be weak for large curvature radii, and objects of constant or near constant curvature[12]. Thus, we expected the performance of algorithm to be affected by 3D surface curvature. We tested detection performance for a set of simple objects that varied in their surface curvature: a sphere, ellipsoid and three different chamfer boxes (Fig. 11). The size of these objects was adjusted to correspond to the *large* object (Fig. 9a). In order to examine the

---

[12] Interestingly, it has been shown that such objects tend to be perceived as less shiny by human observers [71].

Fig. 10 Effect of specular intensity and specular blur. Precision values are shown as *insets* in each subpanel. Overall, the manipulation of both parameters decreases performance. We also find that blurring has a stronger detrimental effect

optimum curvature for detection, we varied the curvature of the chamfer box edges by adjusting the fillet parameter to 20, 30, and 40 % of the box size. The camera motion in this experiment was the same as in Exp. 5.1. Precision values for each object are shown in Fig. 11. We found the detection performance for chamfer box to be largest for an intermediate value of curvature suggesting that there is an optimum curvature for specular object detection. As expected, sphere and ellipsoid were not detected by the algorithm.

Interestingly, human observers perceive objects with simple 3D structure as less shiny compared to undulated, complex objects [20]. Thus, our algorithm is rather appearance-based, rather than physics-based[13].

5.4 Camera motion speed

The detection algorithm depends on the change of SIFT feature vectors on consecutive frames. For very fast camera motion the change in feature vectors will be large, whether or not specularities are present in the scene; therefore, the false alarm rate in detection might increase. In addition, there will be smaller number of matching features between frames. However, size thresholding after connected component labeling is expected to reject spurious regions. To examine the sensitivity of the algorithm to speed, we doubled and tripled the camera speed, for rotation, translation and zoom trajectories, and a large object size. For all three types of motion the precision reduced to 0.96 and 0.94 when doubling and tripling

the camera speed, respectively. Thus, the algorithm appears to be robust against increases in rate of change in appearance. In contrast, an optical flow-based approach would be highly sensitive to camera motion speed; see Sect. 5.8.

5.5 Image noise and motion blur

Rendered images do not suffer from the noise or blurring effects that are inherent in real video sequences, however, SIFT features have shown to have some tolerance to noise and blur [56]. Thus we wished to test the performance of our detection algorithm under more real-life like situations. To this end we introduced additive Gaussian noise (σ, 0–29) and motion blur (length, 0–32 pixels) to our image sequences, varying noise power and motion blur length (Fig. 12). Noise was added to the image before applying motion blur. This is a harder case than the opposite sequence since it creates structured noise in the image. Camera trajectory and object size were the same as in Exp. 5.1. Precision values for this set are shown in Fig. 12. The detection performance breaks down beyond 32 pixels of motion blur, but seems to be robust against additive noise.

We also examined the effect of noise and blur on the simpler object geometries used in Exp. 5.3. Gaussian, additive noise ($\sigma = 7$) and 32 pixel-length motion blur were applied to the chamfer box sequences (Fig. 12). Precision for these sequences drops to zero, whereas it was 33 % for the complex object, suggesting that a more undulated shape may be less sensitive to noise and blur distortions.

---

[13] We suggest below that by complementing our motion-based features with static cues to specularity, e.g., [19], also simple 3D specular shapes may be detected.
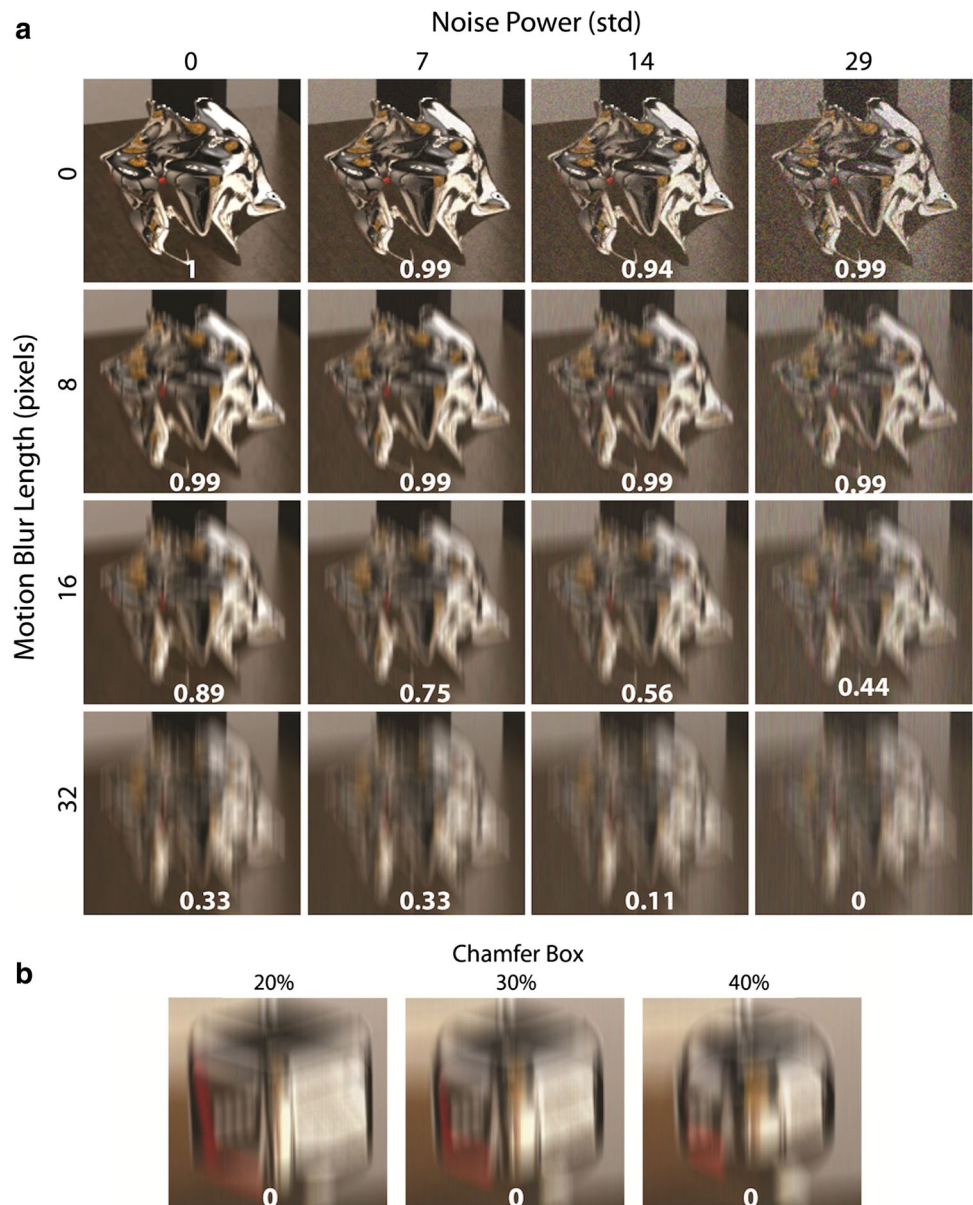
Fig. 11 Effect of 3D shape. Specular object shapes tested: *sphere*, *ellipsoid* and *chamfer box* with three different edge curvatures. Precision of the detection algorithm are shown as *inset* in each subpanel

Fig. 12 Effect of image noise and blur. **a** Shown are frames of a specular object under different additive image noise powers and motion blur length conditions. Precision of the detection algorithm is given as *inset* in every subpanel. **b** Chamfer boxes with added noise and blur ($\sigma = 7$ noise and 32 pixels motion blur). The precision values are all zero as opposed to 33% for the complex object in **a**

5.6 Individual contributions of cues

The specularity field is generated by combining two potential cues for detecting specularity: epipolar deviation which has been suggested in the pervious studies [2,53,63] and appearance distortion, a new measure that we have proposed in this paper. We have noted in Sect. 2.3 that there exist differences between these cues, thus it important to understand the contribution of each cue to detecting specular surfaces. Figure 13 shows the precision-recall curve for the algorithm,
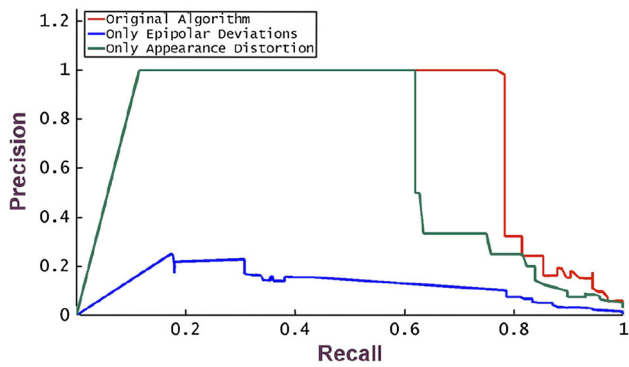
Fig. 13 Individual contributions of cues. A Precision-Recall curve for the algorithm (*red*) and individual cues (*blue* epipolar deviation, *green* appearance distortion). The specular object was of medium size, the camera trajectory, rotation in azimuth

and individual cues for a medium-sized object and a rotating camera motion. While this figure highlights that the appearance distortion cue tends to outperform epipolar deviations, we would like to note that only by combining the cues we significantly boost algorithm performance, as discussed in 2.3 (also see Fig. 14).

5.7 Real video sequences

For real-world experiments we purchased an approximately 30-cm-high specular sculpture. The algorithm had to find this object in four real-world scenes: indoor (office, hallway) and outdoor (street, building). Note, that these movies were taken under everyday conditions (handheld camera, accidental trajectories, etc.), thus should constitute a hardest-case test scenario. To make our tests highly stringent we left the parameters of the algorithm unchanged, i.e., they were optimized for performance in computer-rendered scenes.

Figure 15 shows detection in an indoor and an outdoor scene, as well as the corresponding epipolar deviation-, appearance distortion-, and combined specularity fields. Results for the whole dataset and bounding box videos of detection are available at http://www.bilkent.edu.tr/~katja/specdec. The algorithm is able to detect specular objects in real indoor and outdoor settings, suggesting that performance is quite robust across different environments. Detection success is comparable to synthetic scenes, even though the false alarm rate has increased, probably due to perspective effects, noise, and repetitive patterns. Table 2 shows the precision-recall values for all indoor and outdoor scenes.

Nonrigid motion Nonrigid motion and specular flow have both high levels of appearance distortion and epipolar deviations, thus it is possible that nonrigid motion, such as biological motion may be mis-detected by the algorithm. Yet, we suspected that the appearance distortion caused by specular flow may be higher than those caused by non-rigidly moving agents. To test this issue we also analyzed a real videos


Fig. 14 Performance boost through combination of cues. The combination of epipolar deviation row (**a**) and appearance distortion (**b**) significantly enhances specularity detection performance of the algorithm

(**c**). Epipolar deviation-, appearance distortion-, and combined specularity fields are shown for translational and zoom camera trajectories

|	Epipolar deviations	Appearance distortion	Combination

Fig. 15 Real video sequences. Two frames of indoor (hallway) and outdoor (building) videos are shown in the left most column, along with corresponding epipolar deviation-, appearance distortion-, and combined specularity fields (from *left* to *right*). The *white bounding box* overlaid on the original frame (*left* column) show that the specular object was successfully localized using the combined cue. Color codes are as in Fig. 4. Note, that specular surfaces with constant or low curvature, such as the metal hand rails in the indoor scene or the windows in the outdoor scene, have low appearance distortion and were thus not consistently labeled as specular by the algorithm. Also see Fig. 11. for a general effect of 3D shape on algorithm performance

Table 2 Performance of detection for real images

Scene	Detection Performance	
	Precision	Recall
Hallway	1	0.67
Office	0.4	0.82
Building	0.24	0.54
Street	0.21	0.21

Precision-recall values for indoor and outdoor scenes. Corresponding bounding box videos of detection can be found at http://www.bilkent.edu.tr/~katja/specdec

of a person walking. Figure 16 shows epipolar deviation-, appearance distortion-, and combined specularity fields for one nonrigid sequences. Additional results can be found at http://www.bilkent.edu.tr/~katja/specdec). No bounding box could be estimated for these scenes, thus the algorithm correctly rejected all non-rigid scenarios. Note, however, that a possibility remains that an object which moved into the scene in between the acquisition of two images, might cause sub-

stantial appearance distortion and might thus be mislabeled as specular.

5.8 Optical flow-based detection

It is possible that an algorithm based on optic flow may perform as well in specularity detection, given previous successes of using optic flow-based features for specularity classification [1]. Yet, classification[14] is a rather distinct task from detection and localization, thus the success of optic flow-based features in the former task may not predict success in the latter two. To make this test explicit we used previously identified optical flow field features for specularity detection. Parameters of the optical flow computation [72] and epipolar deviation were identical to [1].

Figure 17 shows the result for one of the computer-rendered image sequences using optic flow features for specularity detection. Evidently, an algorithm based on optic flow fails to correctly localize the specular object. Success of the

---

[14] In [1] images to be classified as matte or shiny contained only a single object and a black background.

Fig. 16 Real video sequences–nonrigid motion. We tested whether the video of a walking person causes false alarms in specularity detection, due to the potentially substantial amount of appearance distortion and epipolar deviations. Yet, the algorithm proved robust under this condition and did not label any regions as specular in these video sequences, i.e., the regions of the filtered specular probability field shown in the right-most column did not pass the thresholding stage (Sect. 3.2.) and thus, no bounding box could be estimated. Also see http://www.bilkent.edu.tr/~katja/specdec for corresponding bounding *box* videos


Fig. 17 Optical flow performance. We adapted the optical flow method proposed in [1] to detect and localize specular objects. Shown is a sample frame from the computer-rendered test set with an overlaid epipolar deviation field, that exhibits very poor specificity. In addition, appearance distortion—based on optic flow—could not be computed reliably. Thus, an algorithm based on optic flow fails to correctly localize specular objects (compare performance to Fig. 7)

epipolar deviation computation is highly dependent on the localization accuracy of SIFT features in order to discriminate matte and specular regions. In fact, the epipolar deviation threshold (Sampson error = 0.02) used in this study was rather stringent compared to other uses of SIFT features, e.g., in 3D reconstruction or visual odometry. Optical flow features, however, were simply not accurate enough for this task, even though the optic flow algorithm itself is recent and with good reported performance [72]. Consequently, the from optic flow-obtained epipolar deviation field was not specific enough for specularity detection.

Appearance distortion based on optical flow suffers from similar problems as the epipolar deviation computation. In particular, increases in camera speed[15] will compromise the optical flow computation, due to an increased difficulty of identifying corresponding pixels, and will thus give rise to an inflated appearance distortion as illustrated in Fig. 17.

---

[15] Compared to the video sequences in [1].

6 Summary and discussion

Surface reflectance is a major factor contributing to an object's appearance, and estimation of surface reflectance is a fundamental problem in computer vision. Recently, image motion has been shown to provide useful information for reflectance classification [1,20] and specularity detection [2,53]. Previous work suggested that there may be several specific image motion cues that signal the existence of specular surfaces [1]. Here, we developed a novel algorithm for the detection and localization of specular objects using image motion, that combines appearance distortion, a novel, particularly strong cue to surface specularity, and epipolar deviations [2].

We have shown that appearance distortion tends to be the more robust image motion feature. Yet, only by combining it with epipolar deviations we significantly improve detection performance. We explain this effect with the complimentary origins and orthogonal visual concepts of these motion cues: feature detection and description. Feature descriptor match-

ing is an essential pre-processing stage for epipolar consistency check and robust fundamental matrix estimation, while epipolar deviations are mainly caused by the localization error during feature position estimation in the feature detector [59]. Therefore, position errors are the main source of deviations. Appearance distortion, however, is directly related to the feature descriptor [60–62], and how it is changing from one frame to the other. How one defines a feature location and how one describes the assigned feature are, as pointed out before, two orthogonal concepts. In fact, it is possible to cross-match different detection and description algorithms [corner feature+image patch, corner feature + histogram of gradients (HOG) descriptor, difference of Gaussians (DoG) feature+image patch, etc.] to generate new detect–describe–match methodologies. In this work, we exploited both of these features and fused them to obtain a reliable specularity detector.

We tested the algorithm under a wide range of conditions. Performance was excellent for all types of camera trajectories (translation, zoom, rotation) and speeds (doubled, tripled). Furthermore, we measured detection performance for different object sizes, less-than-specular reflectance types, various object shapes, as well as under additive noise and motion blur conditions. We find that the algorithm is remarkably robust under noise conditions smaller changes in reflectance. We showed detection to be problematic for small-sized objects as well as spherical shapes. The results on very different object variations, noise, blur, etc., showed the generalization power of the algorithm. Though we tested the algorithm for just one type of rendered scene (office), we think that our results are general as long as there are enough features in the scene to compute fundamental matrix. Our experiments with real videos support this: specular objects were successfully localized in a variety of indoor and outdoor scenes and under varying lighting conditions—without adjusting algorithm parameter for real-world scenarios. However, we found that the precision of the algorithm dropped substantially compared to computer-rendered scenes, probably due to perspective effects, intensity variations and repetitive patterns. The precise influences of these factors should be investigated in future set of experiments.

The algorithm parameters were adjusted conservatively such that the source of precision errors was limited to undetected specular surfaces, rather than false alarm detections. Furthermore, the parameters were optimized for large- and medium-sized objects, with perfect specular (mirror) reflectance and no image noise/blur. Even though the algorithm is not very sensitive to the choice of parameters, in principle, a distinct set of parameters could be obtained to optimize performance in each of the experiments in Sect. 5. A preprocessing step that estimates image noise, blur and camera motion might be beneficial for optimal parameter adjustment. Specifically, real-world outdoor scene sequences

would benefit from a change in parameters. Lastly, the algorithm detected the specular surface in at least one of the frames for nearly all of the tested image sequences. Hence, the algorithm's detection performance may be boosted if the detection results are integrated over multiple frames.

Previous research has shown that moving specular surfaces with small surface curvature variability (ellipsoids) tend to be perceived as matte and non-rigid by human observers [20]. Moreover, appearance distortion, which relies on distinct image motion patterns, has been shown to correlate well with observers' percepts of surface shininess [1], thus is not surprising that we find that our algorithm, which relies on appearance distortion, fails to detect these simple shapes. Surface curvature is a key factor for generating reflectance-specific patterns of image motion [20,50–52], thus simple specular objects that lack surface curvature complexity will not generate these characteristic patterns, thus become "invisible" to the algorithm. We suggest that by combining image motion and static cues to specularity [19,28] we may solve this detection problem in the future[16].

Appendix

Algorithm parameters

- SIFT peak threshold = 3
- SIFT edge threshold = 10
- SIFT feature elimination threshold = 5
- SIFT matching threshold = 2
- RANSAC iteration = 2,000
- Sampson error = 0.02
- Convolution kernel size = 60
- Convolution kernel, Gaussian standard deviation = 30
- Specular field threshold = 1.5×10^{-6}
- Connected component area threshold = 1,000

Optic flow experiments

For the optical flow-based detection experiment, we kept the parameters the same as in [1]. However, we used 5% of the

---

[16] Specular highlights have been suggested as robust features for matching between 2D images and object's 3D representation for pose estimation [73]. This suggests that highlights may also be useful for specular object detection.

optical flow vectors for epipolar deviation computation. The Sampson error, kernel size and standard deviation are identical to the ones used for the SIFT-based method.

References

1. Doerschner, K., Fleming, R., Yilmaz, O., Schrater, P., Hartung, B., Kersten, D.: Visual motion and the perception of surface material. Curr. Biol. **21**(23), 2010–2016 (2011)
2. Swaminathan, R., Kang, S., Szeliski, R., Criminisi, A., Nayar, S.: On the motion and appearance of specularities in image sequences. Lect. Notes Comput. Sci. **2350**, 508–523 (2002)
3. Horn, B.: Shape from shading: A method for obtaining the shape of a smooth opaque object from one view (1970)
4. Horn, B.: Shape from Shading Information. McGraw-Hill, New York (1975)
5. Koenderink, J., Van Doorn, A.: Photometric invariants related to solid shape. Optica Acta **27**, 981–996 (1980)
6. Pentland, A.: Shape information from shading: a theory about human perception. Spatial Vis **4**, 165–182 (1989)
7. Ihrke, I., Kutulakos, K., Magnor, M., Heidrich, W.: EUROGRAPHICS 2008 STAR—state of the art report state of the art in transparent and specular Oobject reconstruction (2008)
8. Wang, Z., Huang, X., Yang, R., Zhang, Y.: Measurement of mirror surfaces using specular reflection and analytical computation. Mach. Vis. Appl. **24**, 289–304 (2013)
9. Saint-Pierre, C.-A., Boisvert, J., Grimard, G., Cheriet, F.: Detection and correction of specular reflections for automatic surgical tool segmentation in thoracoscopic images. Mach. Vis. Appl. **22**, 171–180 (2011)
10. Izadi, S., Kim, D., Hilliges, O., Molyneaux, D., Newcombe, R., Kohli, P., Shotton, J., Hodges, S., Freeman, D., Davison, A., Fitzgibbon, A., Kinectfusion : Real-time 3d reconstruction and interaction using a moving depth camera. In: Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology, UIST '11, ACM, New York, pp. 559–568 (2011)
11. Newcombe, R.A., Davison, A.J., Izadi, S., Kohli, P., Hilliges, O., Shotton, J., Molyneaux, D., Hodges, S., Kim, D., Fitzgibbon, A.: Kinectfusion: real-time dense surface mapping and tracking. In: 2011 10th IEEE international symposium on Mixed and Augmented Reality (ISMAR), pp. 127–136
12. Dörschner, K.: Image Motion and the Appearance of Objects. McGraw-Hill, New York (1975)
13. Shafer, S.: Using color to separate reflection components. Color **10**, 210–218 (1985)
14. Wolff, L., Boult, T.: Constraining object features using a polarization reflectance model. IEEE Trans. Pattern Anal. Mach. Intell. **13**, 635–657 (1991)
15. Nayar, S., Fang, X., Boult, T.: Removal of specularities using color and polarization. In: 1993 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Proceedings CVPR'93, pp. 583–590 (1993)
16. Oren, M., Nayar, S.: A theory of specular surface geometry. Int. J. Comput. Vis. **24**, 105–124 (1997)
17. Roth, S., Black, M.: Specular flow and the recovery of surface structure. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Vol. 2, IEEE, pp. 1869–1876 (2006)
18. Nayar, S., Ikeuchi, K., Kanade, T.: Determining shape and reflectance of Lambertian, specular, and hybrid surfaces using extended sources. In: International Workshop on Industrial Applications of Machine Intelligence and Vision, IEEE, pp. 169–175
19. DelPozo, A., Savarese, S.: Detecting specular surfaces on natural images. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition CVPR '07, pp. 1–8
20. Doerschner, K., Kersten, D., Schrater, P.: Rapid classification of specular and diffuse reflection from image velocities. Pattern Recognit. **44**, 1874–1884 (2011)
21. Ho, Y., Landy, M., Maloney, L.: How direction of illumination affects visually perceived surface roughness. J. Vis. **6**, 8 (2006)
22. Doerschner, K., Boyaci, H., Maloney, L.: Estimating the glossiness transfer function induced by illumination change and testing its transitivity. J. Vis. **10**(4), 1–9 (2010)
23. Doerschner, K., Maloney, L., Boyaci, H.: Perceived glossiness in high dynamic range scenes. J. Vis. **10** (2010)
24. te Pas, S., Pont, S.: A comparison of material and illumination discrimination performance for real rough, real smooth and computer generated smooth spheres. In: Proceedings of the 2nd symposium on Applied perception in graphics and visualization, ACM New York, USA, pp. 75–81
25. Nishida, S., Shinya, M.: Use of image-based information in judgments of surface-reflectance properties. J. Opt. Soc. Am. A Opt. Image Sci. Vis. **15**, 2951–2965 (1998)
26. Dror, R., Adelson, E., Willsky, A.: Estimating surface reflectance properties from images under unknown illumination. In: Human Vision and Electronic Imaging VI, SPIE Photonics West, pp. 231–242 (2001)
27. Matusik, W., Pfister, H., Brand, M., McMillan, L.: A data-driven reflectance model. ACM Trans. Graph. **22**, 759–769 (2003)
28. Fleming, R., Torralba, A., Adelson, E.: Specular reflections and the perception of shape. J. Vis. **4**, 798–820 (2004)
29. Motoyoshi, I., Nishida, S., Sharan, L., Adelson, E.: Image statistics and the perception of surface qualities. Nature (London) **447**, 206–209 (2007)
30. Vangorp, P., Laurijssen, J., Dutré, P.: The influence of shape on the perception of material reflectance. ACM Trans. Graph. (TOG) **26**, 77-es (2007)
31. Olkkonen, M., Brainard, D.: Perceived glossiness and lightness under real-world illumination. J. Vis. **10**, 5 (2010)
32. Kim, J., Anderson, B.: Image statistics and the perception of surface gloss and lightness. J. Vis. **10**(9), 3 (2010)
33. Marlow, P., Kim, J., Anderson, B.: The role of brightness and orientation congruence in the perception of surface gloss. J. Vis. **11**, 16 (2011)
34. Kim, J., Marlow, P., Anderson, B.: The perception of gloss depends on highlight congruence with surface shading. J. Vis. **11**, 4 (2011)
35. Zaidi, Q.: Visual inferences of material changes: color as clue and distraction. Wiley Interdiscip. Rev. Cogn. Sci. **2**(6), 686–700 (2011)
36. te Pas, S., Pont, S., van der Kooij, K.: Both the complexity of illumination and the presence of surrounding objects influence the perception of gloss. J. Vis. **10**, 450–450 (2010)
37. Hartung, B., Kersten, D.: Distinguishing shiny from matte. J. Vis. **2**, 551–551 (2002)
38. Sakano, Y., Ando, H.: Effects of self-motion on gloss perception. Perception **37**, 77 (2008)
39. Wendt, G., Faul, F., Ekroll, V., Mausfeld, R.: Disparity, motion, and color information improve gloss constancy performance. J. Vis. **10**, 7 (2010)
40. Blake, A.: Specular stereo. In: Proceedings of the International Joint Conference on Artificial Intelligence, pp. 973–976
41. Weyrich, T., Lawrence, J., Lensch, H., Rusinkiewicz, S., Zickler, T.: Principles of appearance acquisition and representation. Found. Trends Comput. Graph. Vis. **4**, 75–191 (2009)
42. Klinker, G., Shafer, S., Kanade, T.: A physical approach to color image understanding. Int. J. Comput. Vis. **4**, 7–38 (1990)

43. Bajcsy, R., Lee, S., Leonardis, A.: Detection of diffuse and specular interface reflections and inter-reflections by color image segmentation. Int. J. Comput. Vis. **17**, 241–272 (1996)

44. Tan, R., Ikeuchi, K.: Separating reflection components of textured surfaces using a single image. IEEE Trans. Pattern Anal. Mach. Intell. **27**(2), 178–193 (2005)

45. Mallick, S.P., Zickler, T., Belhumeur, P.N., Kriegman, D.J.: Specularity removal in images and videos: a pde approach. In: Computer Vision-ECCV 2006, pp. 550–563, Springer, Berlin (2006)

46. Angelopoulou, E.: Specular highlight detection based on the fresnel reflection coefficient. In: EEE 11th International Conference on Computer Vision. ICCV 2007. I, pp. 1–8 (2007)

47. Chung, Y., Chang, S., Cherng, S., Chen, S.: Dichromatic reflection separation from a single image. Lect. Notes Comput. Sci. **4679**, 225 (2007)

48. Adato, Y., Ben-Shahar, O.: Specular flow and shape in one shot. In: BMVC, pp. 1–11

49. Szeliski, R.: Computer vision: algorithms and applications. Springer, New York (2010)

50. Adato, Y., Vasilyev, Y., Ben Shahar, O., Zickler, T.: Toward a theory of shape from specular flow. In: ICCV07, pp. 1–8

51. Adato, Y., Zickler, T., Ben-Shahar, O.: Toward robust estimation of specular flow. In: Proceedings of the British Machine Vision Conference, p. 1 (2010)

52. Vasilyev, Y., Adato, Y., Zickler, T., Ben-Shahar, O.: Dense specular shape from multiple specular flows. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR, pp. 1–8 (2008)

53. Oo, T., Kawasaki, H., Ohsawa, Y., Ikeuchi, K.: The separation of reflected and transparent layers from real-world image sequence. Mach. Vis. Appl. **18**, 17–24 (2007)

54. Adato, Y., Vasilyev, Y., Zickler, T., Ben-Shahar, O.: Shape from specular flow. IEEE Trans. Pattern Anal. Mach. Intell. **32**, 2054–2070 (2010)

55. Blake, A., Bulthoff, H.: Shape from specularities: computation and psychophysics. Philos. Trans. Soc. Lond. Ser. B Biol. Sci. **331**, 237–252 (1991)

56. Lowe, D.: Distinctive image features from scale-invariant keypoints. Int. J. Comput. Vis. **60**, 91–110 (2004)

57. Horn, B.K., Schunck, B.G.: Determining optical flow. Artif. Intell. **17**, 185–203 (1981)

58. Adato, Y., Zickler, T., Ben-Shahar, O.: A polar representation of motion and implications for optical flow. In: IEEE Conference on IEEE Computer Vision and Pattern Recognition (CVPR), pp. 1145–1152 (2011)

59. Cordes, K., Muller, O., Rosenhahn, B., Ostermann, J.: Half-sift: high-accurate localized features for sift. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2009. CVPR Workshops. IEEE Computer Society, pp. 31–38 (2009)

60. Toews, M., Wells, W.: Sift-rank: ordinal description for invariant feature correspondence. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2009, pp. 172–177 (2009)

61. Farid, H., Popescu, A.: Blind removal of lens distortions. J. Opt. Soc. Am. **18**, 2072–2078 (2001)

62. Clark, A., Grant, R., Green, R.: Perspective Correction for improved visual registration using natural features. In: 23rd International Conference Image and Vision Computing New Zealand (IVCNZ 2008). IEEE Computer Press, Los Alamitos (2008)

63. Szeliski, R., Avidan, S., Anandan, P.: Layer extraction from multiple images containing reflections and transparency. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Vol. 1, pp. 246–253. IEEE, USA (2000)

64. Vedaldi, A., Fulkerson, B.: Vlfeat : an open and portable library of computer vision algorithms. In: Proceedings of the international conference on Multimedia, ACM, pp. 1469–1472

65. Sankaranarayanan, A.C., Veeraraghavan, A., Tuzel, O., Agrawal, A.: Specular surface reconstruction from sparse reflection correspondences. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, pp. 1245–1252 (2010)

66. Beis, J., Lowe, D.: Shape indexing using approximate nearest-neighbour search in high-dimensional spaces. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 1000–1006 (1997)

67. Fischler, M., Bolles, R.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM **24**, 381–395 (1981)

68. Hartley, R., Gupta, R., Chang, T.: Stereo from uncalibrated cameras. In: Proceedings CVPR'92, IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 761–764 (1992)

69. Sampson, P.: Fitting conic sections to. Comput. Graph. Image Process. **18**, 97–108 (1982)

70. Blinn, J.: Models of light reflection for computer synthesized pictures. In: ACM SIGGRAPH Computer Graphics, Vol. 11, ACM, pp. 192–198

71. Doerschner, K., Kersten, D., Schrater, P.: Analysis of shape-dependent specular motion—predicting shiny and matte appearance. J. Vis. **8**, 594–594 (2008)

72. Gautama, T., Van Hulle, M.: A phase-based approach to the estimation of the optical flow field using spatial filtering. IEEE Trans. Neural Netw. **13**, 1127–1136 (2002)

73. Netz, A., Osadchy, M.: Using specular highlights as pose invariant features for 2d–3d pose estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, pp. 721–728 (2011)

Ozgur Yilmaz received his B.Sc. degree in electrical engineering from Bilkent University, Turkey, in 2003 and Ph.D. degree in electrical and computer engineering in 2007 from University of Houston, USA. Dr. Yilmaz worked as a research scientist in Aselsan Inc., Turkey, for 5 years, where he developed computer vision and image processing algorithms. He had a postdoctoral fellowship in Bilkent University, in which he worked on cognitive vision and medical imaging. Dr. Yilmaz is currently a faculty of engineering in Turgut Ozal University, Ankara, Turkey. His research interests include computer vision, neural networks, cognitive neuroscience and neuromorphic computing.

Katja Doerschner is assistant professor of psychology at Bilkent University and the National Magnetic Resonance Research Center. She received her Ph.D. in experimental psychology from New York University. Her research interests include computational and biological vision, with a focus on the perception of surface qualities.