DOI: 10.2352/ISSN.2470-1173.2016.9.MMRMA-358
Reprinted from
Journal of Imaging Science and Technology® 59(6): 060405-1–060405-9, 2015.

Learning Optimal Incident Illumination using Spectral Bidirectional Reflectance Distribution Function Images for Material Classification

Sandra Skaff▲, Siu-Kei Tin▲, and Manuel Martinello
Canon USA Innovation Center, San Jose, CA, USA
E-mail: sskaff@cusa.canon.com

Abstract. *This article introduces a novel algorithm to learn optimal incident illumination for material classification using spectral bidirectional reflectance distribution function (BRDF) images. The method performs a joint selection of incident angle and spectral band in two steps: (1) clustering and selecting incident angles using statistics on the spectral BRDF images for a specific material, and (2) searching for the optimal angles and spectral bands that maximize material discriminability, which we measure in classification performance. The benefits of reducing the number of incident illumination angles include improving material classification, reducing computational time and storage, and allowing for a less cumbersome and potentially mobile imaging system. The authors show that their approach provides comparable material classification performance when using a reduced number of incident illuminations as compared with when using a larger number. They also compare their approach with prior work.*

INTRODUCTION

Material classification is fundamental to several applications in both consumer and industrial type applications in computer vision and robotics. In consumer applications, material recognition can facilitate visual recognition such as image retrieval. In industrial applications, material recognition can benefit food inspection,[1] recycling,[2] remote sensing,[3] environmental monitoring, and robotic manipulation tasks. Recognizing materials is challenging, as the appearance of a material surface varies with the illumination conditions as well as its color, texture, polarization, and reflectance properties. These properties are characterized by the bidirectional reflectance distribution function (BRDF) of a material.

Capture of a very dense sampling of the BRDF would be ideal;[4] however, such capture requires very precise apparatus as well as accurate registration between images, which can pose unrealistic constraints for real-world applications. Recently, coded illumination has been used to capture spectral BRDF images of uncoated and unpainted raw materials.[5] Additionally, different approaches have been introduced to reduce the capture of material images over a large number of illumination conditions.[6,7]

Inspired by this line of work, we propose to learn the optimal incident illumination that would provide the most discriminability among materials, in turn providing the best material classification performance. The learning is performed jointly for both angle and spectral band in two steps: per-material clustering and selecting incident angles using spectral BRDF images, and searching for the optimal angles and bands that maximize material discriminability. Learning optimal incident illumination would allow for potentially building simpler and less expensive systems which would capture the minimum number of discriminative images needed for material classification. A minimum number of images implies a reduction in computational time and storage, allowing for use in real-time computer vision applications.

Our article makes several contributions. First, we jointly learn optimal incident illumination, both angle and band, using spectral BRDF image slices of raw materials. Second, we extract orientation invariant image-level feature vectors which represent specular, diffuse, and dark components of these slices. Third, we use these feature vectors to cluster angles by spectral BRDF images. This clustering allows for reduction in the search space of optimal angle and band. We perform several experiments, including comparisons with prior work, to demonstrate the effectiveness of our approach.

RELATED WORK

Incident Illumination Learning

While there has been work in the literature on selecting only spectral bands using hyperspectral images[8,9] or learning only angles,[7] learning incident illumination, in terms of both angle and band, for the goal of material classification is very limited. The most recent approach is that of Gu and Liu,[5] which proposes using coded illumination to model the different weightings of angles and LEDs for different material classification tasks. It should be noted that all of the LEDs are used in this case but are assigned different intensities. Moreover, the learning is task dependent, which means that the coded illumination is learned for a combination of materials to be classified. Jehle et al.[7] proposed an algorithm to select incident

▲ IS&T Members.

Received June 30, 2015; accepted for publication Sept. 26, 2015; published online Dec. 03, 2015. Associate Editor: Rita Hofmann-Sievert.

DOI: 10.2352/ISSN.2470-1173.2016.9.MMRMA-358

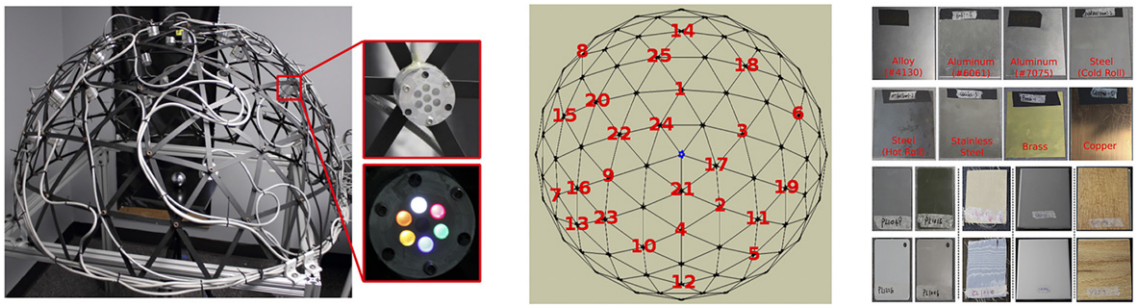

(a) Capture setup (b) Angle positions (c) Materials

Figure 1. Capture setup and materials used by Ref. 5. (a) A photo of the dome and two closeup photos of the six LEDs at one of the 25 angles; (b) the positions of the 25 angles on the top view of the dome; (c) samples of raw materials used in this article: metal (top) and non-metal (bottom).

illumination angles for material classification, which is also task dependent. The authors acquire multiple images of the same scene using an illumination series, which is a set of different illumination conditions. They then select illumination angles that are most important for classification. Moreover, Jehle et al. consider only eight illumination directions which are produced by illuminating a parabolic mirror using parallel light. Their approach is dependent on a parabolic mirror and a very accurate illumination process. Our proposed learning approach, however, comprises two steps, the first one of which learns incident illumination angles based on one material and independently from the classification task. In general, our algorithm is flexible in terms of position and color of incident light sources to select from, and it allows for reduction in size of the imaging setup as well as the captured images and corresponding features. Besides the goal of material classification, there has been work on optimizing both camera and light source positions for measuring the BRDF of 3D objects.[10]

Material Classification

Material classification using texture images has been heavily researched in computer vision.[11–16] Another line of work uses bidirectional texture functions (BTFs) to classify materials.[17,18]

In machine vision, per-pixel material classification has been addressed for various types of materials using polarization[19–21] and near-infrared reflectance,[22] as well as spectral reflectance for printed circuit board inspection.[23] More closely related to this work is material classification using BRDF and BRDF image slices, which is very limited in the computer vision and imaging literature. Hahlweg and Rothe studied material classification using multispectral BRDF imagery;[24] Chen et al. studied material classification using multispectral polarimetric BRDF imagery;[25] Wang et al. used BRDF image slices.[6] Briefly described, the BRDF is a 4D function which specifies the brightness observed in any outgoing direction, in 2D, when light arrives from any incoming direction, in 2D. However, instead of capturing the 4D BRDF, which requires solving for the stereo correspondence between all cameras, the authors capture 2D BRDF slices. The authors then propose several feature vectors and compare their performance in per-pixel material classification. Recently, Shiradkar et al. proposed ink classification using 1D BRDF slices.[26]

As described above, Gu and Liu[5] propose per-pixel material classification using spectral BRDF images captured with coded illumination. More details on their capture system are provided in the next section. Differently from previous work, our work proposes image-based material classification and therefore requires less storage and computation time, which become important factors to consider with large amounts of data.

LEARNING INCIDENT ILLUMINATION

Our goal is to learn the optimal incident illumination that arises in the most discriminative material images for classifying them. We consider a machine-learning-based approach, whereby the angles are learned using the features of the spectral BRDF slice images as taken at the different angles. The use of such a learning approach allows for wide applicability across different imaging capture setups and different incident illumination angles.

In our work, we consider the images and capture setup proposed by Gu and Liu,[5] who showed that imaging of raw material samples using six LED primaries and 25 clusters on a hemispherical geodesic dome with the samples in the center is sufficient for classifying them. Each cluster corresponds to an angle θ on the hemisphere around the sample. At each angle, six images corresponding to differently colored LEDs are captured. Examples of the materials used in image capture as well as the capture setup from Ref. 5 are shown in Figure 1. More details on the setup, the materials, and the capture process are described in the experiments section.

The overall framework of our proposed approach is illustrated in Figure 2 in two steps. In the first step, given two materials, A and B, we learn their K optimal angles from a total set of $N = 25$ angles, $\Theta = \{\theta_1, \ldots, \theta_N\}$ through material-based angle selection. We represent their K corresponding optimal angles by $\Theta^A = \{\theta_1^A, \ldots, \theta_K^A\}$ and $\Theta^B = \{\theta_1^B, \ldots, \theta_K^B\}$, respectively. Each subscript i of θ_i is used to indicate the position of the angle of the light source as labeled in Fig. 1(b). At every θ_i, there are six LEDs, as illustrated in Fig. 1(a), each of which is denoted by its spectral

DOI: 10.2352/ISSN.2470-1173.2016.9.MMRMA-358

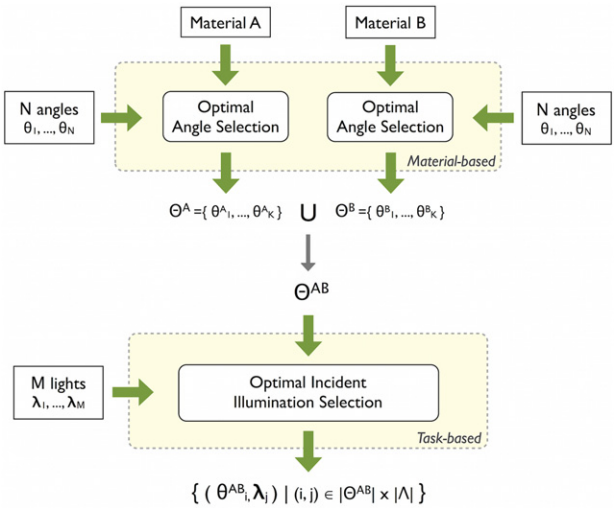

Figure 2. Overall algorithm flowchart in two steps: material-based angle selection, where K is the number of angles selected (top); and task-based optimal incident illumination selection, where n is the number of incident illuminations selected (bottom). The reader is referred to the text for a description of the notations.

band, where $\Lambda = \{\boldsymbol{\lambda}_1, \ldots, \boldsymbol{\lambda}_M\}$ ($M = 6$ in this case) is the set of spectral bands. It should be noted that we use the $\boldsymbol{\lambda}_j$ notation in this case to refer to the wavelength range of LED $j = 1 \ldots M$.

Next, given the optimal angles for each material, our framework performs optimal incident illumination selection to arrive at a set of LEDs, as shown in Fig. 2. If the total set of optimal angles obtained through material-based angle selection (step 1) is denoted by $\Theta^{AB} = \Theta^A \cup \Theta^B$, then the total set of LEDs that are input to the second step of the framework (Fig. 2) can be denoted by $\Theta^{AB} \times \Lambda$. As such, the set of LEDs output by the framework can be denoted by $(\theta_i^{AB}, \boldsymbol{\lambda}_j) \in \Theta^{AB} \times \Lambda$, where $(i, j) \in |\Theta^{AB}| \times |\Lambda|$.

The following subsections are organized as follows. First, we describe the feature vectors computed for spectral BRDF images. Next, we describe the algorithm component for material-based angle selection (top of Fig. 2). In the last subsection, we describe the algorithm component for optimal incident illumination selection, both angle and band (bottom of Fig. 2).

Spectral BRDF Image Feature Representations

We introduce an orientation invariant image-level feature vector representation for spectral BRDF image slices. The advantage of using an image-level representation is that it is robust to noise or artifacts in captured images. Our feature is inspired by the following observation. When we take a picture of a flat material illuminated by a light source, three types of intensity components can be distinguished: specular (S), diffuse (D), and dark (R), as can be seen in Figure 3, which shows BRDF image slices of brass and copper illuminated by six light sources.

The image captured under a specific illumination $(\theta_i, \boldsymbol{\lambda}_j)$, with $(i, j) \in |\Theta| \times |\Lambda|$, is denoted by $I_{\theta_i, \boldsymbol{\lambda}_j}$ and is a grayscale image. In order to consider the different brightness levels of the pixels in the image, we cluster the pixels by brightness using the K-means algorithm. The algorithm labels the pixels belonging to each cluster and learns the means of the clusters in the process. Mathematically, the clustering algorithm aims to partition the pixels x of the image $I_{\theta_i, \boldsymbol{\lambda}_j}$ in N_c clusters, so as to minimize the within-cluster sum of squares as follows:

$$C_1^*, \ldots, C_{N_c}^* = \underset{C_1, \ldots, C_{N_c}}{\operatorname{argmin}} \sum_{k=1}^{N_c} \left[\sum_{x \in C_k} \left\| I_{\theta_i, \boldsymbol{\lambda}_j}(x) - m_k \right\|^2 \right], \tag{1}$$

where $C_1, \ldots, C_{N_c}$ represent the clusters into which the pixels are divided and m_k indicates the mean of the pixel brightnesses in cluster C_k.

Equation (1) is optimized using a two-step iterative process, which includes an assignment step and an update step. The first step assigns a cluster label to each pixel whose

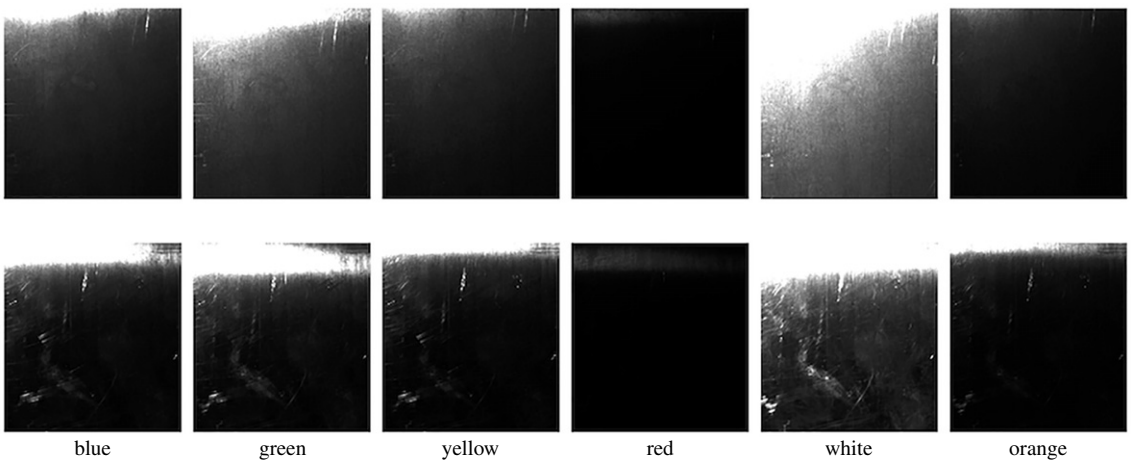

Figure 3. BRDF image slices of brass (top) and copper (bottom) when imaged under blue, green, yellow, red, white, and orange LEDs using Gu and Liu's setup.[5]

©2016 Society for Imaging Science and Technology
DOI: 10.2352/ISSN.2470-1173.2016.9.MMRMA-358

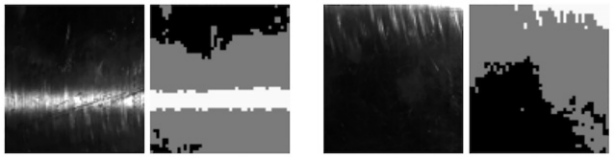

Figure 4. Illumination components used to compute the feature vectors, shown for two image slices of copper.

mean yields the least within-cluster sum of squares:

$$C_k^{(t)} = \{x \in I_{\theta_i,\boldsymbol{\lambda}_j} : \|I_{\theta_i,\boldsymbol{\lambda}_j}(x) - m_k^{(t)}\|^2 \leq \|I_{\theta_i,\boldsymbol{\lambda}_j}(x) - m_l^{(t)}\|^2\} \quad \forall l,\ 1 \leq l \leq N_c, \tag{2}$$

while the update step computes the new means to be the centroids of the pixels in the new clusters:

$$m_k^{(t+1)} = \frac{1}{|C_k^{(t)}|} \sum_{x \in C_k^{(t)}} I_{\theta_i,\boldsymbol{\lambda}_j}(x)\,, \quad \text{with } 1 \leq k \leq N_c, \tag{3}$$

where $m_1^{(t)}, \ldots, m_{N_c}^{(t)}$ indicate the set of N_c means at iteration t. The number of clusters N_c is set by the user, and we set $N_c = 3$. The hypothesis is that three clusters would allow us to group pixels belonging to specular, diffuse, and dark regions of the image. Experimentation also aligned with the observation that three clusters is the optimal number as it provided superior material classification performance to when using 2, 4, 5, or 6 clusters. We denote the image regions or clusters obtained after convergence by $I^S_{\theta_i,\boldsymbol{\lambda}_j}$, $I^D_{\theta_i,\boldsymbol{\lambda}_j}$, and $I^R_{\theta_i,\boldsymbol{\lambda}_j}$, with the corresponding means estimated by the algorithm to $\mu^S_{\theta_i,\boldsymbol{\lambda}_j}$, $\mu^D_{\theta_i,\boldsymbol{\lambda}_j}$, and $\mu^R_{\theta_i,\boldsymbol{\lambda}_j}$, respectively. Figure 4 shows the three areas for two BRDF slices of copper, imaged under different incident illumination angles and spectral bands.

Given the pixel clusters and their relative means, we compute two types of feature vector, denoted $\mathbf{X}_{\theta_i,\boldsymbol{\lambda}_j}$, for each image slice corresponding to illumination $(\theta_i, \boldsymbol{\lambda}_j)$. These are the mean feature and the histogram feature, as described below.

- Mean feature (MF):

$$\mathbf{X}_{\theta_i,\boldsymbol{\lambda}_j} = \left[\mu^S_{\theta_i,\boldsymbol{\lambda}_j}, \mu^D_{\theta_i,\boldsymbol{\lambda}_j}, \mu^R_{\theta_i,\boldsymbol{\lambda}_j}\right]. \tag{4}$$

- Histogram feature (HF):

$$\mathbf{X}_{\theta_i,\boldsymbol{\lambda}_j} = \left[\frac{P^S_{\theta_i,\boldsymbol{\lambda}_j}}{P_{\theta_i,\boldsymbol{\lambda}_j}}, \frac{P^D_{\theta_i,\boldsymbol{\lambda}_j}}{P_{\theta_i,\boldsymbol{\lambda}_j}}, \frac{P^R_{\theta_i,\boldsymbol{\lambda}_j}}{P_{\theta_i,\boldsymbol{\lambda}_j}}\right]. \tag{5}$$

Here, $P_{\theta_i,\boldsymbol{\lambda}_j} = P^S_{\theta_i,\boldsymbol{\lambda}_j} + P^D_{\theta_i,\boldsymbol{\lambda}_j} + P^R_{\theta_i,\boldsymbol{\lambda}_j}$ indicates the total number of pixels in an image slice captured under illumination $(\theta_i, \boldsymbol{\lambda}_j)$. The numbers of pixels in the specular, diffuse, and dark components for the same image slice are indicated by $P^S_{\theta_i,\boldsymbol{\lambda}_j}$, $P^D_{\theta_i,\boldsymbol{\lambda}_j}$, and $P^R_{\theta_i,\boldsymbol{\lambda}_j}$, respectively.

Finally, we can represent the feature vector for material A, comprising multiple spectral BRDF image slices, as such:

$$\mathbf{X}^A = \left[\mathbf{X}_{\theta_1,\boldsymbol{\lambda}_1}, \mathbf{X}_{\theta_1,\boldsymbol{\lambda}_2}, \ldots, \mathbf{X}_{\theta_i,\boldsymbol{\lambda}_j}, \ldots, \mathbf{X}_{\theta_N,\boldsymbol{\lambda}_M}\right], \quad \text{with } (i,j) \in |\theta| \times |\Lambda|. \tag{6}$$

Material-Based Angle Selection

Our algorithm learns incident illumination angles given the set of all possible illumination angles $\theta = \{\theta_1, \ldots, \theta_N\}$, as well as the BRDF slice images of one material. Figure 5 shows the steps of the material-based angle selection algorithm. First, M-dimensional feature vectors are extracted from the image samples corresponding to each of the angles, as denoted with the unfilled circles: each circle represents the feature vector of an image corresponding to one spectral band light source; $N \times \#$ of samples is the number of feature vectors, where N is the total number of light sources. After rearranging, the feature vectors are clustered to obtain K clusters, where K is equal to the number of angles to be selected. Each filled point corresponds to one of the M-dimensional feature vectors. The means, μ_1, μ_2, μ_3, are

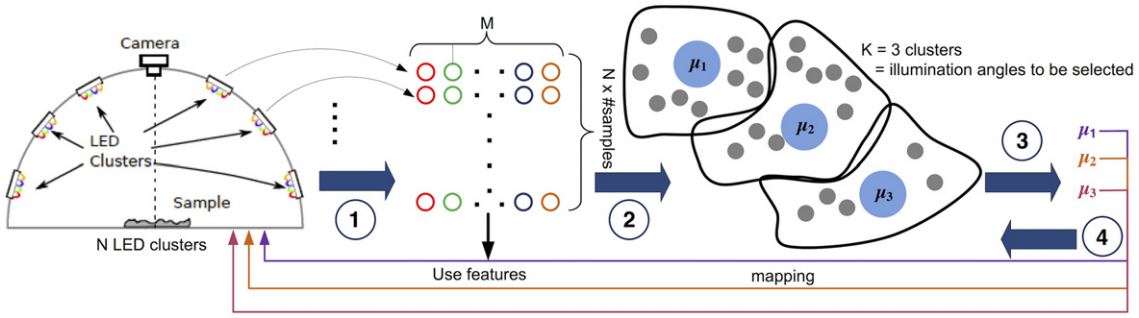

Figure 5. Material-based angle selection framework. From steps 1 to 4. (1) M-dimensional feature vectors are extracted from the image samples corresponding to one angle position (all LEDs), as denoted with the unfilled circles: each circle represents the feature of an image corresponding to one light source; $N \times \#$ of samples is the number of feature vectors, where N is the total number of light sources. (2) The feature vectors are clustered to obtain K clusters equal to the number of angles to be selected; each filled point in space corresponds to the M-dimensional feature vector of a material type. (3) The means, μ_1, μ_2, μ_3, are mapped to the initial feature vectors to select one closest feature vector to each using a distance metric. (4) Each of the three selected feature vectors is mapped to a light source position given the correspondence between the feature vectors and the illumination angles used in capturing the corresponding image. The picture of the dome with the LED clusters is obtained from Ref. 5.

DOI: 10.2352/ISSN.2470-1173.2016.9.MMRMA-358

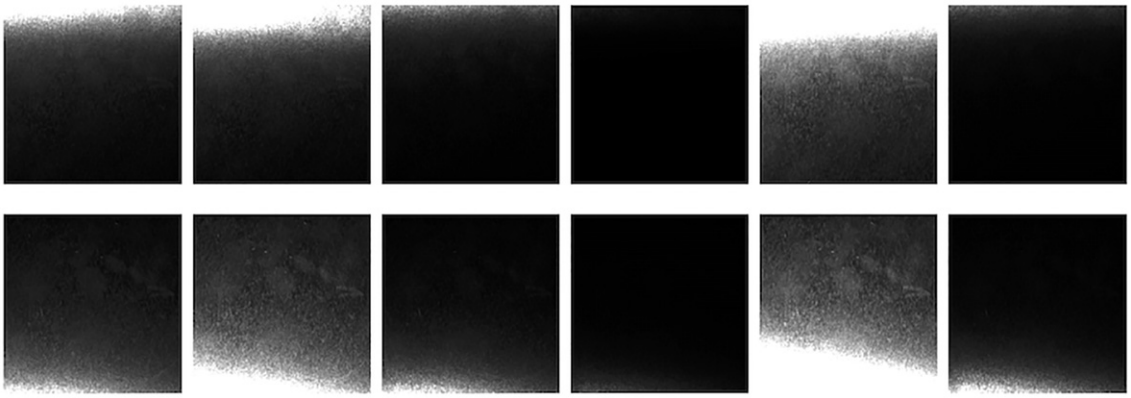

(a) 2 samples of BRDF image slices from cluster #1

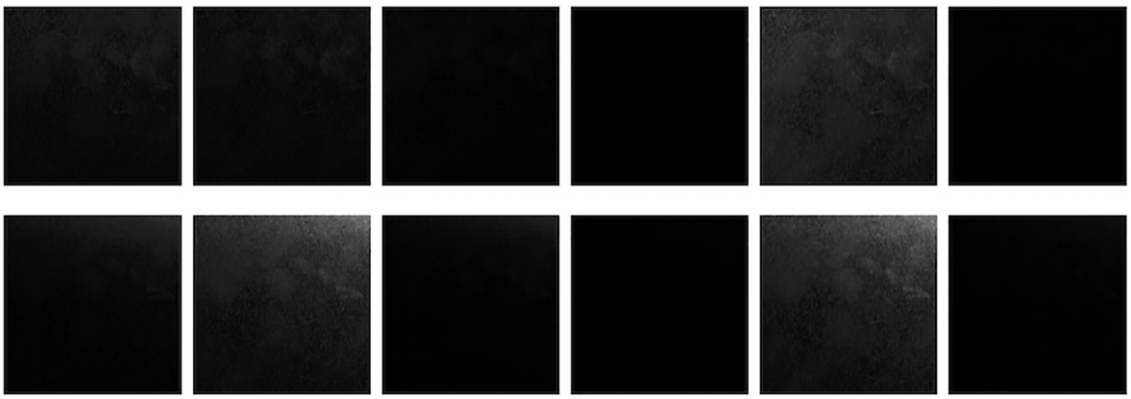

(b) 2 samples of BRDF image slices from cluster #2

Figure 6. Clustering in material-based angle selection for aluminum. The images are BRDF slices for all bands (from left to right) corresponding to two angles which belong to (a) cluster #1 and (b) cluster #2.

mapped to the initial feature vectors to select one closest feature vector to each using the Euclidean distance metric (other metrics can be used). Next, each of the three selected feature vectors is mapped to a light source position given the correspondence between the feature vectors and the illumination angles used in capturing the corresponding image. Figure 6 shows samples of BRDF image slices as grouped into clusters.

Task-based band selection

At each selected optimal angle, light sources of different bands can be placed. Our goal is to select incident illumination with which to illuminate materials such that we can best discriminate among them. Assuming that the task is to classify two materials, A and B, we indicate their learned angles (as described in a previous subsection) by $(\theta_1^A, \theta_2^A, \ldots, \theta_K^A)$ and $(\theta_1^B, \theta_2^B, \ldots, \theta_K^B)$, respectively. The set of learned angles obtained from material-based angle selection can be represented by

$$\Theta^{AB} = \Theta^A \cup \Theta^B = \{\theta_1^A, \theta_2^A, \ldots, \theta_K^A\} \cup \{\theta_1^B, \theta_2^B, \ldots, \theta_K^B\}, \tag{7}$$

with $|\Theta| \leq |\Theta^A| + |\Theta^B|$, since there might be common learned angles among A and B. As mentioned previously, the set of LEDs available for this task is thus represented by $\Theta^{AB} \times \Lambda$, where Λ denotes the set of LEDs at each angle.

Next, we maximize the earth mover's distance D_{EMD}[27] between the image feature vectors of the samples of A and B corresponding to $(\theta_i^{AB}, \boldsymbol{\lambda}_j)$, where $(i, j) \in |\Theta^{AB}| \times |\Lambda|$:

$$D_{EMD}^{AB}(\Theta^{AB}, \Lambda) = \frac{\sum_{u \in A} \sum_{v \in B} f_{uv} \|\mathbf{X}_{\theta_i, \boldsymbol{\lambda}_j}^A(u) - \mathbf{X}_{\theta_i, \boldsymbol{\lambda}_j}^B(v)\|}{\sum_{u \in A} \sum_{v \in B} f_{uv}}. \tag{8}$$

Here, f_{uv} is the flow between $\mathbf{X}_{\theta_i, \boldsymbol{\lambda}_j}^A(u)$ and $\mathbf{X}_{\theta_i, \boldsymbol{\lambda}_j}^B(v)$; u and v represent the samples of A and B, respectively. The flow f_{uv} is computed such that it minimizes the following cost, denoted by C_{EMD}:[27]

$$C_{EMD} = \sum_{u \in A} \sum_{v \in B} f_{uv} \|\mathbf{X}_{\theta_i, \boldsymbol{\lambda}_j}^A(u) - \mathbf{X}_{\theta_i, \boldsymbol{\lambda}_j}^B(v)\|, \tag{9}$$

subject to the following constraints:

$$\begin{aligned} f_{uv} &\geq 0 \quad u \in A, v \in B; \\ \sum_{u \in A} f_{uv} &\leq \mathbf{w}_v^B \quad v \in B; \\ \sum_{v \in B} f_{uv} &\leq \mathbf{w}_u^A \quad u \in A; \\ \sum_{u \in A} \sum_{v \in B} f_{uv} &= \min\left(\sum_{u \in A} \mathbf{w}_u^A, \sum_{v \in B} \mathbf{w}_v^B\right). \end{aligned} \tag{10}$$

DOI: 10.2352/ISSN.2470-1173.2016.9.MMRMA-358

Table I. Mean classification accuracies for metal (left), mixed (center), and non-metal (right) sets over the different numbers n of selected illuminations (θ, λ) or LEDs, ranging from one to five. The accuracies are shown for two types of feature: the MF in the second row and the HF in the third row.

Feature	% accuracy over number of LEDs (n)														
	1			2			3			4			5		
MF	76.8	85.2	80.1	81.2	84.4	81.8	82.5	87.5	79.8	83.8	86.9	78.1	85.2	88.7	77.6
HF	74.6	74.4	71.6	77.3	77.7	75.6	77.3	75.6	78.7	78.7	76.4	78.1	79.6	78.5	78.1

We denote the weight vectors $\mathbf{w}_{\mathbf{X}^A_{\theta_i,\lambda_j}}$ and $\mathbf{w}_{\mathbf{X}^B_{\theta_i,\lambda_j}}$ by $\mathbf{w}^A_u$ and $\mathbf{w}^B_v$, respectively, for simplicity. The first constraint ensures that the flow is positive, the second constraint ensures that the sum of the flows from the feature vectors of samples from A is limited by the sum of weights corresponding to the feature vectors of samples from B, while the third constraint ensures that the sum of the flows from the feature vectors of samples from B is limited by the sum of weights corresponding to the feature vectors of samples from A; the last constraint is the normalization factor of the earth mover's distance (used in Eq. (8)), and is set to the minimum of the sum of weights.

Finally, we choose n incident illuminations $(\theta^{AB}_i, \boldsymbol{\lambda}_j) \in \Theta^{AB} \times \Lambda$ which give the maximum values of D^{AB}_{EMD}. The complexity of computing this distance is $O(|A||B|)$, and therefore reducing the space of (Θ^{AB}, Λ) over which to compute this distance will aid in reducing computational complexity. Additionally, the reduction in search space does not compromise on classification accuracy, as shown in experiment.

MATERIAL CLASSIFICATION EXPERIMENTS

We evaluate the performance of our approach using Gu and Liu's spectral BRDF image database,[5] which comprises ten material categories with four or more samples in each (http://compimg1.cis.rit.edu/data/metal/, Fig. 1(c)). In total, there are ten material classes, six metal ones and four non-metal ones. The metal ones are alloy, copper, brass, stainless steel, two types of steel, and four types of aluminum; the non-metal ones are fabric, ceramic, plastic, and wood. Each material sample is imaged 150 times using each of the LEDs of the imaging dome, yielding 150 spectral BRDF image slices. The LEDs in the dome are grouped into 25 LED clusters, each placed at a different angle. Each cluster comprises six LEDs of the following colors: blue, green, yellow, red, white, and orange. Fig. 3 shows image slices of brass and copper samples imaged using all six LEDs (or bands) at angle location 12.

The experimental paradigm followed in this article can be described as follows. Feature vectors (HF or MF) are extracted for the BRDF image slices. The support vector machine (SVM) was used for classification using the linear kernel and setting the regularization parameter C to 10. A radial kernel was also tested in several cases; however, no improvement in classification accuracy was found. We constructed sets of two material types each from the total set of material categories to obtain 50 classification tasks: 23 sets comprise only metal materials (metal sets), 16 sets comprise a mix of metal and non-metal materials (mixed sets), and 11 sets comprise only non-metal materials (non-metal). We choose different combinations or folds of training and testing data. In total, we consider 16 folds of training and testing data, where 75% of the samples are used for training and 25% are used for testing. The classification accuracy computed is then taken to be the average over the accuracies computed for each of the test sets of the 16 folds.

Incident Illumination Learning Performance

We evaluate material classification given sample images captured under illumination from the selected light sources as described above. Two types of features, the HF and MF, as described previously are extracted for the images. The evaluation is performed for the metal, mixed, and non-metal sets.

Table I shows the results of classification (% accuracy) for two types of feature vector and for different numbers of selected incident illuminations (LEDs in this case), from one to five for the different set types. Additionally, we compare the performance when selecting the LEDs as given by our algorithm with those as given by material-based angle selection (described above) in Table II. Material-based angle selection is in itself task independent; however, in order to evaluate we use images corresponding to the selected angles, for the different sets. Moreover, the table shows a comparison with the case when all incident illuminations are used for benchmarking.

Two general observations can be made. First, the MF provides better performance than the HF. This can be attributed to the fact that the histograms in the HF do not represent reflectance information, while the means in the MF do. Second, our material-based learning provides comparable results to when all angles and bands are used, especially in the metals case. Additionally, to further simplify the setup, we can use one to five LEDs (Table I). Since the MF is the better performing feature, we use it for the remainder of the experimental results presented in this article.

Figures 7(a) and (b) show the mean classification accuracies (%) over 16 folds for three metal classification tasks and three mixed classification tasks. The accuracies are plotted as a function of n, which is the number of incident illuminations or LEDs selected; n is taken to be 1, 2, 3, 4, 5.

In the case of brass/copper (Fig. 7a), we observe that the algorithm provides 99.9% accuracy when using two, three, four, or five LEDs. In the case of two, LEDs at angle

DOI: 10.2352/ISSN.2470-1173.2016.9.MMRMA-358

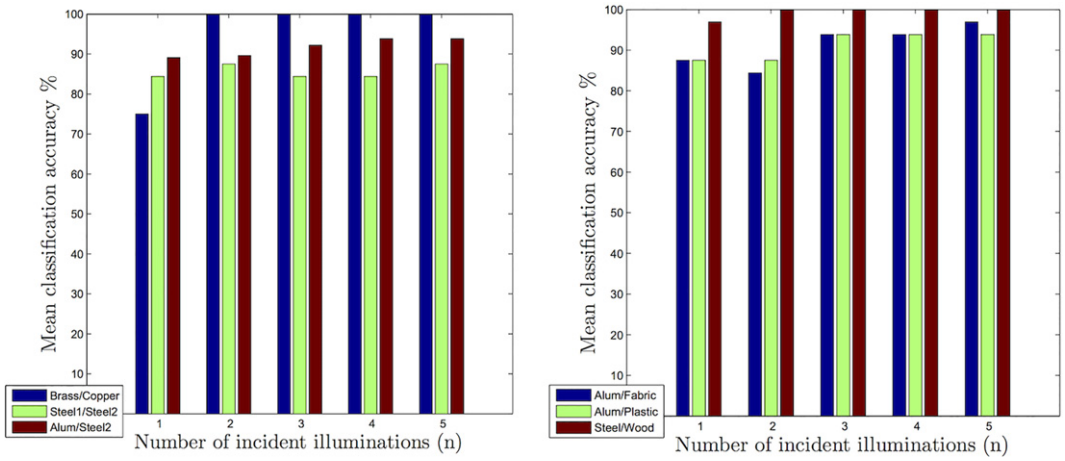

Figure 7. The mean classification accuracies (%) over 16 folds for classifying (a) metals only, namely, brass/copper, aluminum/steel, and two types of steel; and (b) mixed sets, namely, aluminum/fabric, aluminum/plastic, and steel/wood. The accuracies are plotted as a function of the number *n* of incident illuminations selected.

Table II. Mean classification accuracies for different illumination selection algorithms for metal (left), mixed (center), and non-metal (right) when selecting five incident illuminations by angle and spectral band using task-based selection (optimal illumination), three incident illuminations by angle only using material-based selection, and when using all light sources (no selection). The accuracies are shown for two types of feature: the MF in the second row and the HF in the third row.

Feature	% accuracy								
	Optimal illum.			Angle only			None		
MF	85.2	88.7	77.6	87.0	89.1	80.4	87.0	92.4	77.8
HF	79.6	78.5	78.1	81.8	81.0	80.1	81.8	85.6	85.0

positions 18 and 16 (in Fig. 1(b)) are chosen in all folds. Following the notation in the previous section, we can denote the angle at position 18 by θ_{18}; therefore, we denote the elevation and azimuth angles corresponding to this angle by $(\theta^e_{18}, \theta^z_{18})$, and we can write that angle 18 is positioned at $(\theta^e_{18} = 40.6°, \theta^z_{18} = -50.0°)$, while angle 16 is positioned at $(\theta^e_{16} = 45.2°, \theta^z_{16} = 92.7°)$. One of the LEDs is green and the second is orange, which can be expected given the colors of brass and copper. For aluminum/steel, the classification accuracy is 89.1% and 89.6% when using one or two LEDs, respectively, as compared with 93.8% when using four or five LEDs, with the accuracies monotonically increasing from one to five LEDs. In the four LED case, the selected LEDs are at three different angles, which are 8, 16, and 12. Angle 8 is positioned at $(\theta^e_8 = 31.8°, \theta^z_8 = 36.0°)$, angle 16 is positioned at $(\theta^e_{16} = 45.2°, \theta^z_{16} = 92.7°)$, and angle 12 is positioned at $(\theta^e_{12} = 29.7°, \theta^z_{12} = 166.0°)$. In the case of distinguishing two steel types, an accuracy of 87.5% is achieved when using two LEDs at two different angles (8 and 16).

In the case of aluminum/fabric (Fig. 7b), five LEDs are needed to provide the best classification accuracy (96.9%). Among the selected LEDs are two white ones at two different angles. In the case of aluminum/plastic, three LEDs provide maximum classification accuracy (93.8%), where

Table III. Mean classification accuracies when using illumination learned with our method as well as when using incident illumination selected at random. In each case the image slices corresponding to the illumination are used for classification.

Set type	Proposed (% acc.)	Random five (% acc.)
Metal	85.2	67.8
Mixed	88.7	80.3
Non-metal	77.6	76.8

two LEDs are at angle 8 $(\theta^e_8 = 31.8°, \theta^z_8 = 36.0°)$ and one LED is at angle 16 $(\theta^e_{16} = 45.2°, \theta^z_{16} = 92.7°)$. In the case of steel/wood, two LEDs provide maximum classification accuracy (99.9%). The selected LEDs are a green one at angle 18 $(\theta^e_{18} = 40.6°, \theta^z_{18} = -50.0°)$ and an orange one at angle 16 $(\theta^e_{16} = 45.2°, \theta^z_{16} = 92.7°)$, similarly to the brass/copper case, which is expected given the color of wood. In short, the results demonstrate that our algorithm can be a useful tool to build a smaller capture setup while not compromising on classification accuracy, and the best reduction in dimension of the setup is dependent on the classification task.

Benchmarking Performance

We also compare the classification performance when five selected LEDs are used in capture with when using five LEDs selected at random from all possible LEDs. Table III shows the mean classification accuracies for each of the metal, mixed, and non-metal sets corresponding to our method and the different combinations of LEDs selected at random. The table shows that our method outperforms the case when illumination is selected at random significantly for the metal sets (by 17.4% in classification accuracy) and by 8.4% for the mixed sets, but our approach does not do so well for the non-metal sets. We believe that the reason behind these results is that our features are not suitable for representing the texture of non-metals, as they do not retain inter-pixel

©2016 Society for Imaging Science and Technology
DOI: 10.2352/ISSN.2470-1173.2016.9.MMRMA-358

Table IV. Comparing mean classification accuracies (%) on raw materials given by our method and Gu and Liu's approach.[5]

Task	Fisher light[5]	SVM light[5]	Proposed method
Aluminum versus steel	89.9	97.3	93.8
Brass versus copper	98.5	99.0	99.9
Ceramic versus plastic	94.4	95.5	78.1

Table V. Comparing mean classification accuracies (%) on raw materials given by our method and random forest as used by Jehle et al. for learning illumination angles.[7]

Set type	Random forest[7]	Proposed method
Metal	77.1	85.2
Mixed	78.5	88.7
Non-metal	69.8	77.6

information. We expect that using texture features would significantly boost the performance.

We perform two comparisons with prior work. The first one compares the performance of our method with that of Gu and Liu's method,[5] as shown in Table IV. As described in a previous section, Gu and Liu used coded illumination to model the different weightings of angles and LEDs for different material classification tasks. It should be noted that all of the LEDs were used in their case but were assigned different intensities. Moreover, the learning they performed was task dependent, which means that the coded illumination was learned for a combination of materials to be classified. We list the performance of Gu and Liu's method for three binary classification tasks, aluminum/steel, brass/copper, and ceramic/plastic, in two cases: when they use the Fisher light and the SVM light to learn the weights on the LEDs. It should be noted that the accuracies depicted for Ref. 5 are for per-pixel classification, while our method calculates per-image accuracies. Our method provides comparable classification results to Ref. 5 while requiring fewer images and a smaller capture setup (five LEDs in this case) in the case of metals. However, our method does not fare as well for the case of non-metals, as shown in the ceramic/plastic case. Such performance is expected since our approach does not represent texture information, which is important for classifying non-metals.

Second, we compare material classification performance on our data when using our approach and when using random forest[28] as used by Jehle et al.[7] to learn illumination angles. As mentioned in a previous section, the authors in Ref. 7 acquire multiple images of the same scene using an illumination series, which is a set of different illumination conditions. They then select illumination angles that are most important for classification. Moreover, Jehle et al. consider only eight illumination directions which are produced by illuminating a parabolic mirror using parallel light. Their approach is dependent on a parabolic mirror and a very accurate illumination process. It should be noted that the learning performed by Jehle et al. is task dependent, while ours is task independent. Table V shows the mean classification accuracies for the metal, mixed, and non-metal sets for both Ref. 7 and our proposed method. The results show that our method outperforms Jehle et al.'s method on the metal and mixed sets. Again, our method does not fare as well for the non-metal sets, as expected.

CONCLUSIONS

This article presented an approach to jointly learn optimal incident illumination, both angle and spectral band, using spectral BRDF image slices of raw materials. We extracted orientation invariant image-level feature vectors which represent three intensity components of these slices. We demonstrated the effectiveness of these features to learn optimal illumination for material classification. Our joint learning algorithm includes a task-independent angle selection component, where we used the extracted feature vectors to cluster angles. The clustering allows for reduction in the search space of optimal angle and band. We also performed several experiments, including comparisons with prior work. The results demonstrated that our algorithm can achieve the same or better performance for separating metals as well as separating metals from non-metals, while at the same time providing a reduction in the number of light sources.

In the future, we would like to investigate how our approach can be used for separating non-metals, such that the classification performance is similar to that of separating metals among each other or metals from non-metals. Non-metals are best represented using texture features, while our approach uses BRDF-based features. Additionally, our approach is designed to select optimal incident illumination angles for each material in a task. Such a selection is in line with the BRDF models of metals, where a few incident illumination angles can provide enough information for distinguishing metals. However, the same hypothesis does not necessarily hold for non-metals, thus placing these types of materials at a disadvantage in our approach. One way to tackle non-metals would be to understand whether angle selection is needed for non-metals. Perhaps an alternative framework that performs only LED selection can be used instead. Another aspect that can be further investigated is which machine learning framework best suits the purposes of our framework. In our work, we use K-means as a first step and distance-based metrics as a second step. One could think of other approaches such as linear discriminant analysis or Gaussian mixture models instead.

REFERENCES

[1] *Computer Vision Technology for Food Quality Evaluation*, edited by D.-W. Sun (Academic Press, 2007).
[2] J. Hendrik, K. A. Massey, E. Whitham, B. Bras, and M. D. Russell, "Technologies for the identification, separation, and recycling of automotive plastics," Int. J. Environ. Conscious Design Manuf. **6**, 37 (1997).
[3] F. Melgani and L. Bruzzone, "Classification of hyperspectral remote sensing images with support vector machines," IEEE Trans. Geosci. Remote Sens. **42**, 1778 (2004).

DOI: 10.2352/ISSN.2470-1173.2016.9.MMRMA-358

[4] 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 (1999).

[5] J. Gu and C. Liu, "Discriminative illumination: per-pixel classification of raw materials based on optimal projections of spectral BRDF," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2012), pp. 797–804.

[6] O. Wang, P. Gunawardane, S. Scher, and J. Davis, "Material classification using BRDF slices," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2009), pp. 2805–2811.

[7] M. Jehle, C. Sommer, and B. Jähne, "Learning of optimal illumination for material classification," Pattern Recognit. **6376**, 563 (2010).

[8] A. Martinez-Usó, F. Pla, J. Sotoca, and P. Garcia-Sevilla, "Clustering-based hyperspectral band selection using information measures," IEEE Trans. Geosci. Remote Sens. **45**, 4158 (2007).

[9] V. Kumar, J. Hahn, and A. M. Zoubir, "Band selection for hyperspectral images based on self-tuning spectral clustering," *Proc. IEEE European Signal Processing Conf.* (IEEE, Piscataway, NJ, 2013).

[10] H. P. A. Lensch, J. Lang, A. M. Sá, and H.-P. Seidel, "Planned sampling of spatially varying BRDFs," Comput. Graph. Forum **22**, 473 (2003).

[11] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," IEEE Trans. Pattern Anal. Mach. Intell. **24**, 971 (2002).

[12] M. Varma and A. Zisserman, "A statistical approach to texture classification from single images," Int. J. Comput. Vis. **62**, 61 (2005).

[13] M. Varma and A. Zisserman, "A statistical approach to material classification using image patch exemplars," IEEE Trans. Pattern Anal. Mach. Intell. **31**, 2032 (2009).

[14] E. Hayman, B. Caputo, M. Fritz, and J.-O. Eklundh, "On the significance of real-world conditions for material classification," *Proc. European Conf. on Computer Vision* (Springer, Berlin Heidelberg, 2004), pp. 253–266.

[15] B. Caputo, E. Hayman, and P. Mallikarjuna, "Class-specific material categorisation," *Proc. IEEE Int'l. Conf. on Computer Vision* (IEEE, Piscataway, NJ, 2005), pp. 1597–1604.

[16] W. Li and M. Fritz, "Recognizing materials from virtual examples," *Proc. European Conf. on Computer Vision* (Springer, Berlin Heidelberg, 2012), pp. 345–358.

[17] C. Liu, G. Yang, and J. Gu, "Learning discriminative illumination and filters for raw material classification with optimal projections of bidirectional texture functions," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2013), pp. 1430–1437.

[18] M. Weinmann, J. Gall, and R. Klein, "Material classification based on training data synthesized using a BTF database," *Proc. European Conf. on Computer Vision* (Springer, Switzerland, 2014), pp. 156–171.

[19] L. B. Wolff, "Polarization-based material classification from specular reflection," IEEE Trans. Pattern Anal. Mach. Intell. **12**, 1059 (1990).

[20] H. Chen and L. B. Wolff, "Polarization phase-based method for material classification in computer vision," Int. J. Comput. Vis. **28**, 73–83 (1998).

[21] S. Tominaga and A. Kimachi, "Polarization imaging for material classification," Opt. Eng. **47**, 123201 (2008).

[22] N. Salamati, C. Fredembach, and S. Süsstrunk, "Material classification using color and NIR images," *Proc. IS&T/SID's Seventeenth Color Imaging Conf.* (IS&T, Springfield, VA, 2009), pp. 216–222.

[23] A. Ibrahim, S. Tominaga, and T. Horiuchi, "Spectral imaging method for material classification and inspection of printed circuit boards," Opt. Eng. **49**, 057201 (2010).

[24] C. Hahlweg and H. Rothe, "Classification of optical surface properties and material recognition using multi-spectral BRDF data measured with a semi-hemispherical spectro-radiometer in VIS and NIR," Proc. SPIE **5965**, 150–161 (2005).

[25] C. Chen, Y.-Q. Zhao, L. Luo, D. Liu, and Q. Pan, "Robust materials classification based on multispectral polarimetric BRDF imagery," Proc. SPIE **7384**, (2009).

[26] R. Shiradkar, L. Shen, G. Landon, S.-H. Ong, and P. Tan, "A new perspective on material classification and ink identification," *Proc. IEEE Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2014), pp. 2275–2282.

[27] Y. Rubner, C. Tomasi, and L. J. Guibas, "The earth mover's distance as a metric for image retrieval," Int. J. Comput. Vis. **40**, 99 (2000).

[28] A. Bosch, A. Zisserman, and X. Muoz, "Image classification using random forests and ferns," *Proc. IEEE Int'l. Conf. on Computer Vision* (IEEE, Piscataway, NJ, 2007), pp. 1–8.