

Shading Models for Illumination and Reflectance Invariant Shape Detectors

Peter Nillius

Department of Physics
Royal Institute of Technology (KTH)
SE-106 91 Stockholm, Sweden
nillius@mi.physics.kth.se

Josephine Sullivan

Computer Science & Communication
Royal Institute of Technology (KTH)
SE-100 44 Stockholm, Sweden
sullivan@csc.kth.se

Antonis Argyros

Institute of Computer Science - FORTH
Vassilika Vouton, P.O. Box 1385
GR-711 10, Heraklion, Greece
argyros@ics.forth.gr

Abstract

Many objects have smooth surfaces of a fairly uniform color, thereby exhibiting shading patterns that reveal information about its shape, an important clue to the nature of the object. This paper explores extracting this information from images, by creating shape detectors based on shading.

Recent work has derived low-dimensional models of shading that can handle realistic unknown lighting conditions and surface reflectance properties. We extend this theory by also incorporating variations in the surface shape. In doing so it enables the creation of very general models for the 2D appearance of objects, not only coping with variations in illumination and BRDF but also in shape alterations such as small scale and pose changes. Using this framework we propose a scheme to build shading models that can be used for shape detection in a bottom up fashion without any a priori knowledge about the scene.

From the developed theory we construct detectors for two basic shape primitives, spheres and cylinders. Their performance is evaluated by extensive synthetic experiments as well as experiments on real images.

1. Introduction

Many objects have smooth surfaces of a fairly uniform color, thereby exhibiting shading patterns that reveal information about their shape, an important clue to the nature of the object. Research on extracting shape from shading has been ongoing since [9], yet information from shading is rarely used in higher level tasks such as object recognition with a few exceptions, [15, 19].

The aim of this work is to extract information from shading in real images. Rather than trying to recover the surface normal at every position under a number of restrictive assumptions, we try to get a rough estimate of the shape and without any a priori knowledge and very few assumptions about the light source and surface reflectance proper-

ties. This is done by creating detectors for a set of shape primitives, much like [7].

We follow the tradition of using principal component analysis (PCA) to learn appearance models under illumination changes [8] and view changes [12]. Although instead of performing PCA on a set of images, we apply it directly to the image formation equation. Variations in the light source, surface reflectance function (BRDF) and shape define the variations in the image. By performing PCA on an analytic formulation we can rapidly train appearance models for a wide variety of conditions.

2. Shading in Frequency Space

Shading is the gradual change in the image intensity due to variations in the surface normal relative to the light source. The image of a surface is a function of the surface shape, the light source and the surface reflectance properties. Our aim is to create models of shading that are applicable to any kind of image taken with a normal camera (video or still). It is therefore necessary to accurately model all the parts of this process and requires our model to incorporate arbitrary light sources, unknown reflectance functions and variations in the shape.

Fortunately, recent research [1, 17] has enabled the modeling of arbitrary light sources through the use of spherical harmonics, the Fourier basis on the sphere. Surfaces act as low-pass filters on the incident light making the reflected light band-limited in frequency. This allows any illumination condition to be represented with a finite dimensional model in frequency space. In such a model the lighting conditions are represented by their spherical harmonics coefficients.

Also the BRDF can be represented in frequency space as it often is a smooth function. We use the basis for isotropic BRDF's by Koenderink and van Doorn [10] based on the Zernike polynomials, but a basis using spherical harmonics, such as in [18], works equally well. This basis can represent any isotropic BRDF while incorporating the general

Helmholtz’s reciprocity.

The image intensity is proportional to the scene radiance which is the sum of reflected light towards the camera, which in turn is determined by the incident light and the BRDF. By inserting the frequency space representation of the light source and BRDF and solving this integral we get a frequency space basis of the image intensity. The image intensity, I_i , at point i can then be represented as a linear combination of basis functions.

$$I_i = \sum_{k=1}^N c_k E_k(\alpha_i, \beta_i), \quad (1)$$

where (α_i, β_i) is the surface normal, in spherical coordinates, at position (x_i, y_i) in the image. The coefficients $c_k = L_l^m b_{op}^q$, *i.e.* are the products of the spherical harmonics coefficients L_l^m of the light source and the coefficients of the BRDF, b_{op}^q in the Koenderink and van Doorn basis. The indices l, m, o, p and q are given from k by the specific ordering of the basis functions. The basis functions E_k are the results of solving analytically the image formation integral. They are products of the Wigner D-functions (for real spherical harmonics) and the Zernike polynomials. For more details and their explicit form see [14].

The camera projection is approximated as orthographic as the image intensity otherwise would be dependent on the view angle as well. We also assume that all points on the surface are illuminated by the same light source, which is correct if the light source is distant and there are no cast shadows. This assumption usually holds when we are analyzing small image patches as opposed to the whole image at once.

Equation (1) represents the shading in very general conditions, *i.e.* arbitrary unknown light source and unknown BRDF. To be exact an infinite sum is required, but because the surface low-pass filters the reflected light the series can be truncated and still be a good approximation. Even so, while nine basis functions are enough for Lambertian surfaces [1, 17], specular surfaces may require thousands, making it impractical to use directly. One way to reduce this representation is to perform PCA on it.

3. Model-Based PCA of Shading

PCA finds the orthogonal set of directions with maximal variance in a multidimensional dataset. The first principal component is the vector that maximizes the variance of the dataset projected onto the vector. The subsequent principal components maximizes the same variance while being orthogonal to the previous components. By using only the low order principal components we can reduce the dimensionality while maintaining sufficient representational power.

Variations in the light source, surface reflectance and shape define the variations in the image. By performing

PCA on all images of an object we can acquire a low-dimensional basis for representing the appearance of the object. However, the same PCA basis can be computed with model-based PCA by analytically deriving the principal components from the image formation model. Thus, there is no need to capture or render any images. The principal components are computed directly from the variations in the light source and surface reflectance function. One of the advantages is that the light source can be rotated continuously to take into account an infinite number of lighting conditions.

Previous work has analytically derived the principal components of Lambertian surfaces under varying illumination [16, 13] as well as surfaces with variations in the BRDF [14]. In this paper we extend this to incorporate variations in the surface shape as well.

3.1. New Extension to Shape Variations

We begin by defining the image set upon which the PCA will be performed. The image depends on the illumination, the BRDF and the shape. Hence the set of images is defined by the variations of these parts.

Let \mathbf{L} be the spherical harmonics coefficients of the light source and $p_L(\mathbf{L})$ the probability function describing how the illumination conditions vary. Likewise, let \mathbf{b} be the coefficients of the BRDF and $p_b(\mathbf{b})$ their probability function. Finally, define a set of shapes, \mathcal{S} where each shape $s \in \mathcal{S}$ is represented by a set of surface normals $\{(\alpha_{is}, \beta_{is}); i = 1, \dots, n_p\}$, where n_p is the number of pixels in the image. Also define the probability function $p_s(s)$, $s \in \mathcal{S}$. All these probability distributions define the image set. Now we derive the PCA on this set.

We begin by representing the image as a column vector \mathbf{I} . Then equation (1) can be written as a matrix product

$$\mathbf{I} = \mathbf{E}\mathbf{c} \quad (2)$$

where the elements of \mathbf{E} are the irradiance basis functions, $e_{ik} = E_k(\alpha_i, \beta_i)$ and \mathbf{c} contains c_k , the products of the illumination and BRDF coefficients. Note that \mathbf{E} depends on the shape only, while \mathbf{c} depends on the illumination and BRDF, but not the shape.

When computing the image covariance matrix it is important how the image set is centered. In standard PCA this is done by subtracting the mean image. We do not do this. Instead, we center by subtracting the images’ individual mean from each image. This produces better eigenimages as they will contain no DC component and there is no need to subtract a mean image before fitting the PCA basis to an image.

The centered image is the image pixels subtracted by its’

mean. We can write this as another matrix product

$$I_{cent} = \mathbf{I} - \left(\frac{1}{n_p} \sum_{i=1}^{n_p} I_i \right) \mathbf{1} = \mathbf{F} \mathbf{c} \quad (3)$$

where the elements of \mathbf{F} are $f_{ik} = E_k(\alpha_i, \beta_i) - \frac{1}{n_p} \sum_{i=1}^{n_p} E_k(\alpha_i, \beta_i)$.

Now, the covariance matrix with our centering is

$$\Sigma_I = \langle I_{cent} I_{cent}^T \rangle_{L,b,s} \quad (4)$$

where $\langle \cdot \rangle_x = \int_x p_x(x) dx$ is the expectation operator over variable x . Inserting (3) into (4) results in

$$\Sigma_I = \langle \mathbf{F} \mathbf{c} \mathbf{c}^T \mathbf{F}^T \rangle_{L,b,s}. \quad (5)$$

As \mathbf{F} is independent of L and b we can move their expectation operator inside the matrix product. We obtain

$$\Sigma_I = \langle \mathbf{F} \mathbf{V}_c \mathbf{F}^T \rangle_s \quad (6)$$

where $\mathbf{V}_c = \langle \mathbf{c} \mathbf{c}^T \rangle_{L,b}$ is the second moment matrix of \mathbf{c} . This matrix contains the covariances of the illumination and BRDF coefficients. The covariance for a number of lighting conditions and BRDF's have been computed in [13] and [14].

Finally, \mathbf{F} varies with the shape. Let \mathbf{F}_s be the matrix \mathbf{F} computed with the surface normals of shape s . Then

$$\Sigma_I = \sum_{s \in \mathcal{S}} \mathbf{F}_s \mathbf{V}_c \mathbf{F}_s^T p_s(s) \quad (7)$$

The PCA is, as always, performed by computing the singular value decomposition (SVD) of the covariance matrix Σ_I . The resulting eigenvectors \mathbf{U}_k are our basis images.

3.2. An Example

To demonstrate this new framework we have computed the PCA basis for a set of faces. The shape variations are defined by four 3D face models provided by the Max-Planck Institute for Biological Cybernetics in Tuebingen, Germany [3]. The variations in light source are described by a set of captured environment maps (Debevec [5]), each rotated through all possible 3D rotations. Computing the covariance matrix for all rotations of an illumination map is particularly easy, as the cross-covariances between coefficients are all zero and the variances are the sum of squares of the coefficients of the same order, [13].

In the first basis the surface reflectance properties are set to be human skin from the CURET database [4], see Figure 1a. Note the basis images 1-3 and possible 5, which represent the illumination changes, while the other images encode shape variations.

To further demonstrate the framework we have computed a second basis where the Torrance-Sparrow model is used for the surface reflectance. The roughness parameter of the BRDF ranges between 1 and 20, Figure 1b. In contrast to the previous basis, nearly all of the displayed basis images are illumination modes.

4. Shape Detection

The calculated PCA basis can be used for shape recognition by measuring how well the low-dimensional model can reconstruct a given image signal \mathbf{I} . The best estimate of the signal using the first n vectors of the PCA basis is given by

$$\hat{\mathbf{I}} = \sum_{k=1}^n d_k \mathbf{U}_k + \mathbf{1}\bar{I}, \quad \text{with } d_k = \mathbf{U}_k^T \mathbf{I}, \quad (8)$$

where \bar{I} is the image mean.

The residual reconstruction error between the signal and its estimate is given by

$$\epsilon^2(\mathbf{I}) = \|\mathbf{I} - \hat{\mathbf{I}}\|^2 = \sum_{i=n+1}^N d_i^2 = \|\mathbf{I} - \mathbf{1}\bar{I}\|^2 - \sum_{i=1}^n d_i^2. \quad (9)$$

A low value of $\epsilon^2(\mathbf{I})$ is a sign that the data fits the low-dimensional model well. However, it may only indicate that the variance of the considered image patch is low. For instance, a completely flat image will be perfectly recreated by the model and produce a residual error of zero. Therefore, image patches with small values of $\|\mathbf{I} - \mathbf{1}\bar{I}\|^2$ are ignored.

The final score used to indicate the presence of a shape is a normalized version of the residual error

$$\tilde{\epsilon}^2(\mathbf{I}) = \frac{\epsilon^2(\mathbf{I})}{\|\mathbf{I} - \mathbf{1}\bar{I}\|^2} \quad (10)$$

Low values of $\tilde{\epsilon}^2(I)$ signify the likely presence of the surface of interest.

4.1. General Approach

The above principle can be used to detect instances of a shape, but as an object moves and/or changes its pose the appearance can drastically change. It is not practical to let all these changes be handled by a single model. We suggest the following approach which is similar to the view-based approach in [12].

Scale changes are handled by running a detector on a multi-scale pyramid of the test image. Moreover, the detectors are trained to handle the variations in size between the discrete scales in the pyramid. To cope with varying poses a set of models is created to represent the shape in different poses. Each model is then trained to handle the small pose changes between models of adjacent poses.

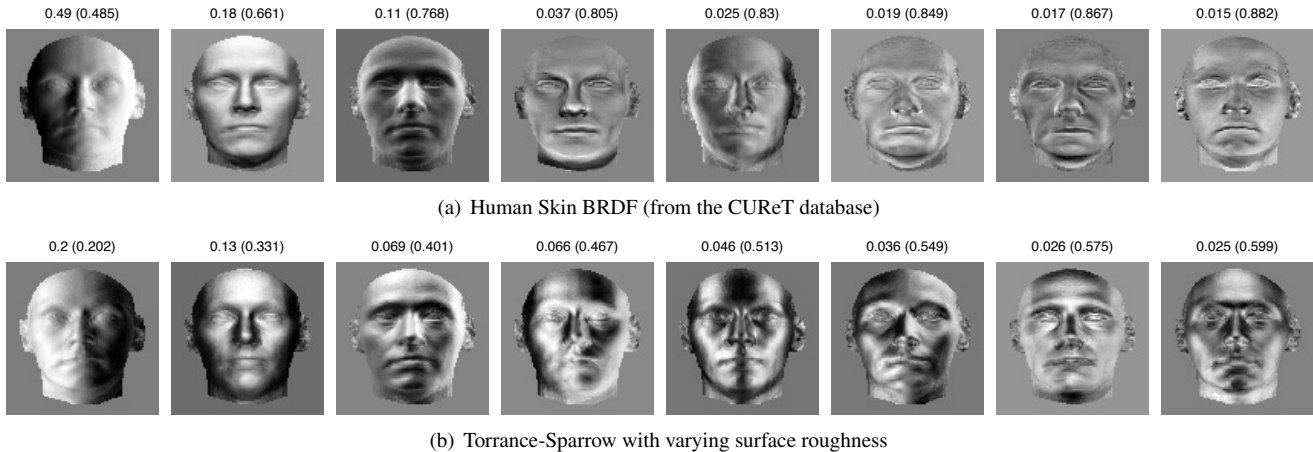


Figure 1. Eigenfaces computed with Model-based PCA from four 3D face models. The light source conditions were modeled with set of captured environment maps rotated through all 3D rotations. The top basis a) was computed with the human skin BRDF from the CURET database. Notice the first three and the fifth eigenimages capture the illumination changes, while the remaining eigenimages encode shape changes. b) The second basis was computed with a Torrance-Sparrow BRDF model of varying roughness. The shinier surfaces create more illumination modes as all displayed eigenimages now encode light source changes.

It is also possible to train the models for shape deformations such as the differences between individuals in a face detector, as demonstrated earlier.

To demonstrate the approach we have created detectors for two basic shape primitives, spheres and cylinders. Due to these shapes’ rotational invariances, the different poses are slightly easier to handle than more complex shapes, but the approach is in principle applicable to any shape or group of shapes. Note also that an image of sphere contains all the visible surface normals and essentially any shape is just a redistribution of those surface normals. Hence, the evaluation of a sphere detector also gives an indication of the performance of shape detection in general.

5. A Sphere Detector

The case of the sphere is relatively simple as it is completely rotationally invariant. Besides the unknown lighting and BRDF, we only need to deal with scale changes and therefore are only required to create one appearance model.

Model-based PCA takes a set of shapes, each in the form of an image containing the surface normals at each pixel. A size of the filter is set. We use 13x13 in this case. The size of the sphere ranges from the size of the filter up to when it interlinks with the filter running on the next scale in the image pyramid. We sample this range and compute a set of surface normals for each size. We also train for sub-pixel translations as this improves detection. This is done by sampling the center position in both x and y coordinates between -0.5 and 0.5 . All these variations need to be done in combination. We sample the size at 10 points and the translation at 11 points per axis resulting in a total of 1210 different shapes (sets of surface normals).




Figure 2. The sphere shading model is trained with shape variations due to scale changes and sub-pixel translations. a) The radius of the sphere is varied to interlink the scales in the image pyramid. b) The center of the sphere is also moved within the space of one pixel. To train the model all combinations of scale and sub-pixels translation are used.

The model also has a mask defining which pixels in the filter actually “see” the object. The smallest sphere sets the mask of the filter and the mask is further reduced by the sub pixels translation as illustrated by Figure 2.

5.1. Lighting and BRDF variations

The lighting conditions are modeled, as in the face bases, with the set of environment maps by Debevec, rotated through all 3D rotations.

The surface reflectance properties were selected to be a set of BRDF’s of Torrance-Sparrow with Beckmann’s facet distribution, mixed with various levels of Lambertian reflectance. Both the specular coefficient k_{spec} and the roughness parameter m where varied as illustrated in Figure 3. The diffuse coefficient k_d was scaled with $1 - k_{spec}$. The BRDF’s were projected onto the BRDF basis and low-pass filtered, as little as possible, to avoid Gibbs ringing. Figure 4 shows the resulting sphere basis.

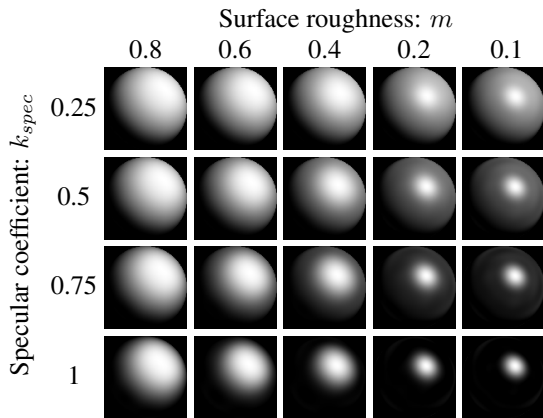


Figure 3. BRDF’s used for training the detectors, here illustrated by rendered spheres. The BRDF’s are a mix of Torrance-Sparrow with Beckmann facet distribution and Lambertian reflectance. Both the surface roughness and specular coefficient is varied as illustrated above. A pure Lambertian BRDF, although not rendered here, is also included in the set.




Figure 4. The ten first basis images for the computed appearance model of a sphere. The numbers show each image’s eigenvalue and within parenthesis the cumulative sum of eigenvalues.

6. A Cylinder Detector

Cylinders are partially rotationally invariant. There are two degrees of freedom that we need to take care of to do 3D cylinder detection. One we deal with and the other we argue we are invariant to.

Rotation in the image, that is rotation around the normal vector of the image plane, is handled by creating a set of models for the different poses.

Rotation into the image plane is not handled. Due to the bas-relief ambiguity [2] there is an inherent ambiguity which cannot be resolved. Although this was proven for Lambertian surfaces only, it also holds to some degree for non-Lambertian reflectance. Think of it this way. Let the light field rotate with the cylinder, essentially the camera is just changing viewpoint. For Lambertian surfaces the shading will be identical (if we disregard the perspective effects). The specular reflection will move as the viewpoint changes, but as we have no knowledge about the light source it is very unlikely that we will be able to differentiate between different rotations into the image plane. For very steep angles the detector is more likely to fail. If one were interested in the exact pose of the cylinder one could analyze the perspective




Figure 5. A set of models is created to represent different rotations of a cylinder. As in the sphere case, each model is trained for scale variations and sub-pixel translations. a) In addition the models are trained for the small rotations between models. The gray squares show the resulting mask due to the variations in pose. b) Computed appearance model for one of the eight directions of a cylinder.

effects at the edges of the cylinder, but that is beyond the scope of this paper.

As in the sphere case, each model is trained to handle small scale variations and sub-pixel translations. Additionally, we train the model to handle the 2D rotations inter-linking the set of models, see Figure 5a. Again, all these variations are applied in all combinations.

Lighting conditions and BRDF’s were modeled as in the sphere case.

To compute the model in Figure 5b we sample the scale variations at 5 points, translation at 11 points per axis and the rotation at 7 points. This results in a total of 4235 shapes.

7. Experiments

7.1. Synthetic experiments

The frequency space framework allows for rapid rendering of images. Thousands of images can be rendered in an instant through a simple matrix multiplication. This allows for extensive testing.

We test the detectors by computing the normalized residual variance in Equation (10) for a huge number of rendered images. This shows how well the PCA basis is able to represent all the possible variations in the image of a shape.

Images are rendered for combinations of illumination, illumination rotation, BRDF, sub-pixel translation and scale changes. Nine illumination maps rotated with 3726 different rotations, 21 BRDF’s, 5x5 sub-pixel and 5 scale changes results in over 88 million images when evaluating the sphere detector. The cylinder detector has another 5 combinations of pose variations resulting in a total of 440 million images!

The normalized residual variance is computed using different numbers of principal components and stored in histograms. Figure 6 shows the mean errors for the sphere detector for a selection of different BRDF’s and number of principal components. As expected Lambertian surfaces has the lowest errors. Smooth (low roughness) and specular materials require more principal components.

But the errors don’t say much on their own. More principal components will reduce the error for spheres, but will


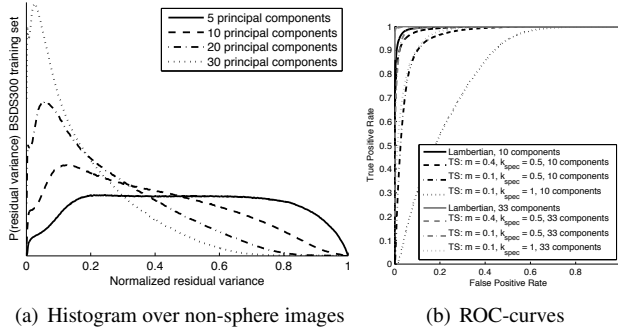


Figure 6. Residual variances for different materials.



(a) Histogram over non-sphere images

(b) ROC-curves

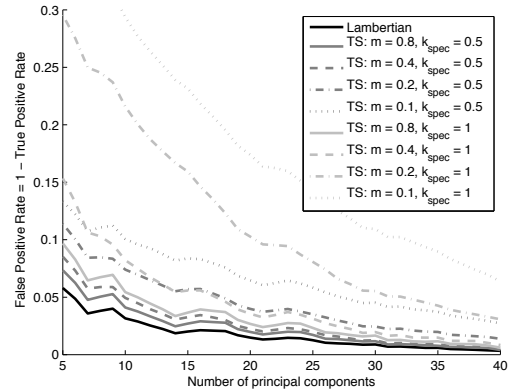
Figure 7. a) Histogram of the residual variance from the sphere detector on the BSDS300 training set. b) ROC-curves for the sphere detector.

likely also reduce the errors for non-spheres. To investigate this we run the detectors on the Berkeley Segmentation Dataset (BSDS300), [11], which to our knowledge contains no spheres, and record their error histograms. Figure 7a shows some of the recorded histograms. Together with the sphere histograms we can plot ROC-curves for the detectors, see Figure 7b.


To further summarize these results we plot the errors at the equal error rate, *i.e.* when the false positive rate equals $1 - \text{the true positive rate}$. Figure 8 shows these graphs for the sphere and cylinder detector. Again, as expected shinier BRDF's require more principal components. Interestingly, more components lead to better discrimination. The basis increasingly fits spheres better than non spheres with an increasing number of components. However this might not be true on real images with other sources of error than the reconstruction error. The graphs also show that it is important to be careful when choosing the number of components. This is most evident in the cylinder case where having six or seven actually reduces the performance compared to five components. This phenomenon is most likely related to the number of modes needed to create rotational invariance in a basis.

7.2. Experiments on real images

To further test the approach we constructed a small dataset of 31 images containing three spheres with differ-



(a) sphere detector



(b) cylinder detector

Figure 8. Predicted false positive rates when threshold is set so that the false positive rate equal $1 - \text{true positive rate}$. These graphs summarize the synthetic experiments quite well. They give an indication on how shiny materials the models can handle and serve as a guide when choosing the number of principal components.




Figure 9. Examples from the sphere image dataset

ent surface reflectance properties. The spheres are made of paper. One is kept as it is, one is painted with a silver color and the third with a gold color than is slightly shinier than the silver paint. The images were taken in a number of different illumination conditions such as sunny, overcast and in various indoor office type environments, see Figure 9 for examples.

The positions of the sphere are manually marked and it is computed at which scale and position the sphere should be detected. At this scale and position the residual variance is recorded. Figure 10 shows the mean errors. Not surprisingly the shinier sphere the higher error for the same number of principal components. Plotted are also some results




Figure 10. Residual variances of the sphere image dataset. The thin lines show some of the results from the synthetic experiments for comparison.

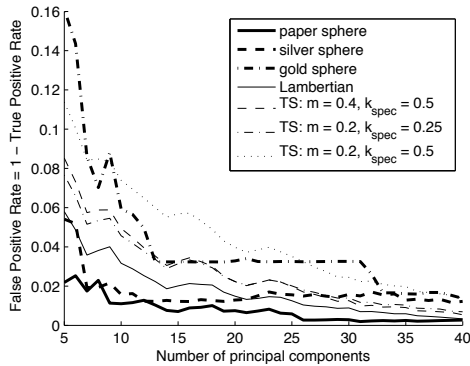


Figure 11. Equal error rates for the sphere dataset. The thin lines show some of the results from the synthetic experiments for comparison.


from the synthetic experiments for comparison. The synthetic errors decrease faster for higher number of principal components. This is most likely because the real images have other types of errors which remain when the reconstruction error is reduced.

To estimate performance of the detector we again compare the errors with the BSDS300 dataset. In Figure 11 the equal error rates are plotted for the sphere dataset as well as some relevant results from the synthetic experiments. It seems that even in real images more principal components increase the performance of the detector, although the silver sphere's curve flattens out at about 12 components.


7.3. Some examples

To get an idea of the performance on more complex images we have created two collages of spheres, Figure 12a and cylinders, Figure 13a. The images were gathered from the web and other source such as the Caltech 256 database, [6].


The detectors were run on each scale in the image pyramid. The number of principal components used were 33 and 5 for the sphere and cylinder respectively. To visualize the errors we merge the results from all scales in a winner-take-



(a)



(b)



(c)

Figure 12. Results from the sphere detector. Please view these in color! The results are discussed in the text.

all fashion. Figures 12b and 13b shows the errors plotted as $e^{-e^2/a}$ to get a suitable contrast. As can be seen most of the spheres and many of cylinders have a high score at the center, but also near the center (but at a smaller scale). Texture and occlusion are sources of errors. Shadows and shape deviations also cause problems. Some of these could be improved *e.g.* by modeling smaller parts of shapes.


Figures 12c, 13c and 13d show the scale and direction for the best fits in the images. These can for the most part be considered correct.

8. Conclusions


We have presented a systematic approach for the creation of shape detectors based on shading. This approach is enabled by extending recent research in frequency space representations of shading and model-based PCA. A novel




(a)



(b)



(c)



(d)

Figure 13. Results from the cylinder detector. Please view these in color! The results are discussed in the text.

derivation of model-based PCA allows the inclusion of shape changes in the formulation. This can be used to handle variations in the object's size and pose as well as other shape variations such as between individual faces as demonstrated in the paper.

We propose to use this extended model-based framework to detect shapes from shading patterns in images. By analyzing the goodness of fit of the appearance models created, it is possible to detect the shape invariant to illumination conditions and the surface reflectance properties. With the new framework the models can also be trained to cope with

small variations in scale and pose. Large scale variations are handled by running the detector on a multi-scale image pyramid. Large pose/view changes are dealt with in the classic view-based way of creating a set of models, one for each pose/view.

The approach is demonstrated by creating detectors for two basic shape primitives, spheres and cylinders. These shapes' rotational invariances make it easier to handle pose changes. More complex shapes would require more work.

The detectors were extensively tested on synthetic images as well as on a dataset of real images, with nice correspondence between the synthetic and real experiments. Matte surfaces are easier to detect than shiny, but shiny surface can be detected as well but requires more principal components.

References

- [1] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *IEEE PAMI*, 25(2):218–233, February 2003.
- [2] P. Belhumeur, D. Kriegman, and A. Yuille. The bas-relief ambiguity. *IJCV*, 35(1):33–44, November 1999.
- [3] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *SIGGraph*, pages 187–194, 1999.
- [4] K. Dana, B. van Ginneken, S. Nayar, and J. Koenderink. Reflectance and texture of real-world surfaces. *ACM ToG*, 18(1):1–34, January 1999.
- [5] P. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *ACM SIGGraph, Computer Graphics*, pages 189–198, 1998.
- [6] G. Griffin, A. Holub, and P. Perona. Caltech-256 object category dataset. Technical report, Caltech, 2007.
- [7] J. Haddon and D. Forsyth. Shape representations from shading primitives. In *ECCV*, pages 415–431, 1998.
- [8] P. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In *IEEE Computer Vision and Pattern Recognition*, pages 995–999, 1994.
- [9] B. Horn. Shape from shading: A method for obtaining the shape of a smooth opaque object from one view. Technical Report AI-TR-232, MIT, 1970.
- [10] J. Koenderink and A. van Doorn. Phenomenological description of bidirectional surface reflection. *JOSA-A*, 15(11):2903–2912, November 1998.
- [11] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its appl. to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, volume 2, pages 416–423, July 2007.
- [12] H. Murase and S. Nayar. Visual learning and recognition of 3-d objects from appearance. *IJCV*, 14(1):5–24, January 1995.
- [13] P. Nillius and J. Eklundh. Low-dimensional representations of shaded surfaces under varying illumination. In *IEEE Computer Vision and Pattern Recognition*, pages II:185–192, 2003.
- [14] P. Nillius and J. Eklundh. Phenomenological eigenfunctions for image irradiance. In *ICCV*, pages 568–575, 2003.
- [15] M. Osadchy, D. Jacobs, and R. Ramamoorthi. Using specularities for recognition. In *ICCV*, pages 1512–1519, 2003.
- [16] R. Ramamoorthi. Analytic pca construction for theoretical analysis of lighting variability in images of a lambertian object. *IEEE PAMI*, 24(10):1322–1333, October 2002.
- [17] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: determining the illumination from images of a convex lambertian object. *JOSA-A*, 18(10):2448–2458, October 2001.
- [18] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for reflection. *ACM ToG*, 23(4):1004–1042, October 2004.
- [19] P. Worthington and E.R.Hancock. Object recognition using shape-from-shading. *PAMI*, 23(5):535–542, 2001.