

Classifying Materials from their Reflectance Properties

Peter Nillius and Jan-Olof Eklundh

Computational Vision & Active Perception Laboratory (CVAP)
Department of Numerical Analysis and Computer Science
Royal Institute of Technology (KTH), S-100 44 Stockholm, Sweden
{nillius,joe}@nada.kth.se

Abstract. We explore the possibility of recognizing the surface material from a single image with unknown illumination, given the shape of the surface.

Model-based PCA is used to create a low-dimensional basis to represent the images. Variations in the illumination create manifolds in the space spanned by this basis. These manifolds are learnt using captured illumination maps and the CURET database. Classification of the material is done by finding the manifold closest to the point representing the image of the material.

Testing on synthetic data shows that the problem is hard. The materials form groups where the materials in a group often are mis-classified as one of the other materials in the group. With a grouping algorithm we find a grouping of the materials in the CURET database. Tests on images of real materials in natural illumination settings show promising results.

1 Introduction

The appearance of a surface depends on its shape, the illumination and the material of the surface. In a normal vision task none of these properties are known a priori. Despite that, human observers are very good at determining the material of an object, even in the absence of texture. The estimation is done purely based on the reflectance properties of the surface. We will explore if this can be done computationally when there is no knowledge about the illumination, but the shape of the object is known.

Recent research [2, 11, 10, 8] has shown that the reflected light from a Lambertian surface can be represented with a low-dimensional model although the variations in illumination are infinite. This is because the surface acts as a low-pass filter on the incident illumination, making the images in practice lie in a low-dimensional subspace. Other work, [12, 9], indicates that this holds for many other types of surface reflectance, e.g. many of the materials in the CURET database, [3].

In this paper we classify the material of an object of known shape from a single image, when the illumination is unknown. In [4] Dror et al recognizes materials under similar assumptions. They use histograms of filter responses

and rely on the structure of the specular reflections to classify the material. Our approach is different in that we represent the images using a generative model, allowing us to discriminate between materials without specular reflections such as felt and velvet.

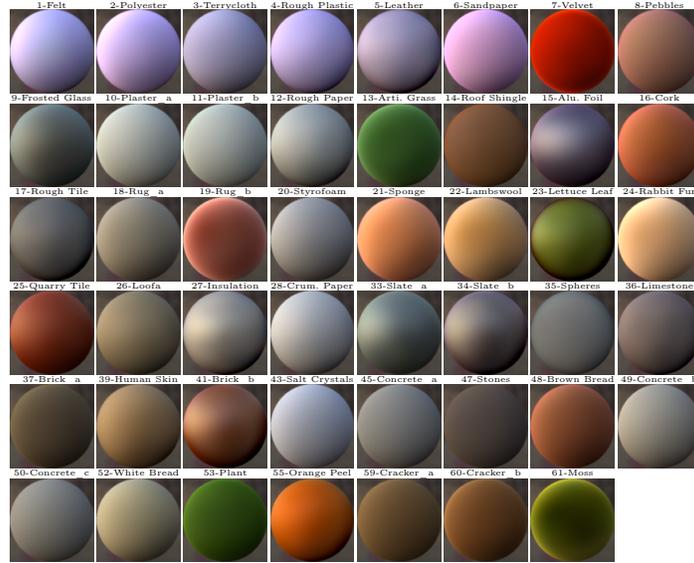


Fig. 1. Rendered images using BRDFs from the CURET database. Classifying materials from their reflectance properties can be very hard, as is in this case. If you disregard the color many materials look very similar.

2 A Low-Dimensional Generative Model for Image Irradiance

To find a basis to represent the images, we use the framework described in [9]. With this framework we can, for a given shape, construct a low-dimensional basis that can represent the images of an object of a wide variety of materials under more or less arbitrary illumination.

The basis is created using model-based PCA. Rather than performing PCA on a set of captured images, the PCA is analytically derived from the image formation model. This makes it possible to create a basis for a wide variety of conditions using a set of captured illumination maps and a database of surface reflectance functions (BRDF's). The illumination maps are undergoing all 3D rotations to take into account every possible lighting configuration.

The BRDF acts as a low-pass filter on the incident illumination, making the reflected light band-limited. Hence, by formulating the image formation in fre-

quency space we can derive a finite dimensional model of the image irradiance even when the illumination is unknown and arbitrary. As in [2, 11] the illumination is represented by its spherical harmonics coefficients, L_l^m . The BRDF is represented by its coefficients, b_{op}^q , in the basis by Koenderink and van Doorn, [5], based on the Zernike polynomials. Computing the image irradiance using these representations leads to a basis for image irradiance E . At this point we approximate the camera projection as orthographic which makes the image irradiance uniquely determined by the surface normal (α, β) . This results in the representation

$$E(\alpha, \beta) = \sum_i c_i E_i(\alpha, \beta) \quad (1)$$

where $c_i = L_l^m b_{op}^q$ (l, m, o, p and q are given by i due to an ordering of the basis functions). The E_i s are the image irradiance basis functions and are products of the Wigner D-functions (for real spherical harmonics) and the Zernike polynomials. See [9] for their explicit form.

The basis can represent the image irradiance from any isotropic surface under any illumination. In the general case an infinite number of basis functions are needed, but for many materials which act as low-pass filters on the illumination, the sum can be truncated and still be an accurate representation. This finite representation allows us to analytically derive the principal components. The variations in the illumination and surface reflectance properties are described by the covariance matrices of their respective coefficients, L_l^m and b_{op}^q . The resulting principal components are linear combinations of the basis functions E_i . See [9] for details.

When we use the PCA basis to represent images we assume that the illumination is the same for each point in the image. This assumption is true if the light source is distant. It is also necessary that there are no cast shadows or local inter-reflections, which is true when the object shape is convex. For non-convex objects this model is an approximation, where the quality of the approximation depends on the concavities and the material/illumination conditions. For instance, bright objects will have stronger inter-reflections than dark objects.

An important property of model-based PCA is that we can relate the principal components to the properties of the illumination and surface reflectance. There is an explicit relation between the coefficients in the PCA basis to the coefficients of the illumination and the BRDF. From the coefficients of illumination and BRDF, the coefficients, \mathbf{d} , in the PCA space can be computed by a simple matrix product

$$\mathbf{d} = \mathbf{A}\mathbf{c}. \quad (2)$$

where the elements of \mathbf{c} are $c_i = L_l^m b_{op}^q$. The PCA basis as well as the matrix \mathbf{A} are computed from the shape of object and the variations in the illumination and BRDF.

Another important aspect of the PCA basis regards robustness. The first basis function is selected so that it maximizes the signal variance of the component it represent. The subsequent basis functions maximize the same variance while being orthogonal to all the previous functions. This means that the first basis

function has the highest signal-to-noise ratio (SNR), on average, hence being the most robust component to estimate. The following components will be less and less robust to estimate. In other words, selecting the number of basis functions to use is not just a question of saving computer memory and computation time, but also a question of robustness and regularization.

3 Material Recognition

Our approach to material recognition is to estimate the coefficients in the basis described in the previous section from the images and compare them to a database of known materials.

Since the illumination is not known we cannot calculate what the corresponding coefficients should be for the materials in the database. We need to take into account all possible illuminations and find the illumination-material pair that best matches the image. For this to be possible it is necessary that the variations in the coefficient space are much smaller than the variations in the illumination (which are infinite). If this is true we can learn the variations in the coefficient space with only a limited amount of training illuminations.

Smooth variations in the illumination result in a manifold of points in the coefficient space. To learn these manifolds we take a set of illumination maps and rotate them over all rotations. To store the manifolds we sample them by sampling the rotation group $SO(3)$ and calculate the coefficients for each sample point for every illumination map and material.

The image is classified by finding the manifold which is closest to the point representing the image. The procedure is very much the same as in [7].

3.1 Learning the Manifolds

The manifold for each material is learnt from a set of illumination maps that are rotated over the full rotation group. The rotation group is sampled and for each rotation (α, β, γ) the spherical harmonic coefficients of the rotated illumination map are calculated. The point on the manifold is given by equation (2).

To sample the rotation group we sample the surface of a sphere and combine it with a circle. The sphere is sampled by starting from an icosahedron inscribed in the sphere. The icosahedron is recursively subdivided by projecting the mid-point of each edge onto the surface of the sphere forming four new triangles for each old triangle, [1]. The circle is sampled at a density as close as possible to the sampling of the sphere.

3.2 Finding the Closest Manifold

To find the closest manifold to a point we simply go through all points on each manifold and calculate the distance to the point to be classified. The distance measure is the sum of squared differences in coefficient space.

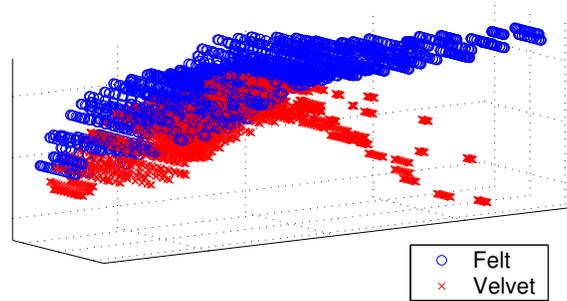


Fig. 2. Sampled manifolds in the coefficient space of materials 1-Felt (blue rings) and 7-Velvet (red crosses) under one of the illumination maps undergoing all 3D rotations, $SO(3)$.

To aid our algorithm in being illumination invariant we take a number of steps. The first element of the point is discarded. It corresponds to the constant function of the basis and captures the variations in the ambient component of the illumination. By discarding it the algorithm becomes independent of such variations.

The remaining elements are normalized to get brightness independence. It also means that we will not be able to differentiate between bright and dark materials, although this could be added at a later stage by comparing the signal variances of the images.

4 Discrimination of Materials in the CURET database

Before we move on to real images we need to assess what can be done. How well can materials be discriminated from their reflectance properties alone? Figure 1 demonstrates that many materials look similar to the human eye.

To test this we will analyze how well the materials in the CURET database can be discriminated in synthetic images, i.e. when there is no noise. The illumination is considered to be unknown. The algorithm is tested on images generated from one of the illumination maps, while the other illumination maps are used to build the manifolds for classification. This is repeated for all nine illumination maps (the leave-one-out principle).

We don't actually need to generate any images. Using the low-dimensional basis framework described in Section 2 we can directly from the illumination and material coefficients compute the coefficients in the low-dimensional basis of the image. This allows for extensive testing. Each of the 48 materials used is tested with nine illumination maps, each under 462 different rotations, summing up to a total of 200 000 images used for testing.

Figure 3 shows the classification rates for the different materials. The correct classification rates, which can be seen in the diagonal, range between 5 and 80 percent. Materials with a high classification rate are 7-velvet and 61-moss which

have particular reflectance properties. Glossy materials have in general a higher recognition rate than matte materials.

What is interesting is that the materials seem to form groups, where a material in a group systematically is being mis-classified as one of the other materials in that same group. This becomes apparent when we order the materials in a particular way. Figure 4 shows the exact same classification rates as Figure 3, but with the materials ordered using a hierarchical grouping algorithm that will be described in the next section. We begin to distinguish blocks in the diagonal of the matrix. There is a large block of matte materials in the top left corner, formed by the materials 1-Felt, 20-Styrofoam, ..., 24-Rabbit Fur. Following the matte materials is a group of glossy materials, 4-Rough Plastic, ..., 15-Foil. Last comes 7-Velvet and a group of velvet-like appearance (asperity scattering), 13-Art. Grass, 19-Rug.b and 61-Moss. Finally we have 35-Painted Spheres which forms a group of its own.

4.1 Visual Grouping of the Materials

It is clear that we cannot expect to distinguish between some of the materials in the CURET database. Instead we can try to find groups in which to classify the materials.

Using the matrix containing the classification rates we group the materials. The grouping is done in a greedy fashion. We start with groups of single materials. Then the two groups that maximize the average recognition rate are joined. This is repeated until the desired number of groups is reached. To select the number of groups one can e.g. look at the ratio between the recognition rates and the rate of selecting the correct material by chance.

Dividing the CURET database into 9 groups results in the grouping in Figure 5. We have labeled the groups according to the characteristics of their members. All matte materials end up in one group. Materials having specular reflectance are split up in three groups. The last five groups are materials that did not fit into any group. These materials have a high recognition rate on their own.

More or less all the groups are sometimes mis-classified as matte materials. This makes sense. In the testing we take all rotations of the illumination into account. This means that sometimes the dominant light source in the scene will be behind the object. Hence, there will be no specularity on the object to differentiate it from a matte material.

5 Classifying the Material in Real Images

To test the algorithm we glued five different real materials onto cylinders, see Figure 6. Cylinders were chosen due to the difficulty of gluing a non-stretchable materials onto a sphere. The cylinders were photographed using a digital camera in different illumination conditions, including outdoor sunny, outdoor cloudy and indoor with indirect light from a window. Before classification the images

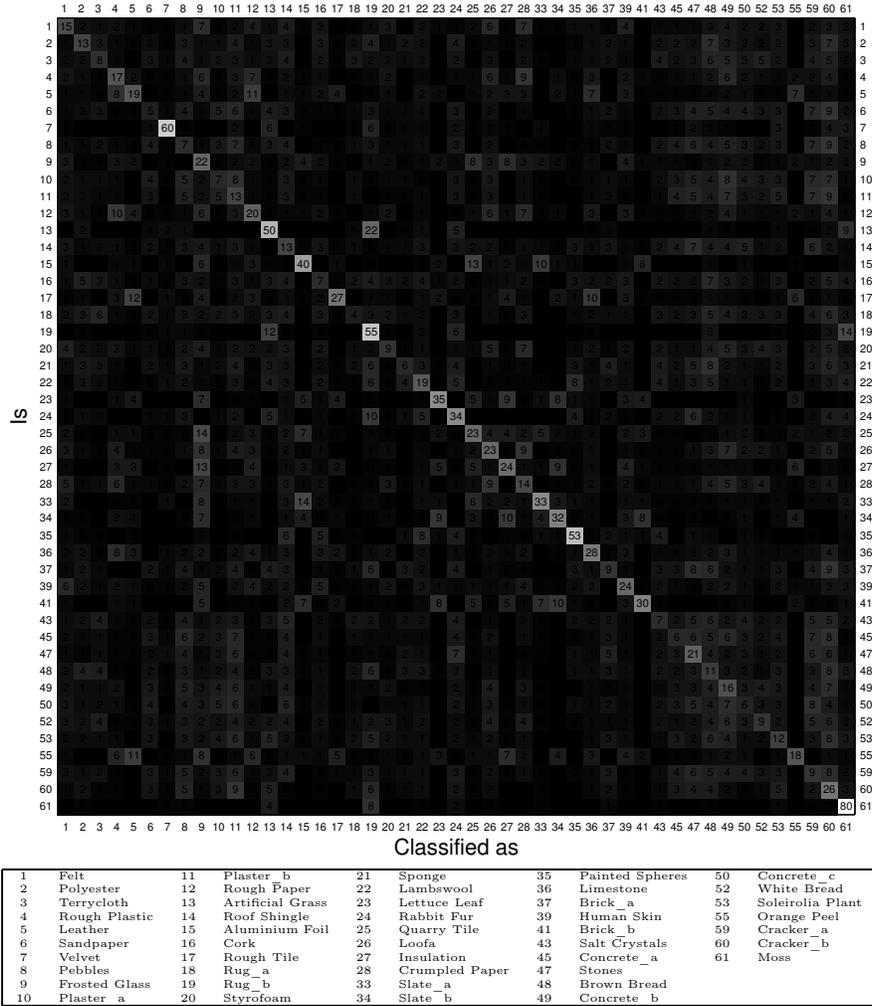


Fig. 3. Recognition rates for the CURET materials. Each row shows the classification rates for a particular materials, e.g. the leftmost element in the first row is the rate that material no. 1 is classified as materials no. 1, the second element is the rate the material one is classified as material no. 2. The diagonal is the correct classification rate. These results are discussed more in the text.

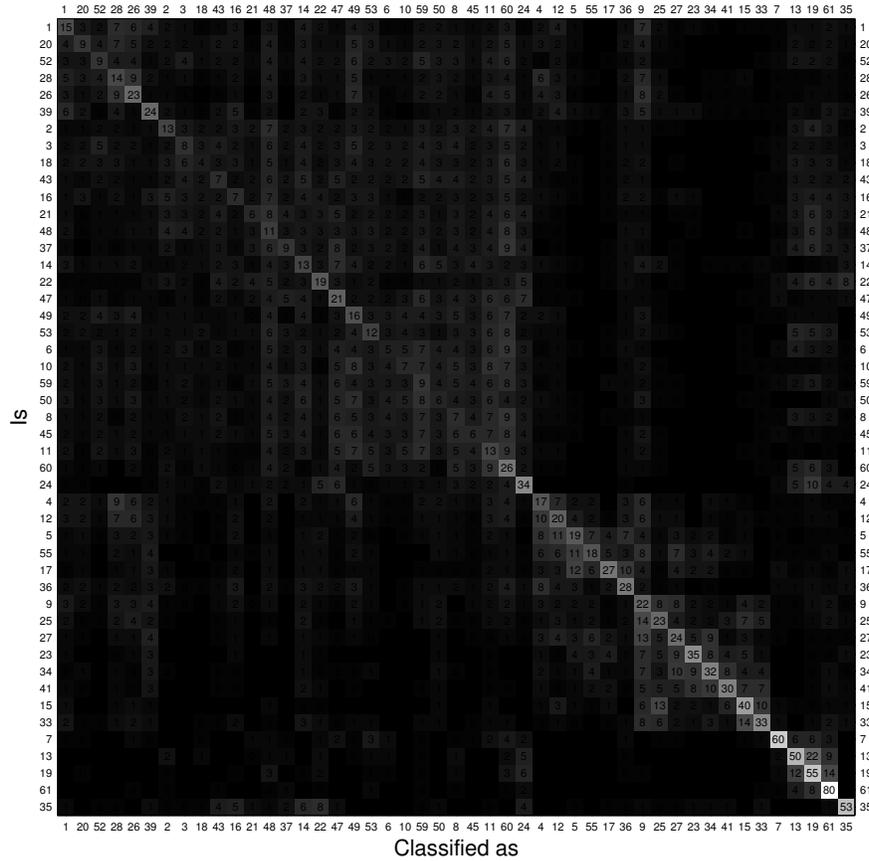


Fig. 4. When the classification rates from Figure 3 are sorted in a particular way a pattern emerge. The materials form groups. Materials within a group often are classified as one of the other materials in the same group. The largest group can be seen as a grey block in the top left corner of the matrix. These are the matte materials, 1-Felt, 20-Styrofoam, ..., 24-Rabbit Fur. After the matte materials comes a group of more glossy materials, 12-Rough Plastic, ..., 36-Limestone. Next comes a group of shiny materials 9-Frosted Glass to 33-Slate.a. Last is a group of materials with asperity type scattering, 7-Velvet, 13-Artificial Grass, 19-Rug.b and 61-Moss.

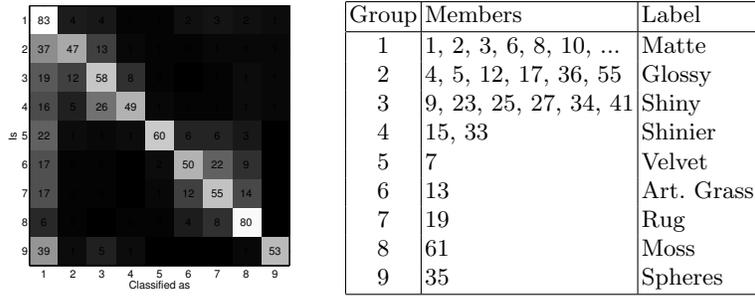


Fig. 5. Classification rates when the materials are grouped into nine groups. Not all members were listed in the matte group due to space limitations, but this group contains all materials that are not in the other groups.

were radiometrically calibrated, using the method in [6]. The geometry of the cylinders were estimated by manually marking where in the image the cylinders were.

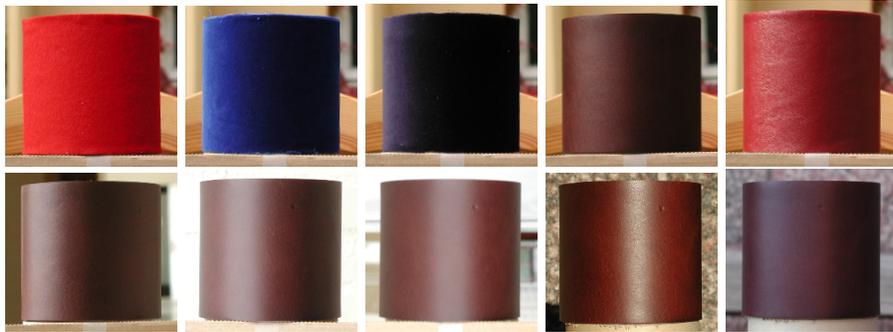


Fig. 6. The algorithm was tested on images of cylinders with the pieces of five different real materials glued onto them. Top row from left to right: felt, velvet 1, velvet 2, leather and imitation leather. Bottom row: leather in five of the different illumination conditions.

Using the framework from Section 2 we computed a basis for the cylinder. A total of six basis functions were used in the experiments. The coefficients for the image were estimated by projecting the image onto the basis. The image was then classified by finding the closest manifold as described in Section 3. The manifolds were this time learned using all nine illumination maps.

Figure 7 shows some of the images being classified. Note how well the basis is able to represent the image irradiance in all cases.

A total of 84 images were used in the experiment. Table 1 summarizes the results. As predicted by the synthetic experiments only a few of the images

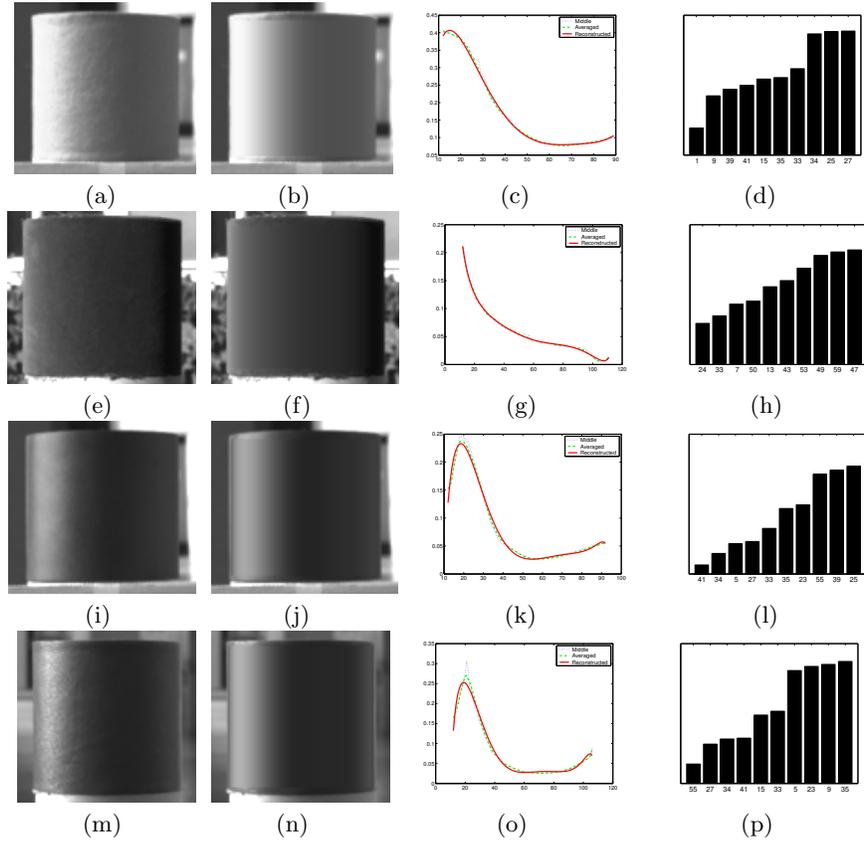


Fig. 7. Examples of classified images: (a)-(d) images for Felt. (a) calibrated gray image, (b) reconstructed gray image (this is what the algorithm “sees”), (c) image and reconstructed intensity profiles. (d) distances to the ten closest materials. Here the material is correctly classified as felt. (e)-(h) show the same images for Velvet 1. The material is here incorrectly classified as 24-Rabbit Fur, 7-Velvet comes third place. (i)-(l) images for leather which in this case is classified as 41-Brick_b, 5-Leather is the third closest material. (m)-(p) imitation leather: classified as 55-Orange, 5-Leather on seventh place. Notice how well the basis represent the irradiance for the different cases.

where correctly classified on an individual basis. Felt and the two velvets have a recognition rate of 5% to 7.7%, which is still several times greater than chance, which is $1/48 \approx 2.1\%$. When using the grouping in Figure 5 the recognition rates are higher. Felt is to a large extent classified as matte. The leather here is classified as Shiny or Shinier, while the leather in the database is categorized as Glossy. This could be because our leather is shinier than the leather in the database. Visually, at least, it appears so. The imitation leather is also mostly classified as Shiny or Shinier.

So far the results match the synthetic results fairly well. The velvet however does not. The synthetic results indicate that velvet should be fairly easy to recognize, but in our experiments the two velvet cylinders are mostly classified as matte. On the other hand, they are also often classified as one the groups Grass, Rug and Moss, which have the same type of surface reflectance as Velvet.

Material	Correct	Matte	Glossy	Shiny	Shinier	Velvet	Grass	Rug	Moss	Spheres
Felt	7.7	77	7.7	0	15	0	0	0	0	0
Leather	0	25	6.2	44	19	6.2	0	0	0	0
Im. Leather	0	10	0	40	35	10	0	0	0	5
Velvet 1	5	55	0	5	5	5	10	10	10	0
Velvet 2	6.7	40	0	33	0	6.7	0	13	6.7	0

Table 1. Classification Rates for the Cylinder Images

6 Conclusions

We have investigated the problem of classifying the surface material from a single image with unknown illumination, given the surface shape.

Recognizing materials from their reflectance properties is hard. We cannot expect to distinguish between many of the materials in the CURET database. Instead we should find groups in which to classify the materials.

The grouping produced by our algorithm suggests that we can expect to distinguish between matte materials, special materials such as velvet and materials of different grades of shininess.

Acknowledgments

This work was done within the EU-IST project IST-2000-29688 Insight2+. The support is gratefully acknowledged.

References

1. D.H. Ballard and C.M. Brown. *Computer Vision*. Prentice-Hall, 1982.

2. R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25(2):218–233, February 2003.
3. K.J. Dana, B. van Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, January 1999.
4. R.O. Dror, E.H. Adelson, and A.S. Willsky. Recognition of surface reflectance properties from a single image under unknown real-world illumination. In *Workshop of recognizing objects under varying illumination*, 2001.
5. J.J. Koenderink and A.J. van Doorn. Phenomenological description of bidirectional surface reflection. *J. Optical Soc. of Am. A*, 15(11):2903–2912, November 1998.
6. T. Mitsunaga and S.K. Nayar. Radiometric self calibration. In *Proc. Computer Vision and Pattern Recognition*, pages I: 374–380, 1999.
7. H. Murase and S.K. Nayar. Visual learning and recognition of 3-d objects from appearance. *Int. Journal of Computer Vision*, 14(1):5–24, January 1995.
8. P. Nillius and J.O. Eklundh. Low-dimensional representations of shaded surfaces under varying illumination. In *Proc. Computer Vision and Pattern Recognition*, pages II:185–192, 2003.
9. P. Nillius and J.O. Eklundh. Phenomenological eigenfunctions for image irradiance. In *International Conference on Computer Vision*, pages 568–575, 2003.
10. R. Ramamoorthi. Analytic pca construction for theoretical analysis of lighting variability in images of a lambertian object. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(10):1322–1333, October 2002.
11. R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: determining the illumination from images of a convex lambertian object. *J. Optical Soc. of Am. A*, 18(10):2448–2458, October 2001.
12. R. Ramamoorthi and Hanrahan P. A signal-processing framework for inverse rendering. In *SIGGRAPH*, 2001.