Learning Discriminative Illumination and Filters for Raw Material Classification with Optimal Projections of Bidirectional Texture Functions

Chao Liu, Gefei Yang, Jinwei Gu
Rochester Institute of Technology

Abstract

We present a computational imaging method for raw material classification using features of Bidirectional Texture Functions (BTF). Texture is an intrinsic feature for many materials, such as wood, fabric, and granite. At appropriate scales, even "uniform" materials will also exhibit texture features that can be helpful for recognition, such as paper, metal, and ceramic. To cope with the high-dimensionality of BTFs, in this paper, we proposed to learn discriminative illumination patterns and texture filters, with which we can directly measure optimal projections of BTFs for classification. We also studied the effects of texture rotation and scale variation for material classification. We built an LED-based multispectral dome, with which we have acquired a BTF database of a variety of materials and demonstrated the effectiveness of the proposed approach for material classification.

1. Introduction

Texture is an important feature for recognizing almost all real-world materials, some of which are intrinsically textures (*e.g.*, fabric, wood, granite) and others exhibits texture at appropriate scales (*e.g.*, metal, paper, plastic, ceramic). Considering lighting and view dependence, a complete representation of texture is the Bidirectional Texture Function (BTF) [6], which is a 6-D function $f(\vec{x}, \vec{\omega}_i, \vec{\omega}_o)$. Subsets of this function have been utilized previously as "3D Textons" for texture classification [5, 17]. The 3D texton approach has shown higher performance for classification compared to conventional texture classification [31, 32]. However, it requires a time-consuming capture of images under multiple lighting/view conditions, which significantly limits its practicality. Can we directly measure discriminative information from BTFs with one or two images?

A second question is texture representation. Most prior work uses the responses to a set of filters (*e.g.*, S-filters, Gabor, steerable, Laws' masks) as the representation [14, 12, 8, 28, 2, 30, 31]. What are "good" filters for classification? We realize the answer is task-dependent. Finding the opti-

mal filters is expected to reduce the dimensionality of the classification task and thus increase performance.

Motivated by these two questions, in this paper, we present an efficient method to directly capture discriminative features of BTFs for material classification with computational lighting and filters optimized for given classification tasks. (1) To reduce the acquisition cost, we observe that each measured image is a linear *projection* of the BTFs to the illumination pattern. Thus, by appropriately setting the illumination pattern, we can directly capture discriminative projections of BTFs without actually measuring the full BTFs. (2) To find good filters for texture classification, we model the filters as linear combinations from a general filter dictionary and optimize the weights by minimizing a discriminant functional. Both the optimal illumination pattern — called "discriminative illumination"— and the weights for filters are optimized iteratively to maximize the distance between different materials.

Our method is closely related to two recent works [9, 13] which also seek optimal illumination for material classification. However, they focused on point-wise material classification using BRDF features only. In this work, we incorporated non-local features — texture — for classification. Experiments showed that our method has better results than theirs. Nevertheless, similar to these prior works, our method also has signal-to-noise benefit due to illumination multiplexing and can also be extended for multi-class and nonlinear classification by using multiple discriminative illumination patterns.

Since BTFs are generally neither rotational-invariant nor scale-invariant, we also studied the effects of rotation and scale to texture classification. We found indeed that both factors degrade the performance of texture classification. One solution we proposed is to incorporate copies of BTFs with multiple orientations and scales in the training stage. Experimental results showed that the optimal illumination and filters learned from this "augmented" training set is much more stable to rotation and scale change for texture classification.

To evaluate our method, we built a LED-based multispectral dome similar to that in [9]. We used this

Figure 1: Samples in the spectral BTF database for material classification. The database includes eight classes of materials (*i.e.*, aluminum, stainless, granite, carpet, paper, plastic and wood) and 90 samples in total. Each sample is imaged under $6 \times 25 = 150$ lighting conditions (6 in wavelength and 25 in angular) with a resolution of 1040×1392. The database is available at http://compimg1.cis.rit.edu/data/texture/.

rig and measured a spectral BTF database covering a range of materials, including metal, plastic, paper, carpet, wood, granite, *etc.*. Some samples are shown in Fig. 1. The database is available at http://compimg1.cis. rit.edu/data/texture/. We evaluated the proposed method for a variety of materials classification tasks. We also compared with prior work on texture classification [31] and BRDF classification [9]. Experimental results showed the effectiveness of the proposed method.

2. Related Work

Material Classification There are two categories of prior work on material classification in computer vision. The passive approach aims to study material perception in order to classify materials from regular images via statistical learning [1, 21, 18]. The active approach, especially in machine vision, employs any useful visual features for classification, such as 2D slices of BRDFs [34, 13], BRDF projections [9], polarization [3], and spectral reflectance [10, 26]. To our knowledge, there has been no prior work using BTF for material classification.

Texture Analysis and Classification Texture analysis and classification has been a long standing research area in computer vision [14, 12, 8, 28, 2, 30, 31]. To deal with the change of orientation and scale, researchers either estimated the dominant orientation/scale of a texture image [16, 4, 31], or utilized rotational-invariant or scale-invariant features [24, 15, 20, 25]. Another body of prior work focused on optimizing filters for texture classification [19, 11, 29]. Unlike these prior works that deal with single texture images, we focus on BTF classification.

Computational Illumination Our work is in the area of computational illumination which uses coded light for efficient material and shape measurement and is also related to task-specific and feature-specific imaging [23, 22], in which the goal of such imaging systems is to maximize the amount of information relevant to given tasks (in our case, material

classification). An essential component in our work is the supervised learning from labeled data sets.

3. Learning Discriminative Filters and Illumination for BTF Classification

Let us define the spectral BTF as a 7D reflectance field $f(p, \omega_i, \omega_o, \lambda)$ with incident light in the direction of ω_i, and reflected light in the direction in ω_o at wavelength λ for a pixel at point p. Consider a sample illuminated by multiple light sources from different directions with different spectra, as shown in Fig.2 (b). At a fixed viewpoint ω_o, the captured image $I(p)$ is given by

$$I(p) = \int_{\lambda, \omega_i} f(p, \omega_i, \omega_o, \lambda) L(\omega_i, \lambda) V(\omega_i) S(\lambda) d\lambda d\omega_i, \tag{1}$$

where $L(\omega_i, \lambda)$ is the incident light in the direction ω_i at wavelength λ, $V(\omega_i) = \max(0, \cos \theta_i)$ is the visibility term, and $S(\lambda)$ is the camera spectral response function.

We represent a textured surface as a vector of the responses with various image filters. Since the appearance of the texture is linear with respect to the incident illumination, the vector of filter responses will also be linear with respect to the incident illumination. Therefore, we can employ a similar approach as in [9] to learn discriminative illumination as optimal projection directions for texture classification. Specifically, let $\mathcal{N}$ denote a local neighborhood (*e.g.*, a 5×5 window). The response of the textured surface with a filter under the incident illumination is:

$$r = \sum_{p \in \mathcal{N}} \alpha_p I(p) = \sum_{p \in \mathcal{N}} \sum_{j=1}^{M} \alpha_p \cdot \hat{f}(p, j) \cdot L_j, \tag{2}$$

where α_p is the coefficient of the filter at pixel p, and $I(p)$ is the pixel intensity of the textured surface under the incident illumination. We note $I(p)$ is a dot product between the *apparent* BTF feature vector of pixel p $[\hat{f}(p, j) = f(p, \omega_i, \omega_o, \lambda) V(\omega_i) S(\lambda)]$ and the incident illumination $[L_j]$, $j = 1, \cdots, M$. For a set of K filters, the

texture descriptor is a $K \times 1$ vector:

$$\mathbf{r} = \left[r^{(1)}, \cdots, r^{(K)}\right]^T = \left(\mathbf{A}^T \cdot \mathbf{F}^T\right) \cdot \mathbf{w}, \qquad (3)$$

where $\mathbf{w} = [L_1, \cdots, L_M]^T$ is a vector of the illumination, $\mathbf{A}$ is an $N \times K$ matrix consisting of the coefficients of the K filters (*i.e.*, $\mathbf{A}_{ij} = \alpha_i^{(j)}$ where $\alpha_i^{(j)}$ is the coefficient of the j-th filter at pixel i, and $\mathbf{F}$ is an $M \times N$ matrix consisting of the apparent BTF feature vectors (*i.e.*, $\mathbf{F}_{ij} = \hat{f}(j, i)$). The illumination vector $\mathbf{w}$ projects the original texture feature $\mathbf{A}^T \cdot \mathbf{F}^T$ of dimension $K \times M$ to a feature vector $\mathbf{r}$ in a lower dimensional subspace of dimension $K \times 1$.

What are the optimal filters to use in $\mathbf{A}$? Instead of using a set of predefined filters, we solve for optimal filters (together with solving for optimal illumination) by maximizing some discriminant functional (*e.g.*, Fisher LDA or SVM with a linear kernel). We model each filter in $\mathbf{A}$ as a linear combination of basis filters from a general dictionary $\mathbf{B}$ consisting of Gabor filters, S-filters, and maximal response filters. We have $\mathbf{A} = \mathbf{BW}$, with vectorized basis filters in columns of $\mathbf{B}$ and weights of basis filters in columns of the weights matrix $\mathbf{W}$. Thus we have

$$\mathbf{r} = \mathbf{W}^T \mathbf{B}^T \mathbf{F}^T \mathbf{w} = \mathbf{W}^T \mathbf{R}^T \mathbf{w}, \qquad (4)$$

where $\mathbf{R} = \mathbf{FB}$ is the filter response matrix for the basis filters, with $\mathbf{R}_{ij}$ being the filter response of the i-th basis filter under the j-th light source.

With a given training set, $\mathbf{R}$ is known. The optimal illumination $\mathbf{w}$ and filters $\mathbf{W}$ for the classification task are learned by maximizing the between-class distance while minimizing the within-class distance (with a unit norm constraint on $\mathbf{w}$ to avoid scale ambiguity), similar to LDA.

$$\max_{\mathbf{W}, \mathbf{w}} J(\mathbf{S}_b, \mathbf{S}_w), \quad \text{st.} \quad ||\mathbf{w}|| = 1,$$

where $\mathbf{S}_b$ and $\mathbf{S}_w$ are the between-class scatter matrix and the within-class scatter matrix for $\mathbf{r}$, respectively,

$$\mathbf{S}_b = \sum_{c=1}^{C} (\bar{\mathbf{r}}_c - \bar{\mathbf{r}})(\bar{\mathbf{r}}_c - \bar{\mathbf{r}})^T$$

$$\mathbf{S}_w = \sum_{c=1}^{C} \sum_{i=1}^{N_c} (\mathbf{r}_{i,c} - \bar{\mathbf{r}}_c)(\mathbf{r}_{i,c} - \bar{\mathbf{r}}_c)^T,$$

where $\mathbf{r}_{i,c}$ is the texture descriptor for the i-th sample in the c-th class, $\bar{\mathbf{r}}_c$ is the average descriptor for class c, and $\bar{\mathbf{r}}$ is the average descriptor for all samples in the C classes. As shown in [7], there are multiple choices of $J(\cdot)$. In this paper we chose $J(\cdot)$ as below

$$\max_{\mathbf{W}, \mathbf{w}} J = \frac{\mathbf{Trace}(\mathbf{S}_b)}{\mathbf{Trace}(\mathbf{S}_w)}, \quad \text{st.} \quad ||\mathbf{w}|| = 1. \qquad (5)$$
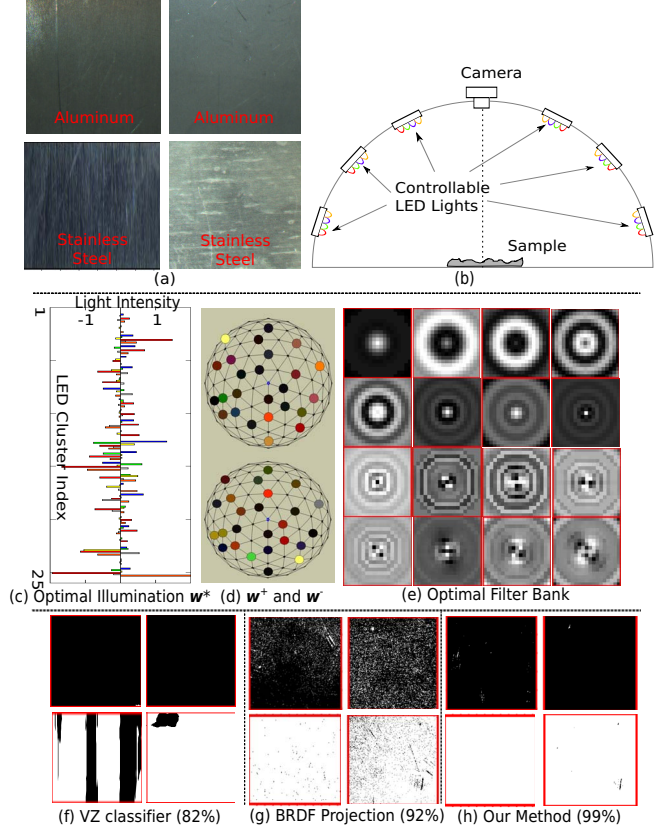

Figure 2: Texture classification with discriminative illumination and filters. (a) Classifying aluminum and stainless steel under conventional lighting with regular color camera is challenging, since they have similar color and gloss. (b) We proposed to capture projections of BTFs for material classification with coded illumination, implemented as a LED-based multispectral dome. (c) and (d) show the optimal illumination. The bar graph shows the learned $\mathbf{w}$, where the 25 bar groups correspond to the 25 LED clusters and the six bars within each group correspond to the six LEDs. This coded light pattern is also shown as $\mathbf{w}^+$ and $\mathbf{w}^-$ where $\mathbf{w} = \mathbf{w}^+ - \mathbf{w}^-$. (e) The optimal filters. (f)(g)(h) show the classification rates on test data using the VZ texture classifier [31], BRDF Projection [9] and our method with the same number of measurements.

We optimize $\mathbf{w}$ and $\mathbf{W}$ in Equation. (5) alternatively. When $\mathbf{W}$ is fixed, finding $\mathbf{w}$ is the well-known LDA problem; when $\mathbf{w}$ is fixed, optimizing $\mathbf{W}$ becomes a trace-ratio problem for dimensionality reduction [33]. Both have algorithms for global optimal solutions. Please refer to the supplementary material for details.

In the classification stage, a test sample is captured when the illumination pattern is set to $\mathbf{w}^*$. Since $\mathbf{w}^*$ often contains negative values, it is implemented as the difference of two nonnegative patterns $\mathbf{w} = \mathbf{w}^+ - \mathbf{w}^-$ where $\mathbf{w}^+ =$

$\max(0, \mathbf{w})$ and $\mathbf{w}^- = -\min(0, \mathbf{w})$. For each point, we use its neighborhood to compute the texture descriptor with the learned optimal filters $\mathbf{A}^* = \mathbf{BW}^*$ for classification.

Figure 2 shows a challenging example for classifying aluminum and stainless steel, which has similar color and gloss. As shown, compared to two state-of-art material classification methods (the VZ classifier [31] and the BRDF projections [9]), our method has better performance on the testing data, with the same number of measurements.

3.1. The Effect of Orientation for BTF Classification

It is well known that the appearance of a BTF is highly dependent on sample orientation and scale, especially for materials with rough surfaces (*e.g.*, carpet, tree bark, fur). Because of self-occlusion and inter-reflection, we often cannot approximate the texture appearance at a different orientation or scale by simply rotating or scaling a captured texture image. How does the change of orientation and scale affect the classification performance? We discuss these effects in the next two subsections.

We first focus on the effect of orientation. To eliminate the contribution of BRDF to classification and study the effect of orientation only to texture, as shown in Fig. 3(a), we prepared two classes of materials which are coated with the same paint — one is paper and the other is grooved clay tile. Thus the material difference will be only due to texture. First, we confirmed that orientation is indeed important for classification. As shown in the top row of Fig. 3(c), if we apply the classifier trained with one orientation to images captured at an unknown orientation, the classification rate is less than 50% for both BRDF and BTF projections.

In order to achieve rotational invariance, we tried two methods. First, we added rotated samples into the training set. Second, we limited the basis filters in the dictionary to be rotational-invariant (*e.g.*, S-filter instead of Gabor filters). We found the latter does not change the results significantly, but adding rotated samples in the training set can effectively increase the robustness to orientation. As shown in Fig. 3(c), the more rotated sampled added, the better the performance. With eight orientations in training, the classification rate is 99.1%. Figure 3(b) shows the trained $K = 16$ filters when we use one orientation, four orientations, and 8 orientations in training, respectively. As shown, the more orientations added in training, the more directional filters are learned. In comparison, adding rotated samples does not change much the classification results of the BRDF projection method [9], as expected.

3.2. The Effect of Scale for BTF Classification

Scale is another important factor in texture classification. Depending on the distance between the sample and the sensor, different geometric and photometric properties will be observed by each pixel and thus result in different appear-
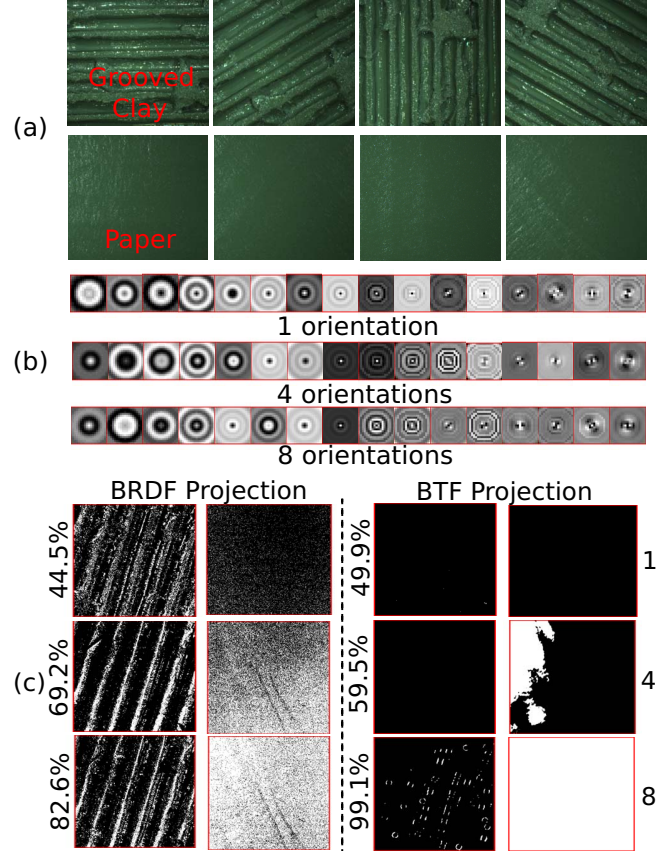

Figure 3: The effect of orientation for BTF classification. (a) We prepared two materials coated with the same blue paint for material classification with texture only. The samples measured at different orientations show the changes in self-shadow and specular lobe caused by surface geometry. (b) By using multiple rotated samples in training, we learned classifiers (*i.e.*, illumination and filters) that are more robust to orientation. (c) As expected, the accuracy of our BTF projection method increases with the number of rotated samples added to the training set, while the BRDF projection method [9] does not vary significantly.

ance. To our knowledge, there is few prior work explicitly models the relationship between scale and appearance.

We performed some preliminary study on this problem. As shown in Fig. 4, we aimed to classify carpet and paper at two scales. Two sets of images are captured for the same samples, one with a lens with focus length 12mm and the other 50mm. As shown, at the two different scales, some details, such as the shadows among knots of the carpet, are lost. As a result, as shown in Fig. 4(g), the BTF-based classifier trained at the lower scale (f=12mm) fails to classify the carpet at the higher scale (f=50mm).

To deal with the scale change, we adapt a similar approach as that in the orientation scenario. We include the

(a) f = 12 mm

(b) f = 50 mm

(c) $\mathbf{w}^+$ and $\bar{\mathbf{w}}$ for One Scale (f=12mm)

(d) $\mathbf{w}^+$ and $\mathbf{w}^-$ for Both Scales

(e) Filters for One Scale (f=12mm)

(f) Filters for Both Scales

(g) Results for Classifiers Trained for One Scale (f=12mm) (77.3%)

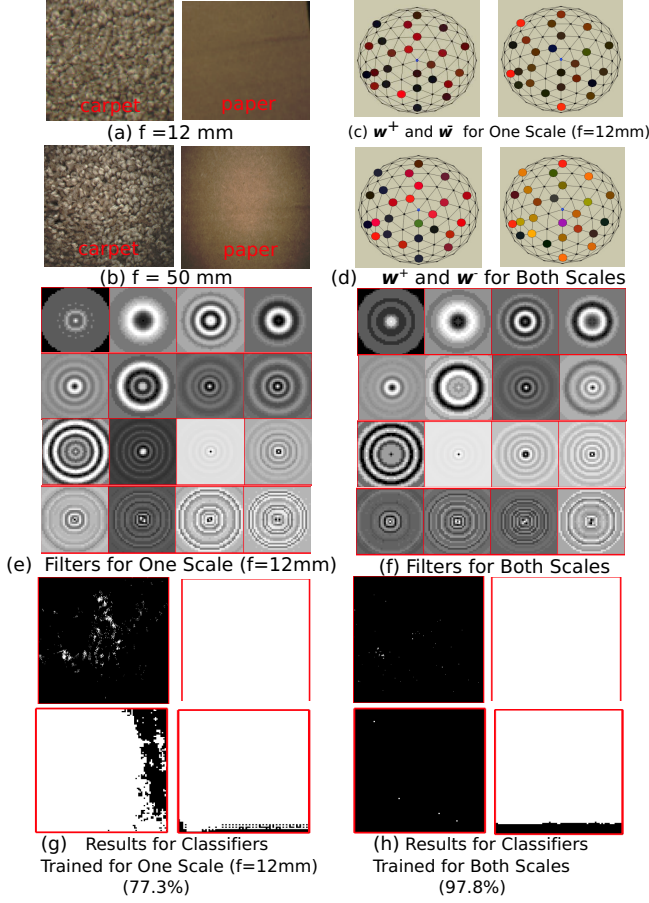(h) Results for Classifiers Trained for Both Scales (97.8%)

Figure 4: The effect of scale for BTF classification. (a)(b) show the images of carpet and paper captured at two different scales. (c)(e): the optimal illumination ($\mathbf{w}^+$, $\mathbf{w}^-$) and filters trained with samples in one scale. (d)(f): the optimal illumination ($\mathbf{w}^+$, $\mathbf{w}^-$) and filters trained with samples in both scales. The differences in the trained illumination and filters confirm that BTF is not scale-invariant. (g): classification results when only samples in one scale are included in the training set. (h) classification results when samples in both scales are included in training set. The classification rate increases as the training sets include both scales.

samples in both scales into the training set. The classification rate increases from 77.3% to 97.8% after we expand the training set to include samples in both scales, as shown in Fig. 4(h). Figures 4(c)(d) and (e)(f) show the corresponding discriminative illumination and filters.

4. Experimental Results

To evaluate our method, we built a LED-based multi-spectral dome similar to the device in [9]. As shown in Fig. 5, the dome consists of 25 clusters of LEDs. Each cluster has six LEDs of different colors: blue, green, am-


Figure 5: The LED-based multi-spectral dome for discriminative illumination and filters for BTF classification.

ber, white, red and orange. The six colors are selected so that the visible spectrum is covered. Each LED cluster is controlled by an Arduino board which controls the intensities of the six LEDs independently. A Lumenera camera is mounted on the top of the dome.

We acquired a spectral BTF database using this dome, covering eight classes of materials (*i.e.*, aluminum, stainless, granite, carpet, paper, plastic and wood) and 90 samples in total. Each sample is imaged under $6 \times 25 = 150$ lighting conditions (6 in wavelength and 25 in angular) with a resolution of 1040×1392. The database, along with source code, is available at

We first run simulation experiments to confirm that both spectral and angular coding are needed and why the six colors and the 25 directions are sufficient. As summarized in Fig. 6, we perform experiments for three classification tasks: aluminum vs. stainless steel, carpet vs. paper, and granite vs. wood. For each sample, we measure $25 \times 6 = 150$ images corresponding to the 150 LEDs in the dome. The left column shows the classification rates with color features only, *i.e.*, we turn on the LEDs of the same color in all the clusters. The middle column shows the classification rates with angular distribution of reflectance for classification, *i.e.*, we use only white LEDs and disable all other color LEDs. We found in both cases, although the performance increases with the number of colors/angular samples used, the classification rate is not as high as that if we use both spectral and angular distribution of reflectance as shown in the right column. These plots also show that six colors and 25 LED clusters are sufficient for these classification tasks.

The general filter dictionary has 541 basis filters, including 13 S-filters [27], 16 circular ring filters, and 512 Gabor filters of different scales and orientations, The circular ring filters are the filters with unit intensity at the pixels with a

Figure 6: Discriminative ability of spectral BTF for material classification. We evaluate the performance of the learned discriminative illumination with different amounts of color, LED clusters and LEDs. The plots show that the spectral and angular distribution of reflectance are complimentary to each other in material classification. The results also demonstrate that six colors and 25 LED clusters are sufficient for these tasks.

particular distance from the center and zeros at other locations. We evaluated the classification performance of the proposed method with respect to the filter size and the number of filters to learn. Details are given in the supplementary material. In all the classification tasks, we set the filter size to be 19×19. We train $K = 16$ filters that are linear combination of the basis filters for each classification task. The iterative algorithm to learn optimal discriminative illumination and filters is usually converged in 3 to 5 iterations.

We compare with several recent related work for texture/material classification in three aspects. First, how does our method compare to methods that use the same number of images as ours? We test the VZ classifier [31] and the BRDF projection method [9]. The VZ classifier is a "bag-of-words" method for regular texture classification with a single image as input. We implemented it for every 49×49 patches to classify the material of the central pixel. The BRDF projection uses discriminative projections of BRDF slices for material classification, which also uses a single coded image as input (implemented as the subtraction of two images since light cannot be negative). As shown in Table 1, with the same number of required images, our method outperforms both methods.

Second, is the joint-learning of the illumination and filters necessary? To answer this question, we compare with a method in which the discriminative illumination and the filters are learned separately. The illumination is first learnt in the same way as in [9]. Then the filters are optimized based on the illumination. Note that this method is a fairly strong competitor, because it is in fact the first iteration of our method with the initial filters being the 2D delta functions[1]. Nevertheless, as shown in Table 1, the results show that the co-learning still improves the performance consistently. Thus we conclude it is necessary to co-learn the illumination and filter patterns for classification.

Third, what is the limit of our method? We hope to understand its performance relative to methods that do not impose any constraints on acquisition cost. One well-known method in this category is the 3D texton [17], in which multiple images under different illuminations are used for texture classification. Filter responses of the stack of images at different orientations and scales are learned with the bag-of-words approach and used as features for texture classification. As shown in Table 1, the 3D texton needs 150 input images and achieves the highest classification rates among all methods for some tasks. This is expected since the feature space of the 3D texton is of much higher dimensionality than those of the other methods. We found our method has quite close performance compared to 3D texton by using only a single coded image. For some tasks, our method performs even better.

Moreover, since the VZ classifier and 3D texton are both "bag-of-words" methods, we also evaluate them with different numbers of words. We found their performance is sensitive to this parameter — which makes them less robust compared to our method. The details and the visualizations of the classification results of the comparisons are included in the supplementary material.

Finally, we can generalize our method for multi-class and non-linear classification similar to that shown in [9]. Table 2 shows the classification rates for some multi-class classification tasks, for which we use the one-vs-all scheme to generalize the binary classifiers. For each task, we compare our method with the BRDF projection method [9] and the VZ classifier [31]. As shown, our method is more accurate for all the tasks.

5. Conclusion and Discussions

In this paper, we presented a method that uses coded illumination to directly measure projections of BTFs for material classification. Optimal illumination patterns and filters are learned via a training stage. We also proposed ways to deal with the change of orientation and scale for material classification. Compared to relevant prior work, the proposed method has higher performance in classification and is economical in measurement cost.

---

[1]Our method converges within fewer than five iterations in most cases

Table 1: Comparison results with the VZ classifier [31], 3D Texton [17], the BRDF projection method [9] and the BRDF projection coupled with optimal filters. The number of required images for each method is shown in the bracket. 1^* means it uses the subtraction of two images since light cannot be negative.

Task	VZ [31] (1)	3D Texton[17] (150)	BRDF Projection [9] (1^*)	BRDF projection + Optimal filters (1^*)	Our method (1^*)
Aluminum vs. Granite	79.12%	**100**%	94.03%	93.37%	98.55%
Aluminum vs. Stainless	84.81%	**99.40**%	86.61%	89.97%	96.12%
Aluminum vs. Wood	81.50%	81.75%	99.56%	99.99%	**100**%
Carpet vs. Wood	80.59%	86.65%	96.09%	95.04%	**99.04**%
Carpet vs. Paper	85.28%	91.48%	93.61%	99.98%	**100**%
Granite vs. Paper	86.83%	**100**%	97.32%	99.99%	99.74%
Paper vs. Wood	73.84%	89.81%	99.20%	100%	**100**%
Plastic vs. Stainless	78.55%	86.97%	99.77%	100%	**100**%
Paper vs. Aluminum	85.63%	99.25%	99.01%	99.08%	**100**%

The proposed method has a few limitations. First, the proposed method works primarily on flat samples (with textures). It still does not decouple 3D shape and material from visual appearance. Second, for complex, nonlinear classification tasks, we may still need to acquire many images under multiple discriminative illumination. Finally, the method we proposed to deal with rotation and scale requires more images in training — a better solution is to explicitly model the relationship between scale/rotation and visual appearance. These are among our future research.

6. Acknowledgement

References

[1] E. H. Adelson. On seeing stuff: The perception of materials by humans and machines. In *Human Vision and Electronic Imaging VI (SPIE)*, volume 4299, pages 1–12, 2001. 2

[2] M. J. Chantler. *The Effect of Variation in Illuminant Direction on Texture Classification*. PhD thesis, Heriot-Watt University, 1994. 1, 2

[3] H. Chen and L. B. Wolff. A polarization phase-based method for material classification in computer vision. *International Journal of Computer Vision (IJCV)*, 28(1):73–83, 1998. 2

[4] F. Cohen, Z. Fan, and M. Patel. Classification of rotated and scaled textured images using gaussian markov random field models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1991. 2

[5] O. G. Cula and K. J. Dana. 3D texture recognition using bidirectional feature histograms. *International Journal of Computer Vision (IJCV)*, 59(1):33–60, August 2005. 1

[6] K. J. Dana, B. V. 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. 1

[7] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press Professional, Inc., San Diego, CA, USA, 2nd edition, 1990. 3

[8] S. E. Grigorescu, N. Petkov, and P. Kruizinga. Comparison of texture features based on gabor filters. *IEEE Transactions on Image Processing*, 11(10):1160–1167, October 2002. 1, 2

[9] J. Gu and C. Liu. Discrminative illumination: Per-pixel classification of raw materials based optimal projections of spectral brdfs. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2012. 1, 2, 3, 4, 5, 6, 7, 8

[10] A. Ibrahim, S. Tominaga, and T. Horiuchi. Spectral imaging method for material classification and inspection of printed circuit boards. *Optical Engineering*, 49:057201, 2010. 2

[11] A. Jain and K. Karu. Learning texture discrimination masks. In *IEEE International Conference on Neural Networks*, 2004. 2

[12] A. K. Jain and F. Farrokhnia. Unsupervised texture segmentation using gabor filters. *Pattern Recognition*, 24(12):1167–1186, 1991. 1, 2

Table 2: Comparison results with the VZ [31] and BRDF Projection method [9] for multi-class classification tasks.

Task	VZ Classifier [31]	BRDF Projection [9]	Our method
Aluminum vs. Granite vs. Plastic	79.25%	76.84%	89.13%
Aluminum vs. Granite vs. Stainless	73.15%	93.23%	97.13%
Aluminum vs. Plastic vs. Stainless	75.09%	92.99%	96.71%
Aluminum vs. Granite vs. Plastic vs. Stainless	73.67%	78.44%	90.49%
Carpet vs. Granite vs. Plastic vs. Stainless	65.98%	64.58%	74.11%
Aluminum vs. Granite vs. Stainless vs. Wood	63.28%	93.75%	97.66%
Paper vs. Plastic vs. Stainless vs. Wood	68.40%	93.12%	96.71%
Granite vs. Paper vs. Plastic vs. Stainless	75.43%	84.44%	91.57%
Aluminum vs. Carpet vs. Granite vs. Paper vs. Plastic	65.54%	74.26%	83.83%

[13] M. Jehle, C. Sommer, and B. Jähne. Learning of optimal illumination for material classification. *Pattern Recognition*, 6376:563–572, 2010. 1, 2

[14] K. I. Laws. *Textured Image Segmentation*. PhD thesis, University of Southern California, 1980. 1, 2

[15] S. Lazebnik, C. Schmid, and J. Ponce. A sparse texture representation using local affine regions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2005. 2

[16] M. Leung and A. Peterson. Scale and rotation invariant texture classification. In *Conference Record of The Twenty-Sixth Asilomar Conference on Systems and Computers*, 1992. 2

[17] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision (IJCV)*, 43(1):29–44, June 2001. 1, 6, 7

[18] C. Liu, L. Sharan, E. H. Adelson, and R. Rosenholtz. Exploring features in a bayesian framework for material recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. 2

[19] A. Mahalanobis and H. Singh. Application of correlation filters for texture recognition. *Applied Optics*, 33(11):2173–2179, Apr 1994. 2

[20] G. McGunnigle and M. J. Chantler. Rough surface classification using point statistics from photometric stereo. *Pattern Recognition Letters*, 21(6-7):593–604, 2000. 2

[21] I. Motoyoshi, S. Nishida, L. Sharan, and E. H. Adelson. Image statistics and the perception of surface qualities. *Nature*, 447:206–209, May 2007. 2

[22] M. A. Neifeld, A. Ashok, and P. K. Baheti. Task-specific information for imaging system analysis. *Journal of Optical Society of America*, 24(12):B25–B41, 2007. 2

[23] M. A. Neifeld and P. Shankar. Feature-specific imaging. *Applied Optics*, 42(17):3379–3389, 2003. 2

[24] T. Ojala, M. Pietikainen, and T. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2002. 2

[25] G. Oxholm, P. Bariya, and K. Nishino. The scale of geometric texture. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2012. 2

[26] N. Salamati, C. Fredembach, and S. Susstrunk. Material classification using color and NIR images. In *Proceedings of IS&T (SID) Color Imaging Conference (CIC)*, 2009. 2

[27] C. Schmid. Constructing models for content-based image retrieval. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001. 5

[28] M. Singh and S. Singh. Spatial texture analysis: A comparative study. In *Proceedings of IEEE International Conference on Pattern Recognition (ICPR)*, volume 1, pages 676–679, December 2002. 1, 2

[29] M. Unser. Local linear transforms for texture measurements. *Signal Processing*, 11(1):61–79, July 1986. 2

[30] M. Varma and A. Zisserman. Classifying images of materials: Achieving viewpoint and illumination independence. In *Proceedings of European Conference on Computer Vision (ECCV)*, volume 3, pages 255–271. Springer-Verlag, May 2002. 1, 2

[31] M. Varma and A. Zisserman. A statistical approach to texture classification from single images. *International Journal of Computer Vision (IJCV)*, 62:61–81, 2005. 1, 2, 3, 4, 6, 7, 8

[32] M. Varma and A. Zisserman. A statistical approach to material classification using image patch exemplars. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(11):2032–2047, November 2009. 1

[33] H. Wang, S. Yan, D. Xu, X. Tang, and T. Huang. Trace ratio vs. ratio trace for dimensionality reduction. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007. 3

[34] O. Wang, P. Gunawardane, S. Scher, and J. Davis. Material classification using BRDF slices. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2805–2811, 2009. 2