

# Discriminative Illumination: Per-Pixel Classification of Raw Materials based on Optimal Projections of Spectral BRDF

Jinwei Gu

Chao Liu

Rochester Institute of Technology, Rochester, NY

## Abstract

*Classifying raw, unpainted materials — metal, plastic, ceramic, fabric, etc.— is an important yet challenging task for computer vision. Previous works measure subsets of surface spectral reflectance as features for classification. However, acquiring the full spectral reflectance is time-consuming and error-prone. In this paper, we propose to use coded illumination to directly measure discriminative features for material classification. Optimal illumination patterns—which we call “discriminative illumination”—are learned from training samples, after projecting to which, the spectral reflectance of different materials are maximally separated. This projection is automatically realized by the integration of incident light for surface reflection. While a single discriminative illumination is capable of linear, two-class classification, we show that multiple discriminative illuminations can be used for nonlinear and multi-class classification. We also show theoretically the proposed method has higher signal-to-noise ratio than previous methods due to light multiplexing. Finally, we construct a LED-based multi-spectral dome and use the discriminative illumination method for classifying a variety of raw materials, including metal (aluminum, alloy, steel, stainless steel, brass and copper), plastic, ceramic, fabric and wood. Experimental results demonstrate the effectiveness of the proposed method.*

## 1. Introduction

Classifying materials — metal, plastic, ceramic, fabric, paint, etc.. — has significant implications for both scientific research and industrial applications across many disciplines, such as remote sensing [16], food inspection [25], mineralogy, and recycling [10]. In computer vision, we primarily focus on uncoated or unpainted raw materials since we are limited to appearance related features, such as color, Bidirectional Reflectance Distribution Function (BRDF), texture, translucency, and polarization. Figure 1 shows some examples of such materials.

Even for uncoated raw materials, appearance-based classification is still challenging, because appearance changes with object shape, illumination, and viewing condition.

Fully describing the appearance of a scene requires a 14-D function, *i.e.*, two plenoptic functions [1]. While some previous work measures subsets of this function for material classification [28, 13], due to the intrinsic high dimensionality, material classification has had limited progress compared to that of object recognition.

In this paper, we focus on per-pixel classification of raw materials based on spectral BRDFs. Instead of first sparsely sampling subsets of this high dimensional function and then performing classification, we propose to use coded illumination to directly measure discriminative features, *i.e.*, projections of spectral BRDFs, for classification. The optimal coded illumination—which we call *discriminative illumination*—is learned from training samples, after projecting to which, the spectral BRDFs of different materials can be maximally separated. The projection operation is automatically realized by the integration of incident light for surface reflection in an imaging system.

While a single discriminative illumination is capable of linear, two-class classification, we show multiple discriminative illumination patterns can be used for multi-class and nonlinear classification. The proposed discriminative illumination method is more economical than conventional methods of using raw material measurements for classification in terms of the number of captured images — this enables the classification of materials that changes with time. In addition, we derive that the discriminative illumination method results in higher Signal-to-Noise Ratio (SNR) than conventional methods thanks to light multiplexing.

We construct a LED-based multi-spectral dome light (as shown in Figs. 2(a)(b)) and use it as a prototype to implement the proposed discriminative illumination for classifying a variety of raw materials, including metal (aluminum, alloy steel, stainless steel, cold roll and hot roll steel, brass, and copper), ceramic, plastic, fabric and wood, as shown in Fig. 1. Experimental results demonstrate the effectiveness of the proposed method (see the supplementary video).

## 2. Related Work

### Per-Pixel Material Classification in Machine Vision

There are several works aiming at per-pixel material classification using various low-level appearance features, such

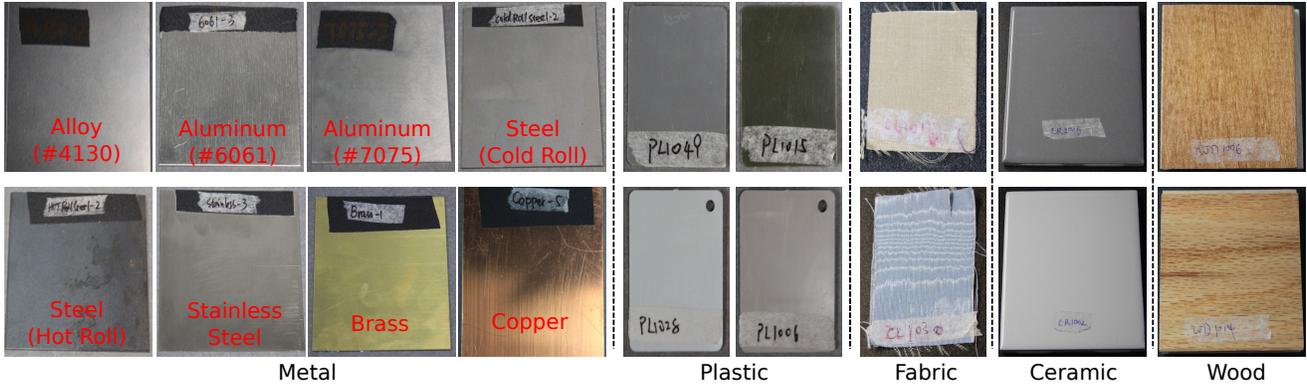


Figure 1. Samples of the database of raw materials tested in this paper. We focus on unpainted, raw materials which are classified based on their surface spectral BRDFs. The database includes metal, plastic, fabric, ceramic, and wood. Within the class of metal, we have samples of alloy (#4130), aluminum (#5052, #6061, #2024, #7075), steel (cold roll and hot roll), stainless steel, brass, and copper. We intentionally choose samples that have similar colors so that most samples cannot easily be classified using only color information. In total, there are 100 sample plates. The measured spatially-varying spectral BRDFs can be downloaded at [www.cis.rit.edu/mcsl](http://www.cis.rit.edu/mcsl).

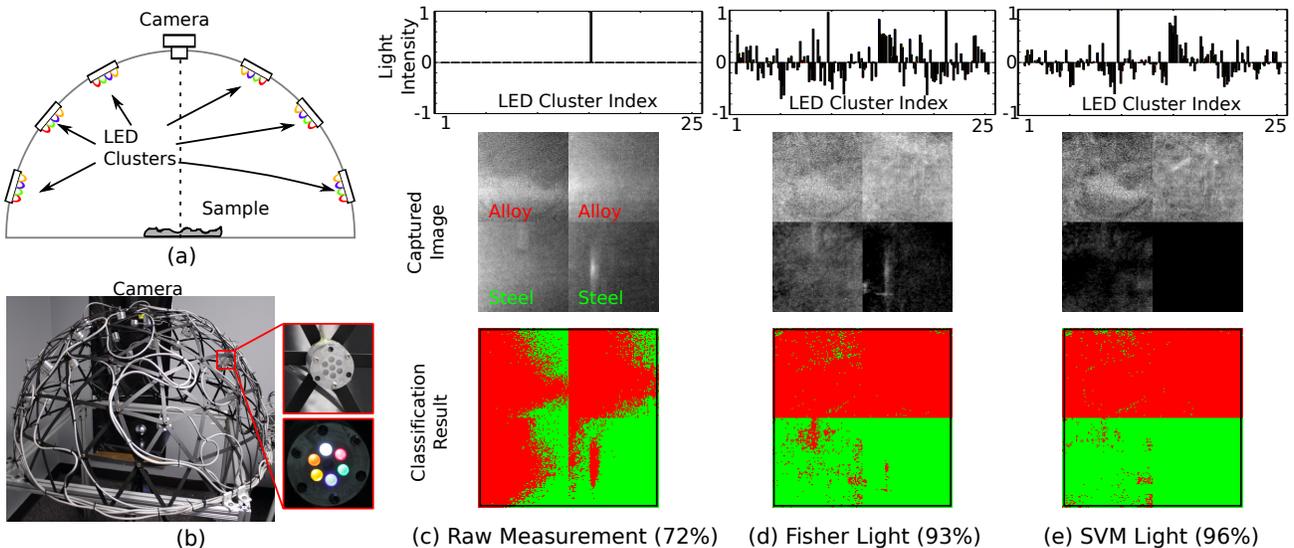


Figure 2. Discriminative illumination for material classification. (a) and (b): we design and build a LED-based multi-spectral dome light for classifying raw materials based on a 2D slice of their spectral BRDFs. The dome has 25 LED clusters. Each LED cluster has six color LEDs which can be weighted individually to create a desired spectrum. We learn optimal illumination patterns from training samples, after projecting to which the spectral BRDFs of different materials can be maximally separated. (c)(d)(e) show an example of alloy-vs-steel classification using three methods: (c) best classification with one of the  $25 \times 6 = 150$  LEDs, (d) classification with Fisher light, and (e) classification with SVM light. **Top:** the corresponding illumination patterns. **Middle:** the captured images. **Bottom:** the classification results, shown as binary images. As shown, the classification rate is at most 72% if we use one of the 150 LEDs. With the learned discriminative illumination (Fisher light or SVM light), we can achieve much higher classification rates (93% or 96%). In addition, we show in Section 4 there is also a SNR benefit of using discriminative illumination. A real-time demo is given in the supplementary video.

as polarization for metal and plastics [29, 3], spectral reflectance for printed circuit board inspection [12] and near infrared reflectance [22] for wood and textiles, and 2D slice of BRDF [28] for paint classification. Unlike these methods, we design imaging systems that use learned coded illumination to directly measure discriminative features in captured images. Recently, Jehle et al. [13] proposed to select *subsets* of basis lights (*i.e.*, rings and sectors) specifically for steel plate classification. In contrast, our method

learns optimal *weighted combinations* of basis illumination for general raw material classification, which offers much more flexibility and SNR benefits due to light multiplexing. We also extend it for multi-class and nonlinear classifications, and use both the spectral and BRDF features.

**Computational Illumination** Our work falls in the area of computational illumination which uses coded light for efficient material and shape measurement [9, 15]. However, instead of seeking for coded light for reconstructing signals

with high SNR, our goal is to find coded light with maximum discriminative ability. This is similar to the relation between EigenFaces and FisherFaces [2].

**Task-Specific and Feature-Specific Imaging** Our work is also related to task-specific and feature-specific imaging [18, 19], in which the goal of such imaging systems is not to capture visually appealing images but 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. Discriminative Illumination: A Physically-based Classifier of Spectral BRDF

For a point on an opaque, unpainted surface, its material property can often be described with a spectral BRDF [20],  $f(\omega_i, \omega_o, \lambda)$ , which is a 5-D function describing the ratio between the incident light in the direction  $\omega_i$  and the reflected light in the direction  $\omega_o$  at the wavelength  $\lambda$ . Although in principle  $f(\omega_i, \omega_o, \lambda)$  itself can be used as a feature for material classification, measuring this 5-D function is time-consuming and error-prone (especially at grazing angles) and thus directly using it for classification is impractical. As mentioned earlier, subsets of spectral BRDF have been used for material classification [28, 13].

#### 3.1. Two-Class Classification

Our approach is to design imaging systems that *directly* measure discriminative features from spectral BRDFs for classification. Consider a canonical problem of a two-class material classification with a linear classifier,

$$\mathbf{w}^T \mathbf{x} + b = \begin{cases} \geq 0 & y \in \text{Class 1} \\ < 0 & y \in \text{Class 2} \end{cases} \quad (1)$$

where  $\mathbf{x} = [f(\omega_i, \omega_o, \lambda)]$  is a vector of the spectral BRDF of a point, and the projection vector  $\mathbf{w}$  and the threshold  $b$  consist of the linear classifier. The key operation here is the projection of the spectral BRDF  $\mathbf{x}$  to the direction  $\mathbf{w}$ .

Instead of measuring the full spectral BRDF and then performing the projection, we can use coded illumination to directly measure the projection from reflected light. Consider illuminating the sample with multiple light sources from different angles with different spectra as shown in Fig. 2(a), the measured reflected light,  $I(\omega_o)$ , is

$$I(\omega_o) = \iint_{\lambda, \omega_i} f(\omega_i, \omega_o, \lambda) L(\omega_i, \lambda) \max(0, \cos \theta_i) S(\lambda) d\omega_i d\lambda, \quad (2)$$

where  $L(\omega_i, \lambda)$  is the incident light in the direction  $\omega_i$  at the wavelength  $\lambda$ ,  $f(\omega_i, \omega_o, \lambda)$  is the spectral BRDF of the sample,  $S(\lambda)$  is the spectral sensitivity of the camera, and  $\max(0, \cos \theta_i)$  is the visibility term.

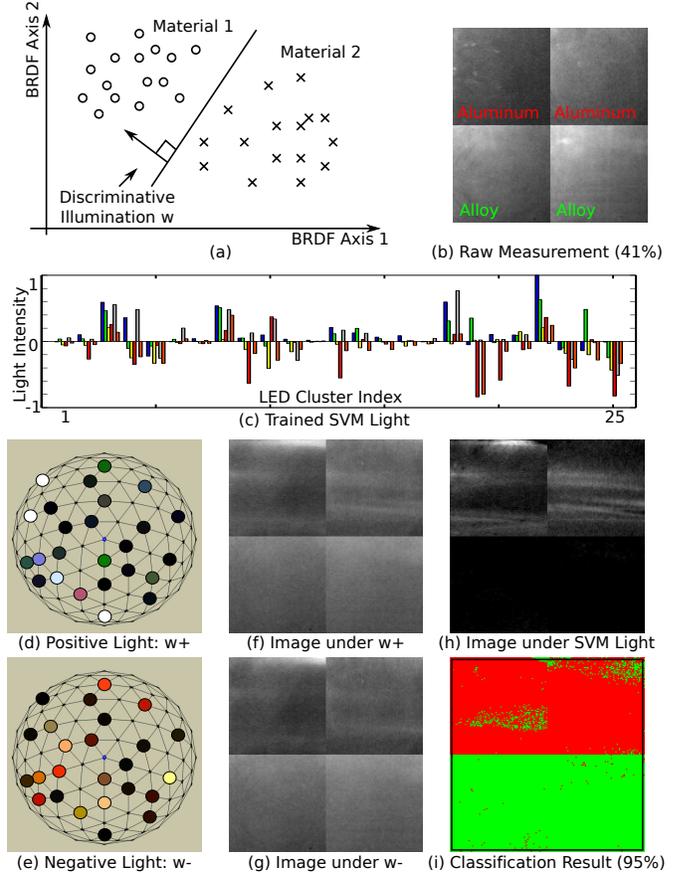


Figure 3. Discriminative illumination as a physically-based linear classifier. (a) A schematic diagram in which coded illumination acts as a linear classifier, after projecting to which the spectral BRDFs of different materials are maximally separated. (b) An example of aluminum-vs-alloy classification. The image is captured by one of the 150 LEDs of the dome which yields the best classification performance on training data. Its classification rate on testing data is 41%. (c) We train a linear kernel SVM classifier from the same training data, with the classification rate of 95% on the testing data. The bar graph shows the learned SVM light,  $\mathbf{w}$ , where the 25 bar groups correspond to the 25 LED clusters and the six bars within each group correspond to the six LEDs. The vertical axis shows the relative brightness of each LED. Since the SVM light,  $\mathbf{w}$ , has negative values, we implement it as the difference of two nonnegative vectors,  $\mathbf{w} = \mathbf{w}^+ - \mathbf{w}^-$ . (d) and (e) show the corresponding light patterns of  $\mathbf{w}^+$  and  $\mathbf{w}^-$  on the top view of the LED dome. The colors of the nodes show the spectra of the LED clusters. (f) and (g) show the corresponding captured images. (h) shows the difference of (f) and (g), which is used for classification. (i) is the classification result, shown as a binary image.

Equation (2) shows that for a given viewing direction (*i.e.*, fixed  $\omega_o$ ), a given camera (*i.e.*, fixed  $S(\lambda)$ ), and a flat sample, the measured reflected light  $I(\omega_o)$  is a *dot product* between the spectral BRDF,  $f(\omega_i, \omega_o, \lambda)$ , and the incident light,  $L(\omega_i, \lambda)$ . More explicitly, if we define  $\hat{f}(\omega_i, \omega_o, \lambda) = f(\omega_i, \omega_o, \lambda) S(\lambda) \max(0, \cos \theta_i)$  as the *spectral BRDF feature vector* for the given viewing direction and the given

camera, we have

$$I(\omega_o) = \mathbf{w}^T \mathbf{x}, \quad (3)$$

where  $\mathbf{x} = [\hat{f}(\omega_i, \omega_o, \lambda)]$ , and  $\mathbf{w} = [L(\omega_i, \lambda)]$ . Thus,  $I(\omega_o)$  directly measures the projection of spectral BRDF. This implies that we can learn the optimal projection vector  $\mathbf{w}$  from training samples for classification, and implement the projection to  $\mathbf{w}$  using coded illumination.

Figure 3(a) shows a schematic diagram of this idea, where the spectral BRDFs of samples are shown as points while the discriminative illumination  $\mathbf{w}$  is shown as a line, after projected to which samples from different classes are maximally separated. The discriminative illumination  $\mathbf{w}$  can be obtained by maximizing the discrimination of materials based on a variety of metrics, such as Fisher’s Linear Discriminant Analysis (LDA) and the Support Vector Machine (SVM) with a linear kernel [5, 4]. Since  $\mathbf{w}$  may have negative values, we implement it as the difference of two nonnegative vectors,<sup>1</sup>  $\mathbf{w} = \mathbf{w}^+ - \mathbf{w}^-$ , where  $\mathbf{w}^+ = \max(0, \mathbf{w})$ , and  $\mathbf{w}^- = -\min(0, \mathbf{w})$ .

Figure 3 shows an example of aluminum-vs-alloy classification. As shown in Fig. 3(b), these two types of metal have very similar color. If we only select two raw measurements from the 150 measurements for classification, the classification rate (on testing samples) is only 41%. Figure 3(c) shows the learned discriminative illumination  $\mathbf{w}$  using a linear kernel SVM — which we call SVM light. The vector  $\mathbf{w}$  is shown as 25 vertical bar groups. Each group has six bars corresponding to the brightness of the six LEDs within a LED cluster. Figures 3(d) and (e) show the two nonnegative light patterns,  $\mathbf{w}^+$  and  $\mathbf{w}^-$ , on the top view of the LED dome. The colors and the brightness of the nodes show the mixed spectra of the LED clusters. As shown,  $\mathbf{w}^+$  has mainly blue and white colors, while  $\mathbf{w}^-$  has mainly red and orange colors.  $\mathbf{w}^+$  has stronger incident light from grazing angles, while  $\mathbf{w}^-$  has stronger incident light from nearly the center of the dome. Figures 3(f) and (g) are the corresponding captured images, and Fig. 3(h) is the difference image of (f) and (g). Figure 3(i) shows the classification result. Using the same number of measurements, SVM light yields much higher classification rate (95%) than using raw measurement for classification (41%).

Figure 2 shows another example of alloy-vs-steel classification. In addition to SVM, we also train discriminative illumination using Fisher LDA. As shown, both the Fisher light and SVM light have higher performance (93% and 96%) than using raw measurement for classification (72%).

### 3.2. Multi-Class Classification

Multiple discriminative illumination patterns can be used for multi-class classification, as shown in Fig. 4(a). There are two common schemes to generalize binary classifiers for

<sup>1</sup>An alternative is to impose a nonnegative constraint on  $\mathbf{w}$  during optimization, with a potential decrease in its discriminative ability.

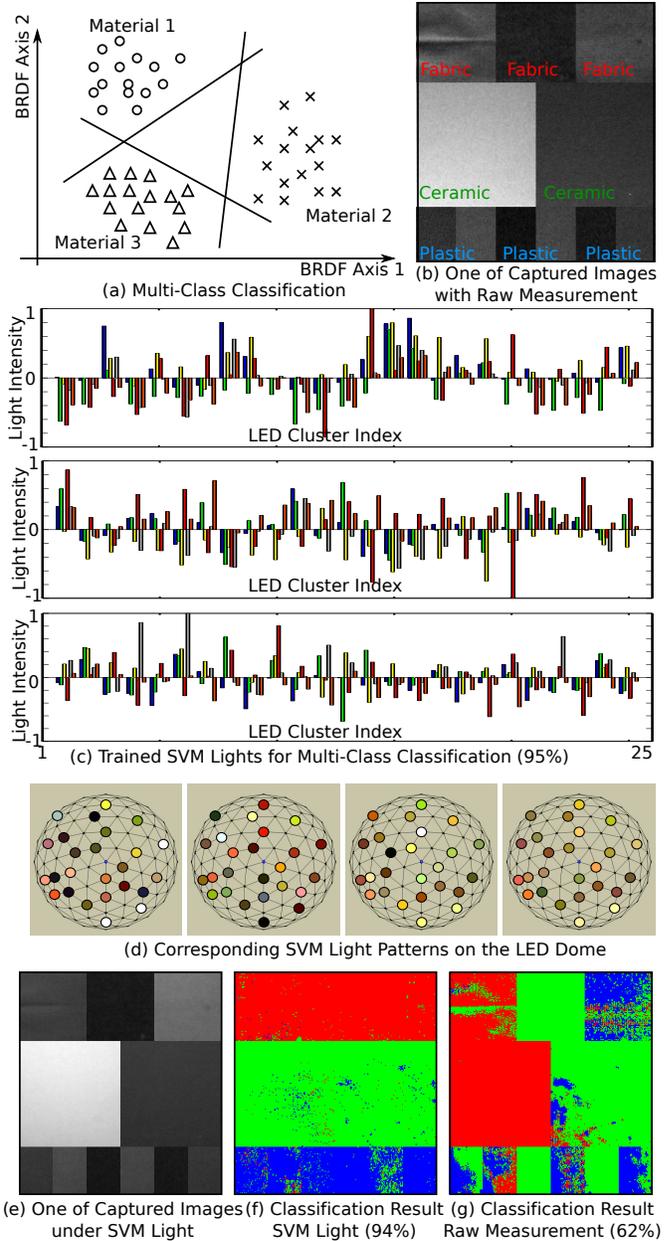


Figure 4. (a) Multiple discriminative illuminations can be used for multi-class classification tasks. We show an example of fabric-vs-ceramic-vs-plastic classification using the one-vs-all strategy. (b) the captured image under one of the 150 LEDs. (c) The learned three SVM light vectors. (d) To handle negative values, we implement the three SVM lights as four nonnegative light patterns. (e) shows one of the four captured images under the SVM lights. (f) the classification result. The classification rate is 94%. (g) In comparison, if we only select three LEDs for classification, we can at most have 62% classification rate.

multi-class tasks: one-vs-all, which has  $N$  binary classifiers for  $N$  classes, and one-vs-one, which needs  $N(N - 1)/2$  binary classifiers for  $N$  classes.

Figure 4 shows an example of three-class classifica-

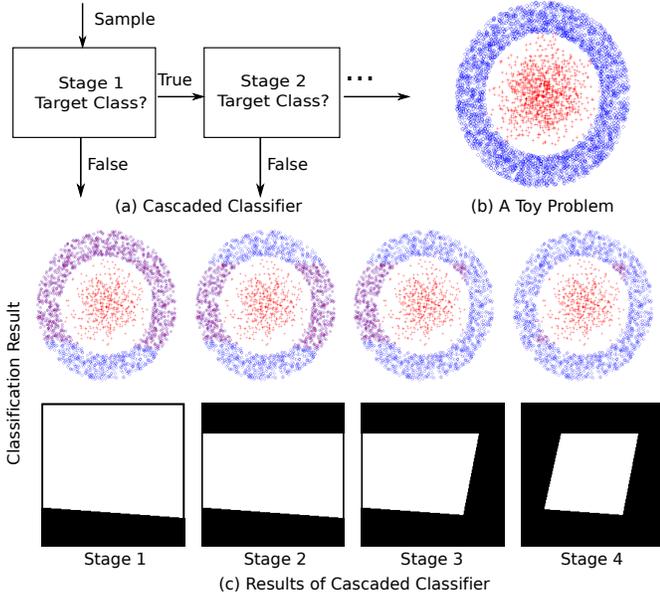


Figure 5. From linear to nonlinear classifier using multiple discriminative illuminations. (a) A cascade classifier for the detection problem, which minimizes false positive rate by adding more stages while maintaining a small given false negative rate for each stage. (b) A toy nonlinear example of detecting red + from blue circles. (c) Classification results of the cascade classifier in which each stage is a linear classifier. **Top**: the classification results on the training samples. **Bottom**: the classification boundaries.

tion, fabric-vs-ceramic-vs-plastic. Samples from these three classes have similar colors, as shown in Figure 4(a). If we select four raw measurements from the 150 measured images for classification, the best classification rate is 62%, as shown in Fig. 4(g). To use discriminative illumination for this task, we implement the one-vs-all method, and thus train three discriminative illuminations with linear kernel SVM classifiers,  $\mathbf{w}_1$ ,  $\mathbf{w}_2$ ,  $\mathbf{w}_3$ , as shown in Fig. 4(c). Since these three discriminative illuminations have negative values, we implement them as four nonnegative illumination patterns,  $\mathbf{w}^- = -\min(0, \mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3)$ , and  $\mathbf{w}_k^+ = \mathbf{w}_k - \mathbf{w}^-$ ,  $k = 1, 2, 3$ , as shown in Fig. 4(d). Figure 4(e) shows one of the four captured images under discriminative illumination, and Fig. 4(f) shows the classification result. With the same number of measurements, we achieve 94% classification rate with discriminative illumination.

### 3.3. Cascade Classifier: From Linear to Nonlinear

Multiple discriminative illumination patterns can also be constructed as an *ensemble classifier* for nonlinear classification [21, 6, 7], such as boosting, bagging, random subspace, and cascade classification.

In this section we focus on training a cascade classifier to solve the detection problem (*i.e.*, one-class classification) [14] where the goal is to distinguish one class of samples from all other possible samples. Often the number of samples of the positive class (*i.e.*, the target class) is

much smaller than that of the negative class (*i.e.*, all non-target classes). As shown in [27], subsets of negative samples and all positive samples are used to train a classifier for each stage of a cascade classifier. We adjust the threshold of each stage to meet a given false negative rate while minimizing the overall false positive rate, as shown in Fig. 5(a). Figure 5(b) shows a toy example of detecting red + (target class) from blue circles. As shown in Fig. 5(c), for this nonlinear classification task, a four-stage cascade classifier (where each stage is linear classification) is sufficient.

We use the following method to train a cascade classifier:

- **Input**: positive sample set  $P^+$ , negative sample set  $P^-$ , false negative rate  $e_0^-$ .
- **Step 0**: Initially set all negative samples as mis-classified:  $Q^- = P^-$ .
- **Step 1**: Randomly select a subset of mis-classified negative samples  $\hat{P}^-$  from  $Q^-$  and make sure  $\text{size}(\hat{P}^-) = \min(\text{size}(P^+), \text{size}(Q^-))$ .
- **Step 2**: Train a classifier based on  $P^+$  and  $\hat{P}^-$ , and adjust the threshold so that the false negative rate of this stage  $e^- \leq e_0^-$ .
- **Step 3**: Classify all negative samples in  $P^-$  using the classifiers of the current and all previous stages. Add all mis-classified negative samples to  $Q^-$ .
- **Step 4**: If the maximum stages have been trained, or  $Q^-$  is empty, break and finish the training. Otherwise, go back to **Step 1**.

For material classification, we choose the task of detecting aluminum as an example and perform classification between aluminum and three other materials (steel, ceramic, and plastic), as shown in Fig. 6. This task has practical implications in recycling since efficiently finding and recycling scrap aluminum from all other waste materials produces significant cost savings over the production of new aluminum [24, 11]. As shown in Fig. 6, we train a four-stage cascade classifier. Figure 6(a) shows one of the raw measurements (under LED #137). Figures 6(b)(c)(d)(e) show the corresponding SVM light  $\mathbf{w}$  (shown as a bar graph), the classification result (shown as a binary image), and the false negative rate and false positive rate for each stage. With four stages of discriminative illumination, we can achieve 4.2% false negative rate and 0.07% false positive rate. In comparison, if we use only a single discriminative illumination (*i.e.*, a linear classifier), with the same false negative rate (4.2%), the false positive rate will be 6%.

## 4. Signal-to-Noise Ratio Analysis

We believe the discriminative illumination has SNR benefits due to light multiplexing [23, 26]. We perform SNR analysis of the discriminative illumination method for material classification. We show that compared to using raw

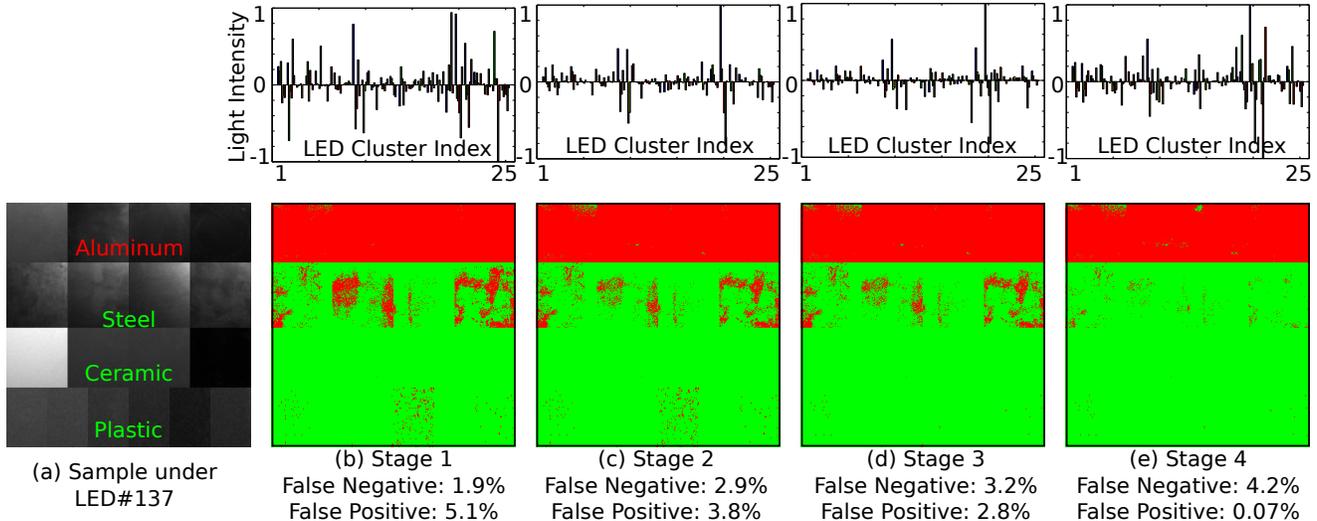


Figure 6. Aluminum detection using a cascade classifier. We train a four-stage cascade classifier to detect aluminum from three other materials (*i.e.*, steel, ceramic, plastic). (a) captured image under one of the 150 LEDs. (b)(c)(d)(e) show the learned classifier and corresponding classification result (as a binary image) for each stage. A linear-kernel SVM classifier is used for each stage. With the false negative rate for each stage to be 2%, the four-stage cascade classifier has false negative rate of 4.2% and false positive rate of 0.07%. In comparison, if we train a single linear light for this problem with the same false negative rate, the false positive rate is 6%.

measurements of spectral BRDF for classification, discriminative illumination not only reduces the number of measurements, but also have SNR benefits in the presence of either read noise or photon noise or both. More explicitly, suppose we have  $M$  raw measurements of spectral BRDF, when read noise dominates, the SNR gain of discriminative illumination is between  $\sqrt{M/2}$  and  $M/\sqrt{2}$ ; when photon noise dominates, the SNR gain is between 1 and  $\sqrt{M}$ . Detailed derivations are given in the supplementary document.

## 5. Experimental Results

As shown in Figs. 2(a)(b), we construct a LED-based multi-spectral dome for material classification. The hemispherical geodesic dome is 1 meter in diameter with 70 nodes, 25 of which are mounted with LED clusters. Each LED cluster has six color LEDs — blue, green, amber, white, red, and orange — plus a white LED in the center, as shown in Fig. 7(a). The six colors are chosen to cover the visible spectrum. UV and near infra-red LEDs can also be used for certain material classification tasks. A thin diffuser is used to uniformly mix the colors within a LED cluster. Each LED cluster is controlled by an Arduino Duemilanove (with ATmega328) board that can provide six 8-bit PWM outputs. At the center of the dome, we have a Lumenera Lu165 monochromatic CCD camera. Figure 7(b) shows a top view of the dome, with labels of the 25 LED clusters. Both the LEDs and the camera have been geometrically and radiometrically calibrated beforehand. We also perform flat fielding to ensure the uniformity of incident illumination on the sample mounted at the center of the dome.

We use the multi-spectral dome as a prototype to implement discriminative illumination for classifying a variety of

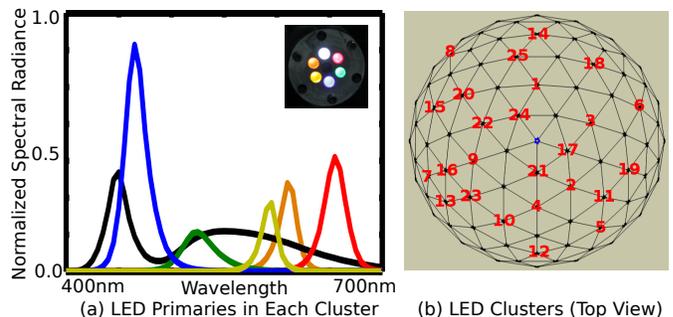


Figure 7. Details of the multi-spectral dome light. (a) We choose the spectra of the six LEDs within each cluster to cover the visible spectrum. (b) The 25 LED clusters on the top-view of the dome.

raw, unpainted materials. As shown in Fig. 1, we collect a database of five material classes: metal, plastic, ceramic, fabric, and wood. For metal, we also have six sub-classes: alloy (#4130 steel), aluminum, steel (hot roll and cold roll), stainless steel, brass, and copper. In total, there are 100 samples. Each sample is a nearly flat plate of size  $4 \times 4$  inches. We use the dome and measure the spatially-varying spectral BRDF data (*i.e.*, 150 images) for each sample plate. This database is available at [www.cis.rit.edu/mcsl](http://www.cis.rit.edu/mcsl).

We first run a series of simulations in order to understand (1) the relative contributions of the spectral and the angular information of reflectance for material classification, and (2) the optimal numbers of LED primaries and LED clusters. Detailed results are referred to the supplementary document. Our conclusions are (1) both the spectral and angular information are important (and indeed complementary), and (2) the six LED primaries and 25 clusters are sufficient for classifying the raw materials in our database.

Table 1. Comparison of classification rates for several raw material classification tasks. In each cell, the top is the classification rate of the training data, and the bottom is that of the testing data.

Task	Fisher Light	SVM Light	Raw Measurement
Aluminum_v _Steel	0.932 0.899	0.979 0.973	0.714 0.582
Brass_v _Copper	0.986 0.985	0.991 0.990	0.712 0.712
Ceramic_v _Plastic	0.945 0.944	0.956 0.955	0.674 0.678
Aluminum_v _Alloy_v _Stainless_v _Steel	0.764 0.746	0.813 0.823	0.452 0.430
Fabric_v _Ceramic_v _Plastic_v _Wood	0.913 0.923	0.947 0.942	0.522 0.525
Brass_v _Fabric_v _Ceramic_v _Plastic_v _Wood	0.885 0.891	0.926 0.924	0.474 0.482

We randomly choose half of the sample plates from each material category as the training data, and use the other half as the testing data. As shown earlier, Fig. 2, Fig. 3, Fig. 4, and Fig. 6 show the classification results for several tasks. Table 1 summarizes classification rates for several other tasks. We evaluate a variety of linear classifiers and find in general Fisher light (*i.e.*, LDA) and SVM light (*i.e.*, SVM with a linear kernel) have better performance. For comparison, we also show the classification results of using the same number of raw measurements. For example, for a three-class classification tasks, the discriminative illumination method (*i.e.*, Fisher light and SVM light) needs four images, and for using raw measurements we select the best four from the 150 images for classification. The experimental results show that the discriminative illumination method in general has higher performance.

We also try discriminative illumination for a challenging but desirable task for recycling aluminum scrap — classifying aluminum by alloy family. Depending on the alloying elements, aluminum alloys include 2000 series (alloyed with copper), 5000 series (alloyed with magnesium), 6000 series (alloyed with magnesium and silicon), 7000 series (alloyed with zinc), *etc.*. Current approaches are mainly based on laser-induced breakdown spectroscopy [8], which is expensive. Here, we use the discriminative illumination to classify four types of aluminum alloys, *i.e.*, #2024, #5052, #6061, and #7075. As shown in Fig. 8, for this challenging task, the discriminative illumination yields reasonably good results (for Fisher light, 71% accuracy and for SVM light, 73% accuracy). In comparison, using raw measurements we can only achieve 37% classification rate. This example shows the effectiveness of the proposed method for raw material classification.

## 6. Limitations and Conclusion

We propose a novel approach of using coded illumination for classifying raw materials based on projections of spectral BRDFs. Optimal illumination patterns are learned from training samples, which directly measure discriminative features for classification. This approach is more efficient than using raw measurements and also has high SNR due to illumination multiplexing. We construct a LED-based multi-spectral dome as a prototype to implement this approach for classifying a variety of raw materials. We show even for some challenging tasks, the discriminative illumination approach can achieve high classification rates.

There are several limitations in our current approach that we plan to address in our future work.

**Surface Normal** So far, we assume flat samples. To apply discriminative illumination for samples with unknown surface normals, one possible solution is to augment the training data set with variants of spectral BRDF feature vectors by tilting flat sample plates at different angles. More details can be found in the supplementary material.

**Texture** In addition to color and BRDF, another important appearance feature is texture. In this work we focus on per-pixel material classification and use only spectral BRDFs. In the future, we plan to look into texture features for material classification.

**Inter-reflection and Subsurface Scattering** Global illumination such as inter-reflection and subsurface scattering can be separated from direct illumination by modulating the discriminative light patterns with high frequency patterns [17]. These global components capture surface roughness (in the case of inter-reflection) and translucency (in the case of subsurface scattering), and can also be used as additional features for material classification.

**Going Beyond the Outer Layer of Materials** Using surface spectral BRDFs, we can only classify raw, unpainted materials. For painted or coated materials, we need to extract features that goes beyond the outer layer, *e.g.*, by using X-rays, or rely on other sensory inputs such as sound, density, hardness, smell, *etc.*. Nevertheless, the proposed methodology of finding optimal projection vectors via training while implementing the projection with computational illumination or imaging may still apply.

## Acknowledgement

We thank Gabrielle Gausted, Oliver Cossairt, and Shree K. Nayar as well as anonymous reviewers for discussions and comments. This work is supported by the RIT VPR Office and the Center for Imaging Science at RIT.

## References

- [1] E. H. Adelson and J. R. Bergen. The plenoptic function and the elements of early vision. *Computational Models of Visual Processing*, pages 3–20, 1991. 1

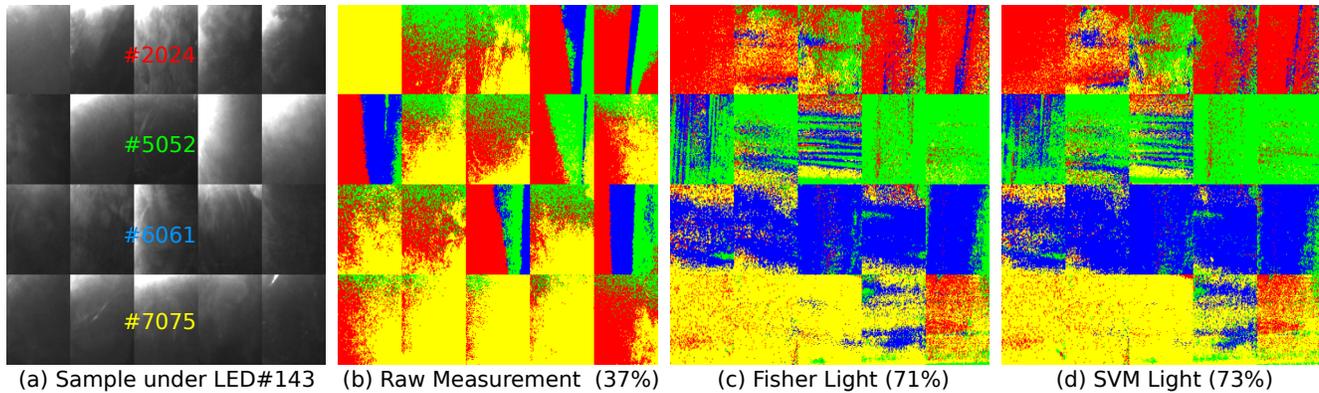


Figure 8. Classification of four types of aluminum (#2024, #5052, #6061, and #7075) using discriminative illumination. (a) We have five samples for each type, and in total 20 samples. These samples have similar color under conventional lighting conditions. The right shows the classification results of the three methods: (b) the best 5 raw measurements selected from the 150 LEDs, with classification rate of 37%, (c) Fisher light with classification rate of 71%, and (d) SVM light with classification rate of 73%. As shown, even for this challenging classification task, the proposed discriminative illumination can effectively improve performance.

- [2] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman. Eigenfaces vs. Fisherfaces: recognition using class specific linear projection. *IEEE TPAMI*, 19(7):711–720, 1997. 3
- [3] H. Chen and L. B. Wolff. A polarization phase-based method for material classification in computer vision. *IJCV*, 28(1): 73–83, 1998. 2
- [4] R. Duda, P. Hart, and D. G. Stork. *Pattern Classification*. Imperial College Press, 2001. 4
- [5] R. A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7:179–188, 1936. 4
- [6] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting, 1995. 5
- [7] J. Gama and P. Brazdil. Cascade generalization. *Machine Learning*, 41(3):315–343, 2000. 5
- [8] A. Gesing. Assuring Continued Recyclability of Automotive Aluminum Alloys: Chemical-Composition-Based Sorting of Wrought and Cast Al Shred. In *TMS*, 2003. 7
- [9] A. Ghosh, S. Achutha, W. Heidrich, and M. O’Toole. BRDF acquisition with basis illumination. In *ICCV*, 2007. 2
- [10] J. Hendrik, K. A. Massey, E. Whitham, B. Bras, and M. D. Russell. Technologies for the identification, separation, and recycling of automotive plastics. *International Journal of Environmentally Conscious Design and Manufacturing*, 6(2):37–50, 1997. 1
- [11] J. Y. Hwang, X. Huang, and Z. Xu. Recovery of metals from aluminum dross and salt cake. *Journal of Minerals and Materials Characterization and Engineering*, 5:47–62, 2006. 5
- [12] 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
- [13] M. Jehle, C. Sommer, and B. Jähne. Learning of optimal illumination for material classification. In *Pattern Recognition*, vol. 6376, pages 563–572, 2010. 1, 2, 3
- [14] S. S. Khan and M. G. Madden. A survey of recent trends in one class classification. In *AICS*, pages 188–197, 2009. 5
- [15] W.-C. Ma, T. Hawkins, P. Peers, C.-F. Chabert, M. Weiss, and P. Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *EGSR*, 2007. 2
- [16] F. Melgani and L. Bruzzone. Classification of hyperspectral remote sensing images with support vector machines. *IEEE Transactions on Geoscience and Remote Sensing*, 42(8):1778–1790, 2004. 1
- [17] S. K. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *SIGGRAPH*, 25(3):935–944, 2006. 7
- [18] M. A. Neifeld and P. Shankar. Feature-specific imaging. *Applied Optics*, 42(17):3379–3389, 2003. 3
- [19] 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. 3
- [20] F. E. Nicodemus, J. C. Richmond, I. W. Ginsberg, and T. Limperis. Geometrical considerations and nomenclature for reflectance. *National Bureau of Standards Monograph*, 160:201–231, October 1977. ISSN 1041-1135. 3
- [21] R. Polikar. Ensemble based systems in decision making. *IEEE Circuits and Devices Magazine*, 6(3):21–45, 2006. 5
- [22] N. Salamati, C. Fredembach, and S. Susstrunk. Material classification using color and NIR images. In *Color Imaging Conference*, 2009. 2
- [23] Y. Y. Schechner, S. K. Nayar, and P. N. Belhumeur. A theory of multiplexed illumination. In *ICCV*, 2003. 5
- [24] M. E. Schlesinger, O. Ilegbusi, M. Iguchi, and W. Wahn-siedler. *Aluminum Recycling*. CRC Press, 2000. 5
- [25] D.-W. Sun. *Computer Vision Technology for Food Quality Evaluation*. Academic Press, 2007. 1
- [26] T. Treibitz and Y. Y. Schechner. Recovery limits in pointwise degradation. In *ICCP*, 2009. 5
- [27] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *CVPR*, 2001. 5
- [28] O. Wang, P. Gunawardane, S. Scher, and J. Davis. Material classification using BRDF slices. In *CVPR*, pages 2805–2811, 2009. 1, 2, 3
- [29] L. B. Wolff. Polarization-based material classification from specular reflection. *IEEE TPAMI*, 12(11):1059–1071, 1990. 2