Material Recognition: Bayesian Inference or SVMs?

Ishrat Badami*
Supervised by: Michael Weinmann,†
Reinhard Klein‡

Institute of Computer Graphics
University of Bonn

Abstract

Material recognition is an important subtask in computer vision. In this paper, we aim for the identification of material categories from a single image captured under unknown illumination and view conditions. Therefore, we use several features which cover various aspects of material appearance and perform supervised classification using Support Vector Machines. We demonstrate the feasibility of our approach by testing on the challenging Flickr Material Database. Based on this dataset, we also carry out a comparison to a previously published work [Liu et al., "Exploring Features in a Bayesian Framework for Material Recognition", CVPR 2010] which uses Bayesian inference and reaches a recognition rate of 44.6% on this dataset and represents the current state-of-the-art. With our SVM approach we obtain 53.1% and hence, significantly outperform this approach.

Keywords: Material recognition, Texture classification, SVMs

1 Introduction

Understanding materials enables us to interact with the real world and influence our decisions in everyday life, e.g. where to drive a bike on a wet muddy road or whether a fabric in textile shop is smooth enough for cushion cover. These daily examples show the importance of material recognition for humans. In the fields of computer vision and computer graphics, one goal is to develop systems which can automatically perform this task. Identifying the respective material of object surfaces for instance allows to handle the object appropriately within a supply chain or to select the corresponding appearance properties for photo-realistic rendering.

For humans material recognition comes naturally. Since one can touch and feel the material surface if it is smooth or rough, hard or soft, take a look from different directions, from close or far distance and observe if it is shiny or dull. The key observation is that the visual appearance of a surface in an image depends on several different factors such as the illumination conditions, the geometric structure of the surface sample at several spatial scales, and the surface reflectance properties, often characterized by the bidirectional reflectance distribution function (BRDF) [23] and its variants [10, 15, 24].

Figure 1: We used Flickr Material Database [27]. This database captures a wide range of appearance of 10 different materials.

In order to capture such characteristics recent investigations [20] combine a large number of different low-level and mid-level features, which are commonly used in related areas such as object and texture recognition tasks, in a Bayesian framework. The authors demonstrated that their approach outperforms previous state-of-the-art methods [29] on a more challenging database [27].

Support Vector Machines (SVMs) [28, 7] have become popular for classification tasks, since they offer advantages such as, ease of generalization of the problem, its ability

*badami@informatik.uni-bonn.de
†mw@cs.uni-bonn.de
‡rk@cs.uni-bonn.de

to handle high-dimensional feature spaces and the absence of local minima [3]. As we want to focus on the comparison between the classification using the Bayesian approach in [20] and one based on SVMs, we combine the idea of using the image features from [20] within a SVM framework and exhaustively compare the achieved classification rates to the ones reported in [20]. For this, we also evaluate our approach on the challenging MIT Flickr Material Database [27]. We observe that with our system the recognition rate improves from 44.6% to 53.1%. We also evaluate our system on the KTH_TIPS2 dataset [4].

The rest of this paper is organized as follows: In Section 2, we describe the present state of methods used in the area. In Section 3, we introduce feature pools used for classification. In Section 4, we explain the support vector machine classification model in the context of material recognition. Finally, in Section 5, we examine our system on the Flickr Material Database and discuss the results. In Section 6, we conclude.

2 Previous work

Learning high-level material categories such as foliage, stone or metal is related to object and texture classification but differs in several aspects. Several approaches address material recognition by focusing on purely texture based image features.

Texture has been defined in terms of dimensions like periodicity, orientedness and randomness [21]. A recent work on 3D textons [19] addresses material recognition using multiple images of varying viewpoint and lightning conditions. Cula and Dana [8] adapted the method of [19] to 2D textons where each histogram is obtained from a single image in the training set. For their evaluation they used the CUReT database [10] consisting of images of 61 different texture samples under 205 different viewing and illumination conditions. A high classification rate of more than 95% is reported in [29] with 2D textons on the CUReT dataset using the NN classifier. In [4] it is shown that the SVM based classifier achieves 98.5% accuracy on the KTH_TIPS2 [4] database consisting of 11 material categories with 4 texture samples in each category photographed under various conditions. Although texture is a characterizing feature cue it is not sufficient for representing material properties completely. It might happen that an object made of different materials has a similar or even the same texture.

Although information about objects can lead us to the right guess concerning the material from which it is made of, sometimes it is really misleading. For example a cup can be made of plastic, metal or glass. In case of an artistic cup it can be carved out of wood or stone as well. This demonstrates the difference between material recognition and object recognition.

The appearance of a material in an image highly depends on the environment illumination and surface reflectance properties described by the BRDF. Material recognition might be trivial in case of a known BRDF. But it is very difficult to estimate the BRDF of the material from a single image without simplifying assumptions [11].

Moreover, the appropriate choice regarding the classification method is also an influencing aspect. While a few approaches make use of a Nearest Neighbor (NN) classifier (e.g. [29]), the methods we consider as state of the art rely on a Bayesian framework [20] or SVMs [4]. SVMs have been proven to consistently achieve good performance in complex real-world problems such as text [16, 12] and image classification [6] and bioinformatics [30] and biosequence [2] analysis. This motivates us to involve SVMs as classifier. However, most of the methods that address material classification are evaluated on datasets such as [10] that are not well-suited for this task as they do not contain the large variations in appearance which occur in real-world scenes and for this, classification rates usually are very high. Hence, the real performance differences cannot be seen in a reliable way. The reason for this is that they have been acquired using a controlled setup. In contrast, [20] show the performance of a Bayesian approach on the significantly more challenging MIT Flickr Material Database [27]. However, they only compare their method against the NN classifier and not against SVMs. We believe that a real comparison between different material classification approaches has to be carried out on such a challenging dataset. Hence, we compare our SVM based technique to the method described in [20].

3 Feature Extraction

For the development of a reliable image based material recognition system it is important to consider image features which are representative and discriminative. However, it seems to be impossible to capture the variety of material characteristics in a single feature descriptor, as commonly used descriptors usually are restricted to a certain material property such as color, texture or reflectance behavior. As the different material properties are not equally descriptive for different material classes, a variety of features have to be considered in order to derive information about the different materials of an object. If the object appears shiny one might think that it is made of glass or metal whereas an object surface covered by minute fibers appear rough and together with the underlying weave pattern leads to a specific textured representation within an image. Wood is recognized usually with its brown color. In order to take several characteristic material properties into account, we follow the idea of [20] and use a pool of features which are covering different aspects of appearance. In general, for a fixed camera and object position, the image can be determined by 1) BRDF, 2) surface structure, 3) color, 4) object shape and 5) environment illumination. As we want to obtain hints on the performance of the SVM classifier in comparison to the Bayesian framework proposed in [20] we extract the same image features.

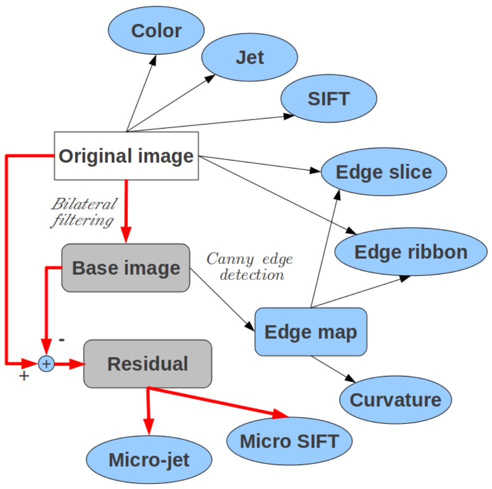

Figure 2: Features used in the classification [20].

Color is an important cue for recognizing materials. For example foliage is green, wood is usually brown and stone has less saturated color whereas fabric, plastic and paper have saturated color. To capture local color information, we store the RGB values in a local 3×3 neighborhood and concatenate them to a vector of dimension 27 as in [20].

Furthermore, every material has a peculiar texture. Wood has a ringing pattern, whereas fabric has a weaving pattern. Like [20] we use two sets of features to characterize texture. The first feature is SIFT [22] which is commonly used as a texture feature and also serves for tasks such as object and scene recognition. The second set of features is responses of an image through a set of Gaussian filters of different scale and orientation, also known as Jet [17]. We consider 2 Gaussian derivative filters at 3 scales and 6 orientations, i.e. $2 \times 3 \times 6 = 36$ rotational variant filters, and 8 Laplacian of Gaussian and 4 Gaussian filters, i.e. 8+4 = 12 rotational invariant filters. We combine all the filter responses in a single vector of size 48.

In addition, it is important to capture features not only on a meso-level but also on a micro-level. The human visual perception system can impressively abstract minute details of texture, e.g. smoothness of metal and glass surfaces, grains in paper and stone, the fibers of the fabric and crinkles in leather. For extracting micro texture of an image we follow the idea in [1] where the image is smoothed by a bilateral filter [13] and the residual is obtained by subtracting the smoothed image from the original image. The obtained residual image is used for further analysis as it reveals the information of texture on a finer scale. We derive descriptors that capture such micro details by computing SIFT and Jet over the residual image. For micro-texture, the Jet filter bank is evaluated on the same set of orientations but on a different set of scales in comparison to the Jet applied to the original image.

Materials can be molded to any arbitrary shape to create different objects, but still the outline shape of an object and its material category are often related, for e.g. fabric and glass have a curvy structure whereas metal, wood, stone can have straight edges and sharp corners. The edges and corners can be acquired from edge maps. We extract such edge maps by applying the canny edge detectors [14] to the base image. Furthermore, we only consider edges having a certain minimum length. Corresponding examples are shown in Figure 3. The curvature along these edges can be used to represent the orientation of the outline shape of a certain material, we calculate this specific descriptor by sampling the edges at three different scales (see Figure 4(a)), which results in a 3D-vector. As we are interested in a dense sampling, we calculate such a descriptor for every second pixel along the edges.

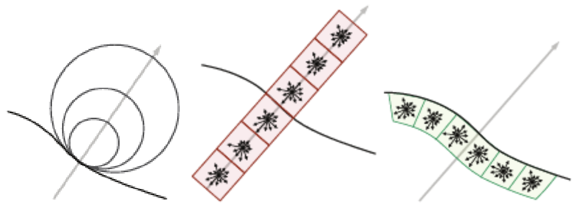

(a) Curvature (b) Edge-Slice(HOG) (c) Edge-Ribbon(HOG)

Figure 4: Curvature is calculated over three different scale, Edge-Slice and Edge-Ribbon are calculated in 6 cells [20] at edges.

Furthermore, reflectance behavior is also an important cue for classifying material categories. Water and glass are translucent, metal is shiny, wood and stone are dull and opaque. Such properties can be observed in form of distinctive intensity changes at the edges in an image. We follow [20] in computing histograms of oriented gradients (HOG) [9] in the vicinity of the edges. More precisely, we first select a slice of a certain width along the normal direction of the edge and compute the gradient at all of the pixels inside the slice. In order to calculate HOG, the slice is divided into 6 cells where the gradient orientation is quantized into 12 bins. We combine the histogram of all 6 cells in a vector of length 72, which will be referred as Edge-Slice [20]. In addition, we use the same method and employ a slice along the tangent direction of the edge in order to obtain the Edge-Ribbon feature [20](see Figure 4(b) and 4(c)).

So far, we described all the features which we use to characterize material appearance due to different properties. Figure 2 shows a flowchart how the features are generated. Among these features color, SIFT and Jet are low level features and can be calculated directly from the image. In contrast, curvature, Edge-Slice, Edge-Ribbon, micro-SIFT and micro-Jet are mid level features which depend on the edge map and the base image respectively (see Figure 3). In order to capture the relevant information appropriately, we calculate color, SIFT, micro-SIFT, Jet and micro-Jet on a evenly sampled grid in the image

(a) Original (b) Base image(bilateral filter) (c) Edge map(Canny) (d) Residual (a)-(b)

Figure 3: **Example images of how features are calculated in our system**. From top to bottom the rows show examples for fabric, glass and wood. On image (a) we apply bilateral filtering [13] to obtain the base image (b). We run the Canny edge detector [14] on the base image and compute edge maps. Curvature, Edge-Slice and Edge-Ribbon are extracted as a feature from the edge map. Subtracting (b) from (a), we get the residual image (d) that depicts micro structures of the materials which are captured by micro-SIFT and micro-Jet features.

areas where the material has been annotated. The remaining features are sampled at every second pixel along the edges.

4 Classification with SVMs

Once the features have been calculated we want to build a robust material recognition system. For this, we first apply a quantization of the features to form characteristics clusters whose centers are denoted as visual words. In the next step, we use SVMs for material classification based on the given inputs. In the following, we will describe these steps in more detail.

4.1 Feature quantization and visual words

Before we start classifying our features we need to group alike features to reduce the massive data into few representative visual words. From training images we estimate visual words which we can expect being present in the test data as well. We use k-means clustering for the quantization. After quantizing individual features into k visual words, the distribution of visual words per image is calculated for all of the different features by assigning each pixel in the image the nearest visual word index and calculating the histogram over the frequency of the visual words. Figure 5 shows some clusters for different feature types.

To generate a common visual word dictionary, for all the different types of features, suppose there are m features in the feature pool (e.g. color, SIFT, Jet) and m corresponding dictionaries $\{D_i\}_{i=1}^{m}$. Each dictionary has V_i codewords (e.g color has 150, SIFT has 250), i.e. $|D_i| = V_i$. Since the features are quantized separately the words generated by the i-th feature are $\{w_1^{(i)}, \ldots\ldots, w_{N_i}^{(i)}\}$, where, $w_j^{(i)}$ is an index representing j-th cluster center of i-th feature, $w_j^{(i)} \in \{1, 2, \ldots V_i\}$ and N_i is the number of words. In order to combine two features, the corresponding dictionaries are simply put together. For example, a document of m sets of words

$$\{w_1^{(1)}, \ldots., w_{N_1}^{(1)}\}, \{w_1^{(2)}, \ldots., w_{N_2}^{(2)}\}, \ldots\ldots, \{w_1^{(m)}, \ldots., w_{N_m}^{(m)}\} \quad (1)$$

can be combined to one set

$$\{w_1^{(1)}, \ldots., w_{N_1}^{(1)}, w_1^{(2)} + V_1, \ldots., w_{N_2}^{(2)} + V_1, \ldots\ldots, w_1^{(m)} + \sum_{i=1}^{m-1} V_i, \ldots., w_{N_m}^{(m)} + \sum_{i=1}^{m-1} V_i\} \quad (2)$$

with a joint dictionary $D = \cup_i D_i, |D| = \sum_{i=1}^{m} V_i$. In case of combining color ($V_1 = 150$) and SIFT ($V_2 = 250$) the

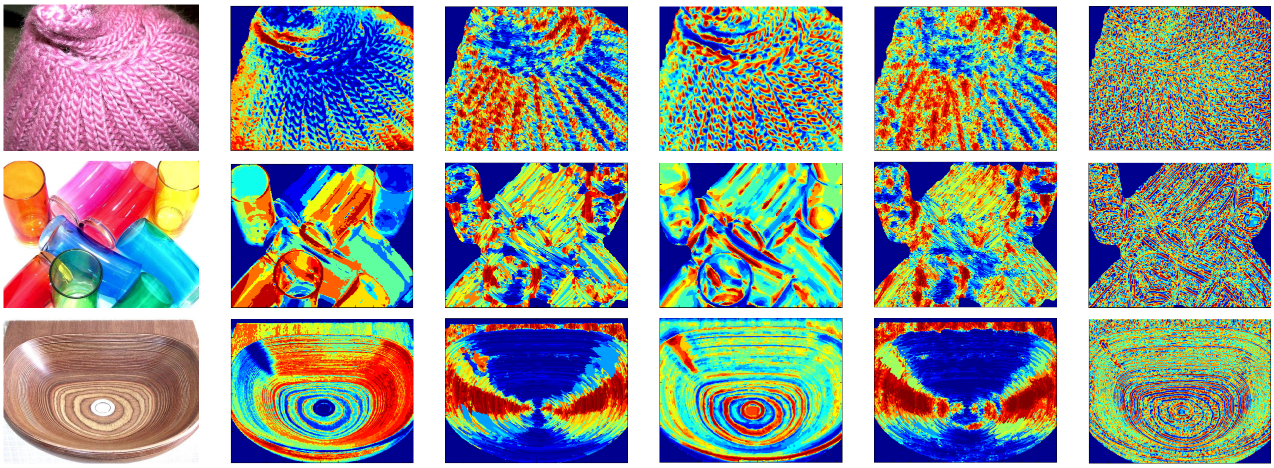

(a) Original (b) Color (c) SIFT (d) Jet (e) micro-SIFT (f) micro-Jet

Figure 5: **Visualization of quantized features**. After finding k cluster centroids of individual features, each pixel is assigned an index of the closest visual word. In order to visualize how the cluster centers are distributed in an image, the corresponding indices are color coded by RGB values. Same colors indicate that feature vectors corresponding to the pixels lie in the same cluster.

first 150 entries of the dictionary (of size $V_1 + V_2 = 400$) are codewords of color and the next 250 entries represent codewords of SIFT. This way we can reduce the multi-dictionary problem to a single dictionary problem.

4.2 Support Vector Machines

Being robust to noise [25] and being capable to generalize in case of a small training set [26], SVMs have become a popular and commonly used classifier for recognition tasks. A single SVM constructs a hyperplane or set of hyperplanes in a high-dimensional space and allows to distinguish between two linearly separable sets of samples. In order to deal with our multiple classes given in the used data sets, there is a need for using an SVM formulation which is capable of dealing with more than two classes. This can be achieved by either using several pairwise classifiers arranged in trees [18], where each of the nodes represents an SVM, or by using an one-vs-others approach, where multiple SVMs are trained and each of them separates a single class from all remaining ones. We follow the first strategy and use the implementation in [5].

As SVMs perform a supervised learning, objects with known class labels are used as samples for the training phase. We use the training data $\{\mathbf{x}_i, yi\}$ with $i = 1,\dots,l$ where $\mathbf{x}_i$ represents the histogram of visual words per image for a single feature type or a feature combination and y_i describes the corresponding material category in $\{fabric, foliage, \dots, water, wood\}$, i.e. $y_i \in \{1,2,\dots,10\}$. As kernel function, we use the Gaussian RBF kernel

$$K(x_i, x_j) = e^{\frac{-\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2}}. \tag{3}$$

5 Results

We run our system on two data sets namely MIT Flickr Material Database [27] and KTH_TIPS2 database [4].

First, we tested our system on the KTH_TIPS2 database. In this database there are 11 different materials namely *Crumpled aluminum foil, Cork, Wool, Lettuce leaf, Corduroy, Linen, Cotton, Brown bread, White bread, Wood, Cracker*. There are 44 different material samples present in the database in total. For each sample, images are taken at 9 scales, 3 poses and 4 different illumination conditions, hence there are $44 \times 9 \times 3 \times 4 = 4752$ images in this database.

We extract the same features and feature combinations as in [20] for this database and the results are plotted in Figure 8. The highest recognition rate is achieved by 99.4%.

As mentioned before, we do not consider this dataset to be challenging enough to derive statements on performance differences within complex scenes. For this, we consider the MIT Flickr Material Database [27], where there are 10 material categories namely *fabric, foliage, glass, leather, metal, plastic, paper, stone, wood, water*. Each category contains 100 images. 50 images show close-up views of the materials and 50 show an object made of the corresponding material. This dataset contains also annotations where this specific material is located in the image domain. Only pixels inside these areas are considered for feature calculation. For training, we randomly choose 50 images per category and test the system on the rest. For reliable results, we take 25 images showing close-up views and 25 images showing the full object made of the material. In addition, the training/testing process is repeated 5 times for different, randomly chosen training sets and the classification rates are averaged.

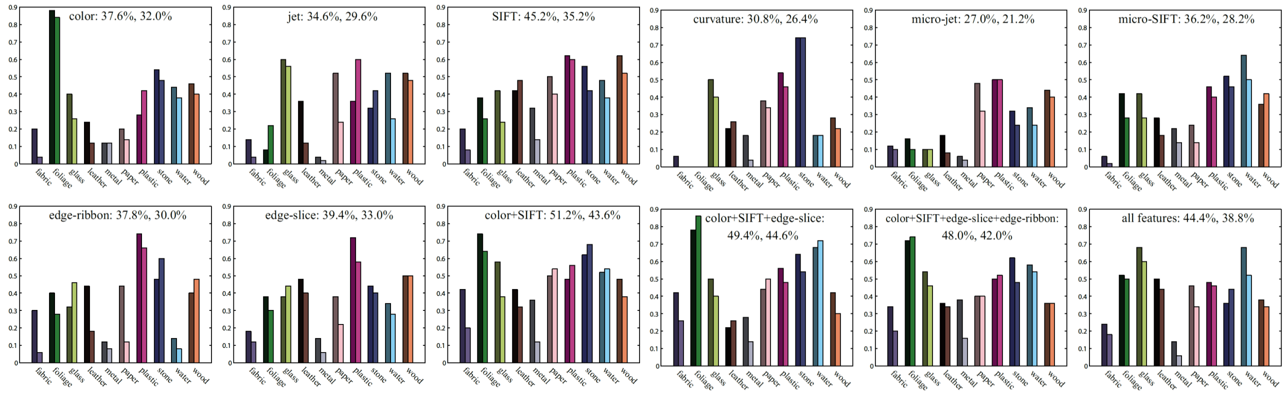

Figure 6: The per-class recognition rate (both training and test) with different sets of features for the MIT Flickr Material Database [28], using the classification approach proposed in [20]. In each plot, the left, darker bar means training, the right and the lighter bar means test. The two numbers right after the feature set label denote the recognition rate on the entire training set and on the entire testing set.

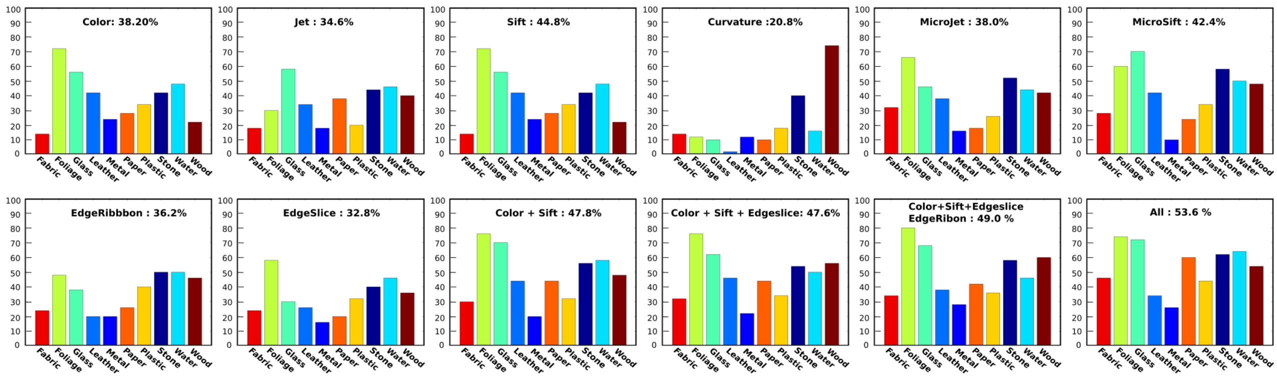

Figure 7: The recognition rate of a randomly chosen split of training-testing shown for each single feature and some of the feature combinations. Only the test rate per material category is mentioned, since the training rate with learned parameters is very high and hence not of our interest. The highest recognition rate (53.6%) is achieved when all the features are used.

We extract the features color, SIFT, Jet, micro-SIFT and micro-Jet on every 5-th pixel inside the given mask where the material is present. Edge-Slice, Edge-Ribbon and curvature are calculated on every second pixel on the edges in the edge map. After calculating the features we perform k-means clustering separately for each feature. We use exactly the same number of clusters for each feature type (150 for color, 250 for SIFT, 200 for Jet, 250 for micro-SIFT, 200 for micro-Jet, 100 for curvature, 200 for Edge-Slice, 200 for Edge-Ribbon).

After forming the dictionary for each feature we learn visual word histograms based on the training images using SVMs. First we train and test with single features and then we combine the features as in [20]. We observe that our best performing single averaged feature (SIFT = 42.2%) has a recognition rate which is very close to the best performance (44.6%) stated in [20]. Among all features and feature combinations, in every trail we achieve the best rates (53.2%, 53.6%, 53.8%, 53.6%, 51.6%) when all the features are combined. A comparison of average recognition rates of individual features and feature combinations is tabulated in Table 1. It can be seen that for most of the features there is an improvement in recognition accuracy of about 4 to 5 %. For the individual features, micro-SIFT and micro-Jet and the combination of all the features there is a significant improvement in the classification accuracy. The confusion matrix in Figure 9 shows the accuracy of the classification of individual material categories. The rates of misclassification are also shown.

6 Conclusions

In this paper, we addressed the problem of material recognition using various image features in combination with a SVM framework and compared it to the Bayesian approach proposed in [20]. The recognition rate achieved by our system is 53.1% in average on the MIT Flickr Material Database and 99.4% on the KTH_TIPS2 database. The reason for the huge difference in recognition rate between the two datasets is due to the larger intra-class variations

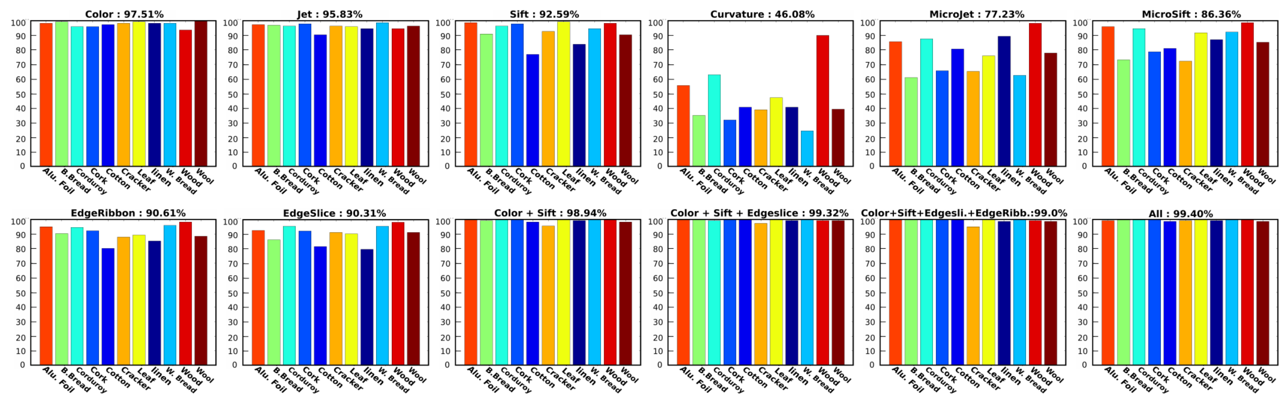

Figure 8: The recognition rate of a randomly chosen split of training-testing for KTH_TIPS2 database is shown for each single feature and some of the feature combinations. Only the test rate per material category is mentioned, since the training rate with learned parameters is nearly 100% for this dataset. The highest recognition rate (99.4%) is achieved when all the features are used.

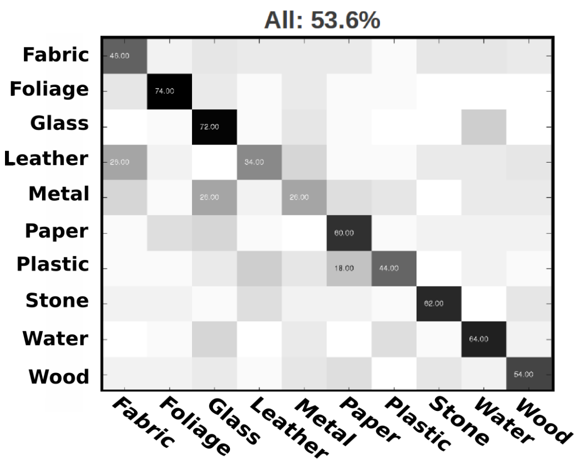

Figure 9: Confusion matrix for the Flickr Material Database. Diagonal entries shows the percentage with which each category is recognized. Rates are color coded using gray scale values (black = 100% and white = 0%). Each row sums to 100%.

Feature	Ce Liu et.al	Our model (avg. of 5 iterations)
Color	32.6 %	37.6%
Jet	29.6 %	34.0%
SIFT	35.2 %	42.2%
Curvature	26.4 %	21.6%
Micro-Jet	21.2 %	36.5%
Micro-SIFT	28.2 %	42.0%
Edge-Ribbon	30.0 %	36%
Edge-Slice	33.0 %	34.6%
Color+SIFT	43.6 %	48.6%
Color+SIFT+Edge-Slice	**44.6** %	49.1%
Color+SIFT+Edge-Slice +Edge-Ribbon	42.0 %	49%
All	38.8 %	**53.1**%

Table 1: Performance comparison between [20] and our system.

in the Flickr Material Database. In contrast, KTH_TIPS2 comprises images of the same material taken in different view and lighting conditions in the same material category. We showed that with the same pool of features as given in [20], SVM classifies materials with a higher rate in comparison to the Bayesian approach of [20]. We have also analyzed the contribution of each feature in our system to the performance gain. Future developments should consider exploring different and better characteristic features for materials. Improvements by integration of different classification techniques can also be investigated.

References

[1] S. Bae, S. Paris, and F. Durand. Two-scale tone management for photographic look. *ACM Trans. Graph.*, 25:637–645, July 2006.

[2] G. Bejerano. *Automata learning and stochastic modeling for biosequence analysis*. 2003.

[3] C.J.C. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2:121–167, 1998.

[4] B. Caputo, E. Hayman, M. Fritz, and J. Eklundh. Classifying materials in the real world. *Image Vision Comput.*, 28:150–163, January 2010.

[5] C. Chang and C. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011.

[6] O. Chapelle, P. Haffner, and V. Vapnik. Support vector machines for histogram-based image classification. *IEEE Transactions on Neural Networks*, 10(5):1055–1064, 1999.

[7] C. Cortes and V. Vapnik. Support vector networks. *Machine Learning*, 20:273 – 297, 1995.

[8] O. G. Cula and K. J. Dana. Compact Representation of Bidirectional Texture Functions. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1041–1047, 2001.

[9] N. Dalal and B. Triggs. Object detection using histograms of oriented gradients. *European Conference on Computer Vision Workshop on Pascal VOC*, 06, 2006.

[10] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.*, 18:1–34, January 1999.

[11] P. Debevec, T. Hawkins, C. Tchou, H.P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '00, pages 145–156, 2000.

[12] S. Dumais, J. Platt, D. Heckerman, and M. Sahami. Inductive learning algorithms and representations for text categorization. In *Proceedings of the seventh international conference on Information and knowledge management*, CIKM '98, pages 148–155, 1998.

[13] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '02, pages 257–266, New York, NY, USA, 2002. ACM.

[14] Canny J. A computational approach to edge detection. *IEEE Trans. Pattern Analysis Machine Intelligence*, 8(6):679–698, 1986.

[15] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan. A practical model for subsurface light transport. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '01, pages 511–518, 2001.

[16] T. Joachims. Text categorization with support vector machines : Learning with many relevant features. *Machine Learning ECML98*, 1398(23):2–7, 1998.

[17] J. J. Koenderink and A. van Doorn. Representation of local geometry in the visual system. *Biological Cybernetics*, 55:367–375, 1987.

[18] Ulrich H.-G. Kreßel. *Pairwise classification and support vector machines*, pages 255–268. MIT Press, Cambridge, MA, USA, 1999.

[19] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *Int. J. Comput. Vision*, 43:29–44, June 2001.

[20] C. Liu, L. Sharan, E. H. Adelson, and R. Rosenholtz. In *CVPR*.

[21] F. Liu and R.W. Picard. Periodicity, directionality, and randomness: Wold features for image modeling and retrieval. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 18:722–733, 1996.

[22] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60:91–110, 2004.

[23] F. E. Nicodemus. Directional reflectance and emissivity of an opaque surface. *Applied Optics*, 4(7):767–775, 1965.

[24] S. C. Pont and J. J. Koenderink. Bidirectional texture contrast function. *Int. J. Comput. Vision*, 62:17–34, April 2005.

[25] M. Pontil and A. Verri. Properties of Support Vector Machines. Technical Report AIM-1612, 1997.

[26] D. Roobaert, M. Zillich, and J. Eklundh. A Pure Learning Approach to Background-Invariant Object Recognition Using Pedagogical Support Vector Learning. In *Proceedings of the 2001 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Kauai, Hawaii*, pages 351–357, 2001.

[27] L. Sharan, R. Rosenholtz, and E. H. Adelson. Material perception: What can you see in a brief glance? [abstract], 2009.

[28] V. Vapnik. *The nature of statistical learning theory*. Springer, New York, 1995.

[29] M. Varma and A. Zisserman. A statistical approach to material classification using image patch exemplars. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31:2032–2047, November 2009.

[30] J. Wang. Application of Support Vector Machines in Bioinformatics. Master's thesis, 2002.