Division of Informatics, University of Edinburgh

Institute of Perception, Action and Behaviour

Comparison of HK and SC curvature description methods

by

Helmut Cantzler, Robert Fisher

Informatics Research Report EDI-INF-RR-0081

Division of Informatics **May 2001**
http://www.informatics.ed.ac.uk/

Comparison of HK and SC curvature description methods

Helmut Cantzler, Robert Fisher

Informatics Research Report EDI-INF-RR-0081

DIVISION *of* INFORMATICS
Institute of Perception, Action and Behaviour

May 2001

Proc. 3rd Int. Conf. on 3-D Digital Imaging and Modeling (3DIM01), Montreal, Canada, pp 285-291, May 2001.

Abstract :

Keywords :

Comparison of HK and SC curvature description methods

H. Cantzler and R. B. Fisher

Machine Vision Unit,
Institute for Perception, Action and Behaviour,
Division of Informatics, University of Edinburgh,
Edinburgh, EH1 2QL, UK
helmutc@dai.ed.ac.uk

Abstract

This paper compares two different local surface shape description methods. The general goal of surface shape description methods is to classify different surface shapes from range data. One well-known method to classify patches of various shapes is the HK classification space [2, 1, 10]. Another way to classify patches is the SC method introduced by Koenderink [9]. This paper presents several experiments designed to show the (1) qualitatively different classification, (2) the impact of thresholds and (3) the impact of different noise levels. We conclude that Koenderink's approach has some advantages at low thresholds, complex scenes and at dealing with noise.

1 Description of the algorithms

Gaussian (K) and Mean (H) curvatures are the most widely used indicators for surface shape classification in range image analysis. The HK segmentation [2, 1, 10] was introduced by Besl in 1986. He used Gaussian and Mean curvatures, which are calculated from the two principal curvatures κ_1 and κ_2. The Gaussian curvature equals the product of the principal curvatures.

$$K = \kappa_1 * \kappa_2 \tag{1}$$

The Mean curvature equals the arithmetic average of the principal curvatures.

$$H = \frac{\kappa_1 + \kappa_2}{2} \tag{2}$$

Image points can be labelled as belonging to a viewpoint independent surface shape class type based on the combination of the signs from the Gaussian and Mean curvatures as shown in Table 1. We found that it is not necessary to differentiate between the different kinds of saddles (ie. K $<$ 0). Therefore, we classify all saddle points as hyperboloid (Hy) points.

	K $<$ 0	K = 0	K $>$ 0
H $<$ 0	Saddle Valley (Sv Hy)	Concave Cylinder (-Cy)	Concave Ellipsoid (-El)
H = 0	Minimal (M Hy)	Plane (Pl)	Impossible
H $>$ 0	Saddle Ridge (Sr Hy)	Convex Cylinder (+Cy)	Convex Ellipsoid (+El)

Table 1: Classification for the HK segmentation based on the sign of the Mean (H) and Gaussian (K) curvatures

Koenderink defined an alternative curvature representation [9]. His approach (SC classification) decouples the shape and the magnitude of the curvedness. The surface in terms of relative curvature remains invariant under changes in scale. He defined a shape index S, which is a number in the range [-1,1]. The index covers all shapes except for the planar shape which has an indeterminate shape index ($\kappa_1 = \kappa_2 = 0$). The shape index provides a continuous gradation between shapes, such as concave shapes (-1 $<$ S $<$ -1/2), hyperboloid shapes (-1/2 $<$ S $<$ 1/2) and convex shapes (1/2 $<$ S $<$ 1). The image points can be classified as shown in Table 2. We use the positive principal curvatures ($\kappa_{1,2} > 0$) for convex objects.

$$S = \frac{2}{\pi} * \arctan\left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2}\right) \quad \kappa_1 \geq \kappa_2 \tag{3}$$

Beside the shape index, Koenderink introduced the positive value C for describing the magnitude of the curvedness at a point. It is a measure of how highly or gently curved a point is. At a point that has no curvedness the value becomes zero. Therefore, this variable can be used to recognise a planar surface.

$$C = \sqrt{\frac{\kappa_1^2 + \kappa_2^2}{2}} \quad (4)$$

The SC classification was used in several publications as a replacement [5, 4] or an enrichment [8] of the HK segmentation. Although both algorithms are used, the literature contains as far as we know no comparison of these algorithms. Guid compared Koenderink's approach in an application-based view with several other surface interrogation methods in terms of determining surface fairness, run-time, sensitivity to changes and invariance to geometric transformations [7]. His conclusion is that this approach is appropriate for detecting waveness and searching extreme curvatures. Inflection points and convex/concave regions may be properly detected. In his comparison the SC classification is one of the slowest algorithms. The algorithm has a fair sensitivity to changes in surface and is invariant to geometric transformations.

2 Thresholds

Both the HK and SC methods have the problem that it is impossible to have an exact zero value, because of image noise and small shape variations. Therefore, zero-thresholds are used to decide if a value is zero or not. Everything below a certain threshold is recognised as zero. Figure 1 shows the classification regions for the HK segmentation (top) and the SC segmentation (bottom) for the same threshold value. The graphs are drawn to the same scale. Regions are labelled using the shapes from Tables 1 & 2. Data-points are located in the κ_1/κ_2-graph like in Figure 6. The classification schema are put over the graph to classify the points.

The HK segmentation uses two thresholds. The Gaussian τ_K and the Mean τ_H threshold are used to classify plane surfaces. With the following formula the Gaussian threshold can be calculated from the Mean threshold [3].

$$\tau_K = \tau_H * (\tau_H + 2 * \max|H(x)|) \quad x \in Image \quad (5)$$

Koenderink's approach uses the threshold τ_C to classify plane surfaces patches (if $C \leq \tau_C$). We choose $\tau_C = \tau_H$ for comparison of the algorithms.

Shape	Index range
Concave Ellipsoid (-El)	S ∈ [-1,-5/8)
Concave Cylinder (-Cy)	S ∈ [-5/8,-3/8)
Hyperboloid (Hy)	S ∈ [-3/8,3/8)
Convex Cylinder (+Cy)	S ∈ [3/8,5/8)
Convex Ellipsoid (+El)	S ∈ [5/8,1]

Table 2: Classification for Koenderink's approach based on the shape index (S)

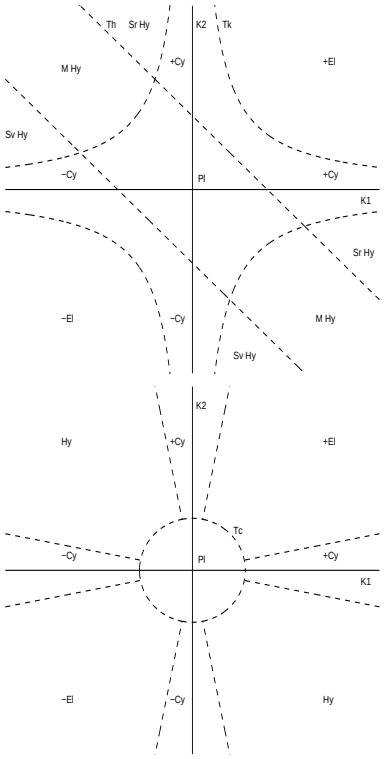

Figure 1: Classification for the HK (top) and SC (bottom) methods. Dashed lines are determined by the threshold values and separate the classification regions

Some implementations use a low and a high threshold. "Each curvature value is classified as Negative, Zero, Positive or Unknown based on the values of 'inner' and 'outer' thresholds" [6]. Everything below the low threshold is recognised as zero. A value between the low and high threshold is recognised as unknown (classification later according to local context) and everything above the high threshold is recognised as a normal value. To make the comparison easier to perform here the low and high thresholds are the same.

3 Comparison

We separate the comparison into three sections. In each of the sections below we answer a question:

Section 3.1: Do the algorithms make any qualitatively different shape classifications?

Section 3.2: Do the two algorithms have any qualitatively or quantitatively different behaviour as the classification threshold varies?

Section 3.3: How do the two algorithms compare as noise levels increase?

3.1 Mathematical comparison

The classification of shapes is implemented quite differently in the two algorithms. The HK segmentation relies on the the right values of the two thresholds. Because of the thresholds, the plane area in the HK scheme is not symmetric (see Figure 1). Furthermore, the cylinder area contains a part of the ellipsoid area ($\kappa_1 = \kappa_2$). Beside this, the cylinder area gets narrower for high curvatures. On the other hand, with the SC classification the plane area is symmetric and the cylinder area uses a constant range of the shape index.

The following shows that the region for the cylinder is narrower for highly curved objects in the HK algorithm. For a cylinder just at the classification threshold (assume $\kappa_1 > \kappa_2 > 0$) we have:

$$\kappa_2 = \frac{\tau_K}{\kappa_1} \tag{6}$$

If we plot the cross-section ($z = 0$) of ellipsoids ($\frac{1}{2}\kappa_1^2 x^2 + \frac{1}{2}\kappa_2^2 y^2 + \frac{1}{2}\kappa_1^2 z^2 = 1$) for 3 decreasing values of κ_1 we see a family of shapes as shown in Figure 2. We use the threshold $\tau_K = 0.008$ to calculate the shapes. The allowable "cylindrical shapes" must lie inside these limit shapes for a given extent in the x direction.

On the other hand, the region for the cylinder is bigger for Koenderink's approach at the shape border ($\frac{5}{8}$):

$$\frac{5}{8} = \frac{2}{\pi} * \arctan\left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2}\right) \tag{7}$$

$$\kappa_2 = \kappa_1 * \frac{1 - tan(\frac{5\pi}{16})}{1 + tan(\frac{5\pi}{16})} \approx \kappa_1 * 0.1989 \tag{8}$$

If we plot a similar family of cross-sections as above for the same 3 decreasing values of κ_1 we see this family of shapes as shown in Figure 2. As above, the allowable "cylindrical shapes" lie inside these limit shapes for a given extent in the x direction.

What these two figures show us is that as κ_1 becomes smaller (ie the radius in the x direction becomes bigger) the HK algorithm becomes more strict about what is cylindrical by requiring more elongated shapes to have thinner radii, whereas the SC algorithm allows the maximal radius to scale up as the shape elongates.

3.2 Empirical comparison

Both classification methods were implemented in our range image segmentation program. The curvatures are estimated by using the 1st and 2nd fundamental forms of the surface pixels [1]. The segmentation program was compared with other approaches by Hoover in [6]. The classification methods label each point of a range image after smoothing it once as one of six shapes. The shapes are plane, concave/convex ellipsoid, concave/convex cylinder and hyperboloid. The results (the classified images) are compared with the ground truth of the images to count the number of mislabelled points. The ground truth is a hand segmentation, which outlines the boundary of each surface patch.

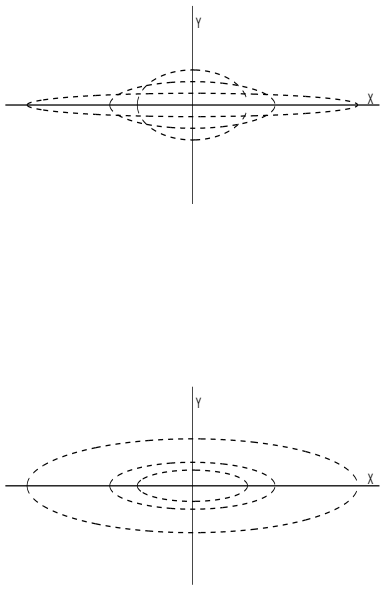

Figure 2: The shape family at the border between the positive cylinder and the positive ellipsoid for HK (top) and SC (bottom). The shapes of each image are calculated with the three decreasing κ_1 values 0.06, 0.04 and 0.02. Allowable cylinders lie inside the limit shape for a given extent in the x direction. The cylinder is rotationally symmetric about the X axis.

For testing two kinds of test images were used. First, some real test images which contain only one type of shape were used. Then, two more complex test images, which contain several shapes, were used. Each test image was segmented 25 times with different thresholds for each method.

Figures 3 till 5 show graphs of the scores of the two algorithms applied to real, but simple, shapes. The graphs describe the percentage of mislabelled points (vertical axis) for different thresholds. Some points are worth noticing. First, Koenderink's approach is slightly shifted to the right side on all test graphs, because the same threshold creates

a larger central/planar classification region for the HK segmentation in the κ_1/κ_2-coordinate-system. This difference is not relevant as one could adjust thresholds accordingly. Second, for the images that contain only planar or hyperbolic surfaces (Figures 3 & 4) and also for elliptical surfaces the best result occurs at low or high thresholds, because the threshold can be used to make the plane region cover all data-points, or shrink the plane & cylinder region to zero. Figure 4 shows classification results on a real hyperbolic surface. As the threshold increases, points are mislabelled as planar. As the threshold decreases, correct classification increases, but again levels off as noise limits classification. In this case the HK algorithm does slightly better because its cylindrical classification region disappears, allowing more noisy points to be classified correctly. Thus, the difference between the two algorithms appear for points having locally cylindrical shape.

The results for the convex cylinder (see Figure 5) are interesting. At high threshold levels, all points are misclassified as being planar. What is more interesting is at lower threshold levels. With the HK classification, as the threshold decreases, more points are labelled as cylindrical until a critical level when points start being misclassified as elliptical, because the cylindrical region has been shrunk too much. This trade off does not occur with the SC classification. Of note in both cases is the high level of misclassification arising from surface shape noise. Figure 6 shows a scatter plot of the estimated curvatures on the cylindrical surface, which are clearly very noisy even though the data itself is reasonably good. This sensitivity arises from the need to estimate local 2^{nd} derivatives. One can easily see why many points are misclassified. A few points belong to a negative cylinder. These points are located at the border of the positive cylinder.

Figure 7 shows the results from a more complex object (see Figure 8) containing multiple shape classes (planar, elliptical/spherical and cylindrical surfaces). The classification results here are from a combination of surface shapes. What seems the case here is the SC algorithm offers a broader range of threshold values where the classification performance is best. If optimal threshold values can be found, the HK algorithm can get slightly better classification results. However, as seen in Figure 5 & 7, this optimal value may be hard to find when no ground truth is known.

A second complex scene were used to compare the approaches. The scene is a factory scene and contains planes and cylinders with different radii (see Figure 10). Figure 9 shows that the SC algorithm performs slightly better. The small dips in the graph mark the thresholds that discriminate between the different cylinders.

3.3 Behaviour for different noise levels

Image noise is an important issue in vision computing. To measure the impact of noise on both algorithms we performed another series of experiments with noisy data. Firstly, we created several synthetic images with all primitive surfaces (plane, cylinder, ellipsoid and hyperboloid). The images have the sizes 128x128 points for a single surface images and 192x192 for images with multiple surface types. The test data were formed by adding uncorrelated Gaussian distributed noise with zero mean and variable standard deviation (for the different noise levels) to the images. The images are not smoothed. Figures 11 until 16 show the minimum percentage of mislabelled points for the best possible threshold for a certain noise level. Notice that we can not do this experiment with a single planer surface, since the best threshold for a planer surface is ∞ and therefore the number of mislabelled points is always 0.

Mislabelled points are increasing along with the noise level. The noise has an impact on the data after a noise level of 10^{-3}. One could think that the noise would impact the data already at a level of 10^{-4} (resolution of some laser scanners). Because of the synthetic nature of the images, the image data is concentrated at a single point in the κ_1/κ_2-graph. Therefore, the noise does not have as strong of an impact on the results. The best threshold for the surfaces is zero, except for the cylindrical surface in conjunction with the HK algorithm, because the classification region of the cylinder vanishes when using the zero value. The HK algorithm has an advantage with the hyperboloid and the ellipsoid surfaces (see Figure 11 & 12), because the algorithm can shrink the cylindrical region. This creates a bigger classification region where noisy points are still recognised as the correct shape. We performed two experiments with two different cylindrical data sets. One data set contains a circular cylinder with uniform curvedness. The other data set contains four different curved cylinders. The HK classification has an advantage with the uniform curvedness (see Figure 13). This result certainly depends on the fixed shape border of the SC classification. With the other data set (see Figure 14), the performance differences are slightly smaller at lower noise levels.

We did one more test with an image that contains equal large plane, cylinder, ellipsoid and hyperboloid surfaces. In images with multiple surfaces the borders between the surfaces cause irregular shapes. To avoid the labelling of the irregular shapes the shape borders are eliminated in the ground truth file and have therefore no impact on the calculation of the mislabelled points. This reduced the total number of mislabelled points in comparison to the previous test images. We performed the experiment with one test image that contains the four shapes and another test images which contains the four shape regions but the three

curved surfaces are each replaced by four different cylinder, ellipsoid and hyperboloid surfaces. The SC segmentation performs slightly better in the first test image (see Figure 15). The methods perform differently in the second test image (see Figure 16), where the HK algorithm is affected by the noise at much lower levels than in the single surface images. The performance of the SC method is significantly better over all noise levels.

4 Conclusion

The performance of both methods on images containing single surfaces is basically the same. One difference is that the HK segmentation is using the zero-threshold to recognise cylindrical surfaces. What leads to the effect that cylinder points vanish from the image at low thresholds. Therefore, the SC classification is more stable at low thresholds on scenes containing cylinders. For the SC algorithm a slightly higher threshold should be used for the same error rate as the HK algorithm (but this effect is largely unimportant in terms of classification performance). In our noise tests, the SC algorithm can deal better with image noise in images which contain different surfaces, because the HK segmentation cannot focus with an optimal threshold on a single surface. Thus we conclude Koenderink's SC classification scheme has a slight advantage (5-10% lower error rate) when dealing with real scenes containing multiple surfaces and moderate noise.

References

[1] P.J. Besl. *Surfaces in range image understanding*. Springer, 1988.

[2] P.J. Besl and R.C. Jain. Invariant surface characteristics for 3-d object recognition in range images. *Comput. Vision Graphics Image Proc.*, 33:33–80, 1986.

[3] L.-D. Cai. *Scale-based surface understanding using diffusion smoothing*. PhD thesis, University of Edinburgh, Jan 1990.

[4] P. Clarysse, D. Friboulet, and IE. Magnin. Tracking geometrical descriptors on 3-d deformable surfaces: Application to the left-ventricular surface of the heart. *IEEE Transactions on medical imaging*, 16(4):392–404, 1997.

[5] C. Dorai and A.K. Jain. Cosmos - a representation scheme for 3d free-form objects. *IEEE Trans. Pat. Anal. and Mach. Intel.*, 19(10):1115–1130, 1997.

[6] A. Hoover etal. An experimental comparison of range segmentation algorithms. *IEEE Trans. Pat. Anal. and Mach. Intel.*, 18(7):673–689, July 1996.

[7] N. Guid, C. Oblonsek, and B. Zalik. Surface interrogation methods. *Computers and Graphics*, 19(4):557–574, 1995.

[8] Y. Kawata, N. Niki, H. Ohmatsu, R. Kakinuma, K. Eguchi, M. Kaneko, and N. Moriyama. Quantitative surface characterization of pulmonary nodules based on thin-section ct images. *IEEE Transactions on nuclear science*, 45(4):2132–2138, 1998.

[9] J. Koenderink and A. van Doorn. Surface shape and curvature scales. *Image and vision computing*, 10(8):557–565, 1992.

[10] E. Trucco and R. Fisher. Experiments in curvature-based segmentation of range data. *IEEE Trans. Pat. Anal. and Mach. Intel.*, 17(2):177–182, Feb 1995.

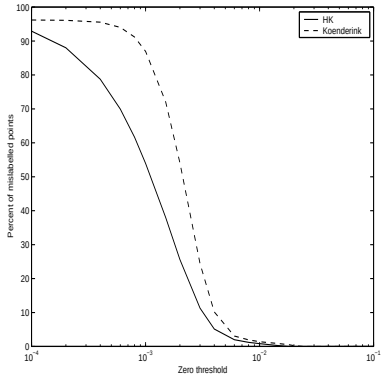

Figure 3: Mislabelled points versus zero threshold for the plane surface.

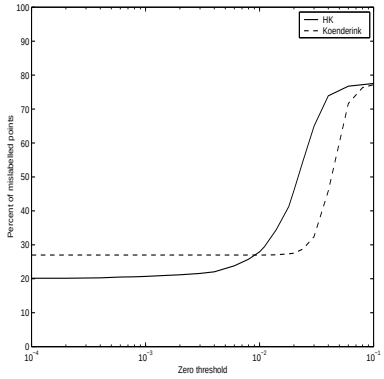

Figure 4: Mislabelled points versus zero threshold for the hyperboloid surface.

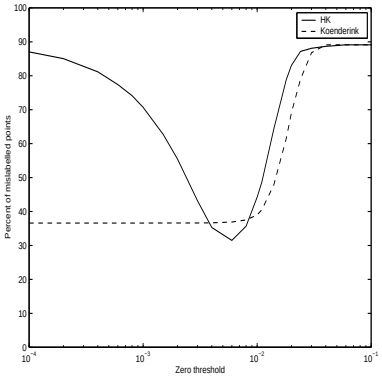

Figure 5: Mislabelled points versus zero threshold for the convex cylinder surface.

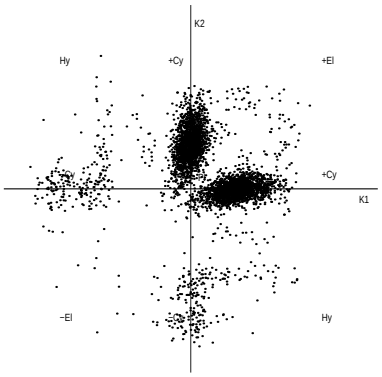

Figure 6: Scatter plot for the convex cylinder.

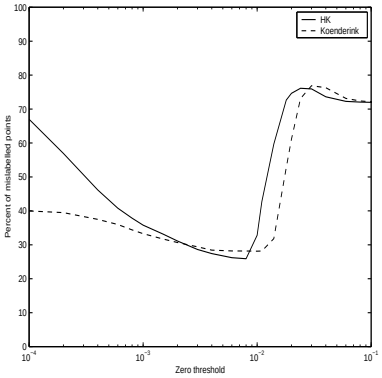

Figure 7: Mislabelled points versus zero threshold for the bomb.

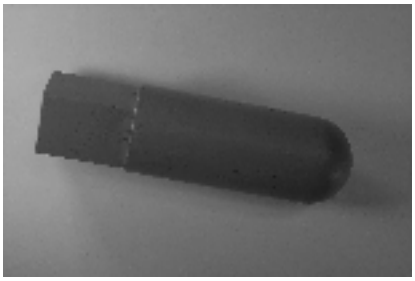

Figure 8: The bomb.

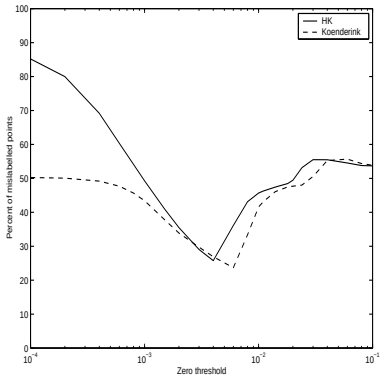

Figure 9: Mislabelled points versus zero threshold for the factory scene.

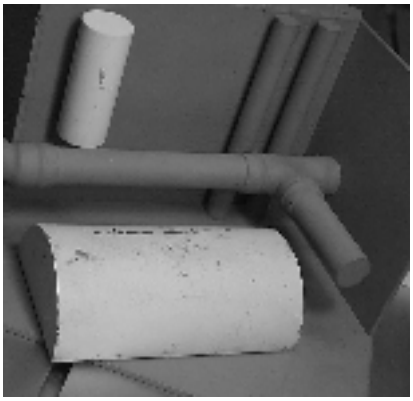

Figure 10: The factory scene.

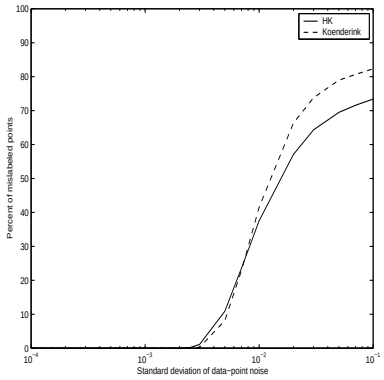

Figure 11: Mislabelled points versus noise level for the convex ellipsoid surface.

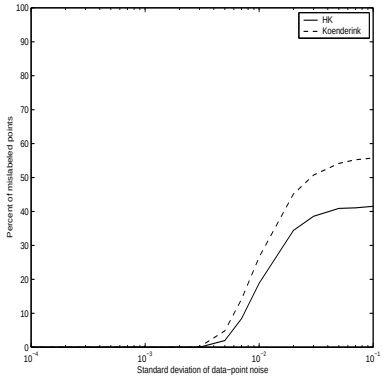

Figure 12: Mislabelled points versus noise level for the hyperboloid surface.

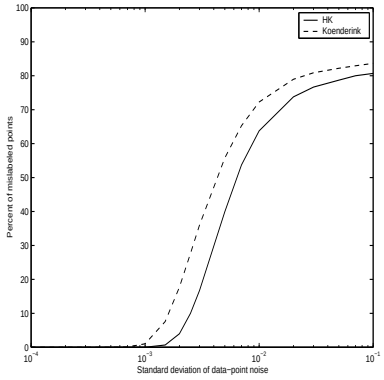

Figure 13: Mislabelled points versus noise level for one cylinder surface.

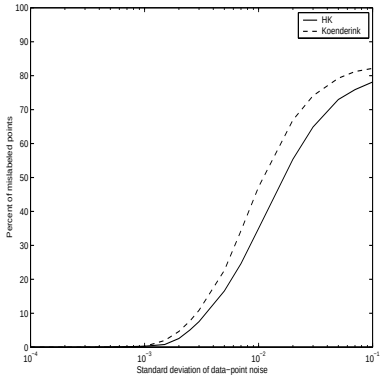

Figure 14: Mislabelled points versus noise level for four different cylinder surfaces.

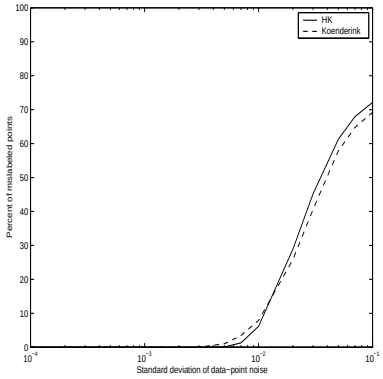

Figure 15: Mislabelled points versus noise level for an image which contains a single large equal-area large plane, cylinder, ellipsoid and hyperboloid surface.

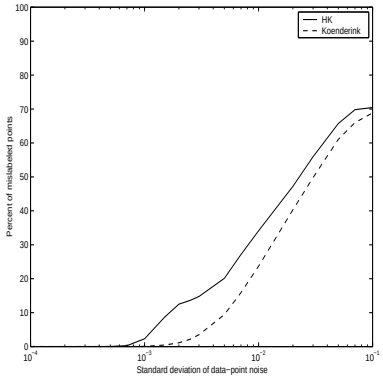

Figure 16: Mislabelled points versus noise level for an image with the single curved surfaces from figure 15 substituted with four differently curved surfaces for the cylinder, ellipsoid and hyperboloid regions.