

# Coded Illumination and Imaging for Fluorescence Based Classification

Yuta Asano<sup>1</sup>, Misaki Meguro<sup>1</sup>, Chao Wang<sup>2</sup>, Antony Lam<sup>3</sup>, Yinqiang Zheng<sup>4</sup>, Takahiro Okabe<sup>2</sup>, and Imari Sato<sup>1,4</sup>

<sup>1</sup> Tokyo Institute of Technology  
{asano.y.ac, meguro.m.ab}@m.titech.ac.jp

<sup>2</sup> Kyushu Institute of Technology  
{c.wang@pluto., okabe@}ai.kyutech.ac.jp

<sup>3</sup> Saitama University  
antonylam@cv.ics.saitama-u.ac.jp

<sup>4</sup> National Institute of Informatics  
{yqzheng, imarik}@nii.ac.jp

**Abstract.** The quick detection of specific substances in objects such as produce items via non-destructive visual cues is vital to ensuring the quality and safety of consumer products. At the same time, it is well-known that the fluorescence excitation-emission characteristics of many organic objects can serve as a kind of “fingerprint” for detecting the presence of specific substances in classification tasks such as determining if something is safe to consume. However, conventional capture of the fluorescence excitation-emission matrix can take on the order of minutes and can only be done for point measurements. In this paper, we propose a coded illumination approach whereby light spectra are learned such that key visual fluorescent features can be easily seen for material classification. We show that under a single coded illuminant, we can capture one RGB image and perform pixel-level classifications of materials at high accuracy. This is demonstrated through effective classification of different types of honey and alcohol using real images.

**Keywords:** Fluorescence · Coded Illumination · Classification

## 1 Introduction

The detection of specific substances in objects such as produce items via non-destructive visual cues is vital to applications for ensuring the quality and safety of consumer products. For example, in a factory setting, we may need to evaluate the quality of food products and whether they have been contaminated with harmful bacteria and substances. A promising approach is to use coded illumination, in which controlled, active lighting makes the distinctive features of different materials visually apparent. In fact, a number of coded illumination approaches for material classification have been proposed [1–4].




Fig. 1: In many cases, visually distinguishing something like different types of honey is difficult. In the top portion of the image, the left-most two vials contain honey made from acacia flowers. The right-most two vials contain Canadian clover honey. On the bottom, we have illuminated the samples with learned illuminants that make the fluorescent emissions of the substances in the honey show visually distinct appearances.

These aforementioned approaches are all promising but they do not consider fluorescent effects, which have been shown to be especially effective in the analysis of organic substances. In short, fluorescence is a process by which an incident wavelength of light excites a substance and causes it to emit light of typically longer wavelengths. Thus for a given substance, if we were to excite it with the right kind of incident light, we would clearly see its distinctive features (Fig. 1). Indeed, the distinctive excitation and emission characteristics from the fluorescent component of various materials have been used for effective detection of substances and classification tasks. For example, Sugiyama et al. [5] showed that the fluorescence excitation-emission matrix can be used as a kind of “fluorescence fingerprint” for detecting the presence of Mycotoxin in wheat (known to cause vomiting, diarrhea, and headaches) and aerobic bacteria on beef. Fluorescence has also been used to identify cheeses [6] and wines [7], differentiate between fresh or aged fish [8], determine the botanical origin of different types of honey [9], and more.

However, conventional fluorescence-based analysis setups can only make point measurements of the target object and are often slow. For example, [5] indicates that capture of the excitation-emission matrix for a single point takes on the or-

der of minutes. On the other hand, a number of techniques for capturing the reflective and fluorescent spectral components of entire scenes have been proposed [10–13] but these either require multiple images or at least one hyperspectral image, which limits their applicability in machine vision applications.

In this paper, we propose directly learning optimal coded illuminants and weightings of the RGB channels in a camera to make fluorescent features for classification visually apparent in images. We explicitly model reflective and fluorescent effects and cast our formulation into an SVM framework [14] to jointly learn the illuminants and RGB channel weights in an alternating optimization scheme. We show that our final system is able to perform single-shot, pixel-level classification of organic materials, so our system is suited to fast quality control applications in settings such as factories. We demonstrate real sample applications in the classification of different types of honey and alcohol. To our knowledge, ours is the first approach for coded illumination-based classification using fluorescence.

## 2 Related Work

### 2.1 Material Classification Using Coded Illumination

The use of coded illumination to highlight discriminative features of material surfaces has shown great promise for machine vision classification applications. In their early work, Gu and Liu [1] proposed a per-pixel material classification approach using spectral bidirectional reflectance distribution functions (BRDFs). In their setup, they used formulations such as SVM or Fisher LDA to optimize the intensities of multispectral and multidirectional light sources for binary classification. They showed effective classification but their setup required capturing two grayscale images because they needed to simulate negative intensities via image subtraction. They also showed multiclass classification was possible by solving a set of one-versus-one classification problems but this required  $K(K-1)/2+1$  grayscale images for  $K$  classes. Later, Liu and Gu [2] extended their work to use RGB images. Using the same lighting setup but with a three-channel camera, they then used the binary or multiclass Fisher LDA formulations to find the 3-D feature space that maximizes the ratio of the between-class to within-class scattering. However, they still needed to capture two RGB images to simulate negative intensities via image subtraction.

In Wang and Okabe [3], they proposed a coded illumination approach that would only require a single image for per-pixel material classification. This provided a great advantage because single-shot systems are well suited to situations where the objects are in motion. In a factory setting, one may expect objects to be moving along quickly on a conveyer belt. The single-shot capability of their system was made possible by enforcing non-negative constraints on the learned coded illuminants so that a second image for simulating negative intensities would not be needed. They also showed that it was possible to capture a scene using one fixed set of coded illuminants and an RGB camera but in post-processing, achieve multiclass classification. This was made possible by jointly

learning a single set of non-negative coded illuminants with multiple postprocessing grayscale conversions of the RGB image. The multiple grayscale images generated from a single captured RGB image would then highlight features effective for multiple binary classification decisions.

In Blasinski et al. [4], they also proposed a non-negative coded illumination approach to material classification. Specifically, they learned multiple illuminant spectra based on an SVM formulation or non-negative PCA. They then captured scenes using RGB camera spectral responses under the illuminants and show effective per-pixel classification in test scenes with different fruits. In general, they reported that about 3–4 illuminants gave good performance with only modest gains if more coded illuminants were added. Their paper differs from the previously mentioned papers in that they do not use multidirectional light but rather vary the illuminants primarily in the spectral domain.

## 2.2 Fluorescence for Classification and Detection of Substances

The previously mentioned coded illumination approaches showed promising results. However, they all assumed scenes to be purely reflective and did not consider fluorescent effects. We now briefly describe the difference between reflectance and fluorescence. In summary, reflectance is when both the incident and reflected light from the material are of the same wavelength. On the other hand, fluorescence is when light of a typically shorter wavelength “excites” a substance and then typically longer wavelengths of light are “emitted”.

It is well-known that fluorescence can reveal a lot about the state of objects. In particular, organic objects exhibit distinctive fluorescent characteristics based on what kinds of substances and/or bacteria are present. For example, Sugiyama et al. [5] used a fluorescence spectrometer to make point measurements to determine the spectral excitation-emission matrix of different organic objects. They showed that the excitation-emission matrix could be treated as a kind of “fluorescence fingerprint” to identify the presence of Mycotoxin in wheat (known to cause vomiting, diarrhea, and headaches). They were also able to detect aerobic bacteria on beef. As mentioned earlier, fluorescence has also been used for varied tasks such as identifying different types of cheeses, wines, honey [6, 7, 9] and even to tell the difference between fresh and aged fish [8]. It is well-known that observing fluorescence is an effective means of analyzing various materials but conventional measurements such as that of [5] do not capture the entire scene and take on the order of minutes to capture the entire excitation-emission matrix. This precludes applications to settings such as quality control in a factory where numerous products could be moving quickly along a conveyer belt.

## 2.3 Fluorescence Imaging and Classification

In recent years, there have been a number of proposed techniques for capturing fluorescence spectral components for entire scenes. Lam and Sato [10] proposed using a sparse set of narrowband illuminants and images combined with basis vectors to estimate the fluorescence spectral components. Fu et al. [11] proposed

capturing hyperspectral images under two high frequency light spectra to estimate the fluorescence spectral components. Later, Fu et al. [12] estimated the components using an RGB camera and multiple active illuminants. Zheng et al. [13] devised a means to estimate all the fluorescence spectral components using only a single hyperspectral image. All the aforementioned approaches require either multiple images or at least one hyperspectral image, which limits applications to static scenes.

In this paper, we propose directly learning coded illuminants for material classification tasks. We explicitly model fluorescence and derive a formulation that can be cast into an SVM learning framework. In doing so, we create illuminants that excite the fluorescent components of specific substances such that their distinctive features are easily seen under an RGB camera. Furthermore, our proposed system only requires a single image and so is applicable to scenes with moving objects (such as in a factory with conveyer belts). We demonstrate our system with real applications in classifying different types of honey and alcohol. In summary, our contributions are as follows:

1. We explicitly model the images of reflective-fluorescent materials under an RGB camera and show that this formulation can be cast into an SVM learning framework for optimizing coded illuminants.
2. We demonstrate that the resultant coded illuminants can make it so that visual features from the fluorescent components of substances are easily seen.
3. We provide a comparison between coded illuminants and standard illuminants in classification tasks to demonstrate the benefits of our proposed approach.
4. To our knowledge, we are the first to propose coded illuminants that leverage fluorescence for classification tasks—despite the well-known observation that fluorescence provides highly distinctive cues for detecting the presence of substances.

### 3 Coded Spectral Response and Illumination for Fluorescence-based Classification

#### 3.1 Imaging Model

Most fluorescent materials actually have a combination of reflectance and fluorescence. So we start with presenting a model for how reflective-fluorescent materials are observed under a given illumination spectrum for a single channel camera. It is well-known that the image of any given reflective-fluorescent material is a linear combination of the reflected incident light and the emitted light from the fluorescent component. This emitted light is typically shifted to longer wavelengths than the incident light. Thus for a given camera, the outgoing wavelength  $\lambda_o$  for a reflective-fluorescent material illuminated by incident light at wavelength  $\lambda_i$  can be modeled as

$$P(\lambda_o, \lambda_i) = R(\lambda_o)L(\lambda_i)\delta(\lambda_o - \lambda_i)C(\lambda_o) + E_m(\lambda_o)C(\lambda_o)E_x(\lambda_i)L(\lambda_i) \quad (1)$$

where  $R(\lambda_o)$  is the reflectance at wavelength  $\lambda_o$ ,  $L(\lambda_i)$  is the illuminant at wavelength  $\lambda_i$ , and  $C(\lambda_o)$  is the camera spectral response at wavelength  $\lambda_o$ .  $E_m(\lambda_o)$  and  $E_x(\lambda_i)$  are the emission and excitation of the fluorescent component at their respective wavelengths. The excitation term  $E_x(\lambda_i)$  determines how much the energy from incident light at wavelength  $\lambda_i$  is able to excite the fluorescent component. On the other hand, the emission term  $E_m(\lambda_o)$  determines how much light at wavelength  $\lambda_o$ , the fluorescent component is able to emit relative to the amount of energy from the excitation.  $\delta(\lambda_o - \lambda_i)$  is the unit impulse function where  $\delta(0) = 1$  and  $\delta(x) = 0$  for  $x \neq 0$ . The unit impulse function ensures only the incident wavelength is reflected for the reflective component.

Then to determine the image of the material under wideband light for a wideband camera, we can simply sum over all the possible combinations of wavelengths  $\lambda_o$  and  $\lambda_i$ :

$$I = \iint P(\lambda_o, \lambda_i) d\lambda_o d\lambda_i \approx \sum_{m=1}^M \sum_{x=1}^X P(\lambda_o^{(m)}, \lambda_i^{(x)}) \Delta\lambda_o \Delta\lambda_i. \quad (2)$$

The right hand side of Eq. 2 is the discrete approximation that is used in practice. In our setup, we calculate the  $P$  term at intervals of 10 nm for both wavelength parameters.

## 4 Learning the Coded Illumination

We now describe how the imaging model in Eq. 2 can be used to formulate a framework for learning an illuminant spectrum and weighting for the RGB channels so that distinctive fluorescent features for classification are easily seen. We could then perform pixel-wise classification of the types of materials present with just a single image.

For convenience, let  $T(\lambda_o, \lambda_i) = R(\lambda_o)L(\lambda_i)\delta(\lambda_o - \lambda_i) + E_m(\lambda_o)E_x(\lambda_i)$ . Then Eq. 2 can be written in matrix form as

$$\begin{aligned} I &= \left( C(\lambda_o^{(1)}) \dots C(\lambda_o^{(M)}) \right) \begin{pmatrix} T(\lambda_o^{(1)}, \lambda_i^{(1)}) & \dots & T(\lambda_o^{(1)}, \lambda_i^{(X)}) \\ \vdots & & \vdots \\ T(\lambda_o^{(M)}, \lambda_i^{(1)}) & \dots & T(\lambda_o^{(M)}, \lambda_i^{(X)}) \end{pmatrix} \begin{pmatrix} L(\lambda_i^{(1)}) \\ \vdots \\ L(\lambda_i^{(X)}) \end{pmatrix} \\ &= \mathbf{c}^T \mathbf{T} \mathbf{l} = \mathbf{f}^T \mathbf{l}. \end{aligned} \quad (3)$$

Note that  $\mathbf{T}$  is basically the fluorescence excitation-emission matrix of the material (but with reflectance terms added in),  $\mathbf{c}$  is the vector representing the camera spectral response, and  $\mathbf{l}$  is the vector representing the illuminant spectrum. We define vector  $\mathbf{f}$  as the reflective-fluorescent feature of the given material under camera spectral response  $\mathbf{c}$ . Thus for a given camera spectral response and material's  $\mathbf{T}$  matrix, the image of the material under illuminant  $\mathbf{l}$  is the inner-product between reflective-fluorescent feature  $\mathbf{f}$  and illuminant  $\mathbf{l}$ .

In the case of an RGB camera, we have three channels. So for a single illuminant, the reflective-fluorescent material's image would consist of three values computed as

$$\begin{pmatrix} I_r \\ I_g \\ I_b \end{pmatrix} = \begin{pmatrix} \mathbf{f}_r^T \\ \mathbf{f}_g^T \\ \mathbf{f}_b^T \end{pmatrix} \mathbf{l}. \quad (4)$$

For discussion purposes, we also define the weighting of the RGB values:

$$I = (w_r \ w_g \ w_b) \begin{pmatrix} I_r \\ I_g \\ I_b \end{pmatrix} = w_r \mathbf{l}^T \mathbf{f}_r + w_g \mathbf{l}^T \mathbf{f}_g + w_b \mathbf{l}^T \mathbf{f}_b \quad (5)$$

where  $w_r$ ,  $w_g$ , and  $w_b$  are weights in the summation of the image of the materials under each RGB channel.

From Eq. 5, we can see that for an RGB camera, the combination of illuminant spectrum  $\mathbf{l}$  and RGB weighting values  $w_r$ ,  $w_g$ , and  $w_b$  constitute a linear discriminant hyperplane of the form

$$I + b = w_r \mathbf{l}^T \mathbf{f}_r + w_g \mathbf{l}^T \mathbf{f}_g + w_b \mathbf{l}^T \mathbf{f}_b + b = 0 \quad (6)$$

where  $b$  is a bias term. Then given a set of features  $\mathbf{f}_r$ ,  $\mathbf{f}_g$ ,  $\mathbf{f}_b$ , and class labels  $y \in \{1, -1\}$ , we might try to learn an appropriate hyperplane using a soft-margin SVM [14]. This is similar to previous work that used soft-margin SVM optimization inspired approaches to learn coded illuminants [1–4]. However, past approaches have only considered the reflectance of incident light but not fluorescence excitation-emissions as we do here. Going back to our discussion, Eq. 6 shows that we have unknown illuminant spectrum  $\mathbf{l}$  and unknown RGB weighting values  $w_r$ ,  $w_g$ , and  $w_b$ . In addition, the first three terms in the summation are all dependent on illuminant spectrum  $\mathbf{l}$ . Thus the standard SVM soft-margin optimization procedure cannot be used. Fortunately, we have found that although Wang and Okabe [3] worked in the domain of reflectance BRDFs and did not optimize light spectra, their reformulated SVM soft-margin optimization can be used in the spectral domain for learning our proposed fluorescence-based coded illuminants. For clarity, we present the optimization formulation with our fluorescence terms integrated here:

$$\begin{aligned} \min_{\mathbf{l}, w_r, w_g, w_b, b, \xi_n} \quad & \frac{1}{2} |\mathbf{l}|^2 (w_r^2 + w_g^2 + w_b^2) + \beta \sum_{n=1}^N \xi_n \\ \text{s.t.} \quad & y_n [\mathbf{l}^T (w_r \mathbf{f}_{nr} + w_g \mathbf{f}_{ng} + w_b \mathbf{f}_{nb}) + b] \geq 1 - \xi_n \quad (n = 1, 2, \dots, N), \\ & \xi_n \geq 0 \quad (n = 1, 2, \dots, N), \\ & l_k \geq 0 \quad (k = 1, 2, \dots, K). \end{aligned} \quad (7)$$

$N$  is the number of training samples,  $\mathbf{f}_{nm}$  denotes the  $n^{\text{th}}$  reflective-fluorescent training sample for camera color channel  $m$ .  $\xi_n$  is the slack variable and  $\beta$  is the weight penalty term. In our setup, we use coded illuminants ranging from 350 nm - 640 nm in increments of 10 nm so  $K = 30$ .




Fig. 2: Consumer products collected for our dataset. Various types of honey (top) and alcohol (bottom) spanning different brands are shown.

The above formulation has an unknown illuminant spectrum  $\mathbf{l}$  and unknown set of RGB weighting values  $w_r, w_g, w_b$ . In our setup, we solve for the RGB weights and then the illuminant spectrum and bias using alternating iterations of quadratic programming. Specifically, we initialize the illuminant spectrum  $\mathbf{l} = (1 \ 1 \ \dots \ 1)^T$  and bias  $b = 1$  and solve for the RGB weighting values. Then the RGB weights are fixed and we solve for the illuminant spectrum and bias. The iterations are repeated until convergence or a preset maximum number of iterations is reached.

## 5 Experiments

### 5.1 Data Collection

We built a dataset consisting of various types of honey and alcohol (Fig. 2). Specifically, we obtained acacia honey (4 brands), Canadian clover honey (4 brands), orange honey (3 brands), whisky (3 brands), scotch (3 brands), bourbon (3 brands), brandy (2 brands), and cognac (2 brands). For each product, we used a fluorescence spectrometer to capture  $20 \times 20$  hyperspectral images of the given sample at multiple narrowbands ranging from 350 nm - 640 nm in increments of 10 nm. The narrowband lights were all normalized in postprocessing so that they would have equal intensity.

### 5.2 Experiment Setup and Classification Tasks

**System Setup:** Our proposed system consists of an RGB camera and coded illuminant spectrum as the light source. We use a PointGrey GS3-U3-23S6 camera with color filters as our RGB camera. Note that this camera has a linear



response function with manual settings for gamma correction and white balance. Thus our setup assumes a linear response function. However, we can still use an sRGB camera with a non-linear response function by first obtaining its response function in advance and then converting it to give a linear response image. For the first part of our classification experiments, we take the 20 x 20 hyperspectral images from our dataset and simulate the image of each sample under our RGB camera’s known set of RGB spectral response functions and coded illuminants (ranging from 350 nm - 640 nm in increments of 10 nm). In the next phase of our tests, we demonstrate an implementation of our system using our PointGrey GS3-U3-23S6 RGB camera and a Nikon ELS programmable light source for generating coded illuminant spectra. For a given coded illuminant we can then capture an RGB image and classify each pixel using the discriminant hyperplane defined in Eq. 6.

**Classification Tasks:** For the classification tasks, our aim is to differentiate between different types of honey and alcohol in a one-versus-one manner. As an example, consider the problem of classifying acacia honey versus Canadian clover honey. In this case, we use two samples (one sample from each type of honey) and learn a coded illuminant spectrum to separate them. (Each sample is an image that consists of 20 x 20 pixels so this means we have 400 datapoints per sample to learn the coded illuminant.) The coded illuminant spectrum is then used as the light source for the 2D RGB images of all instances of acacia honey and Canadian clover honey in our dataset that were not used in the training data. We note that all these different instances of honey come from different brands.

**Training-Testing Splits:** In our example on acacia versus Canadian clover honey, we described a single training-testing split in our classification tests. To thoroughly test the classification of acacia honey versus Canadian clover honey, we exhaustively try all combinations of training-testing splits of the data in which the training set always consists of one sample of each type of honey. We then determine the average accuracy of all the pixel-level classifications on the test set (containing exclusively different brands from the training set) and report them. We also repeat the same test procedure for various combinations of one-versus-one classification problems for different types of honey as well as different types of alcohol.

**Comparisons to Non-coded Illuminants:** We also repeat our experiments with three conventional illuminants (Fig. 3). We conduct these tests by using the same formulation as Eq. 7 to learn a discriminant hyperplane but the illuminant spectrum is kept fixed. In other words, only the bias term and RGB weighting values are learned for the classification task. Then for testing, the fixed standard illuminant, learned bias term, and learned RGB weights are used for pixel-level classification as is done in the coded illumination tests.




Fig. 3: Conventional (non-coded) illuminants used for comparisons to coded illuminants for classification tasks.

## 6 Results

We report the average accuracies on classifying different types of honey based on botanical origin (which kinds of flowers they were made from) in Table 1. We see that the proposed coded illuminants can be used for effective classification. In Fig. 4, we show examples of the excitation-emission matrices of different instances of honey and their categories. We also include examples of learned coded illuminants for separating classes based on these excitation-emission characteristics. We can see that the coded illuminants will emphasize the range of wavelengths where the material exhibits high excitation and emission. As mentioned earlier, for tests with conventional illuminants, we used Eq. 7 to learn weights for the RGB channels and the bias term but kept the illuminant fixed. We found that the resultant classifiers using conventional illuminants would output the same class label for input test data in almost all the cases. Thus many of the average accuracies in Table 1 appear to be the same (e.g. 40% appears because the number of instances in the testing set consisted of 60% of one class versus 40% of the other class). On the other hand, using our proposed coded illuminant approach, we achieved effective classification of different types of honey despite only training on two samples of particular brands of honey and then testing on multiple brands. We note though, that in the case of Canadian clover honey versus orange honey, our accuracy was lower. However, our proposed approach still allowed for effective discrimination of other types of honey whereas conventional non-coded lighting could not classify any honey in most of the cases.

In Table 2, we can see classification results on various types of whisky versus brandy. Since whisky is distilled beer and brandy is distilled wine, we would expect various categories of whisky to be separable from brandy. Indeed, we can see in Table 2, that the classification accuracies using our coded illumination approach indicate we can differentiate between different types of whisky and brandy. In Table 3, we show results from tests on various types of whisky versus various types of whisky. Since they are more similar to each other than in the case of whisky versus brandy, the overall classification accuracies are lower. In Fig. 5, we show examples of the types of coded illuminants learned for classifying alcohol. There is a good amount of variety in their characteristics.

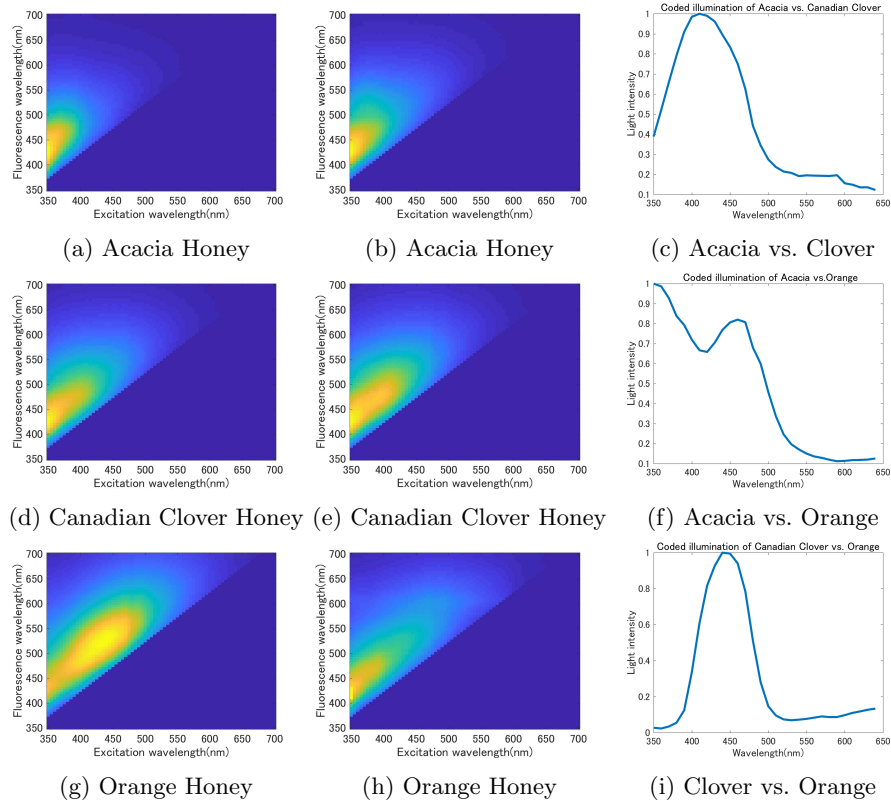


Fig. 4: Examples of Excitation-Emission Matrices and Coded Illuminants for Differentiating Between Types of Honey

Up to this point, we have presented binary classification results but our formulation also allows for multiclass classification. Typically, multiclass classification is performed by using multiple binary classifiers to decide on class membership. Our formulation actually allows for obtaining  $V$  linear discriminant hyperplanes with only a single image. This is because it is possible to take a single image under only *one coded illuminant* and then learn a set of RGB weights  $w_{vr}$ ,  $w_{vg}$ ,  $w_{vb}$ , and biases  $b_v$  for each linear discriminant hyperplane  $v$ . Thus using basically the same optimization formulation as binary classification, we start with a single fixed illuminant spectrum and biases  $b_v$ . We then iteratively update each set of RGB weights for each binary classification problem. Then the multiple sets of RGB weights are all fixed and we update the single illuminant spectrum and biases  $b_v$ . This alternating process is repeated until convergence or a preset number of iterations is reached. For final classification, these multiple hyperplanes can then be used to vote for the class labels of test cases. In Table 4, we can see results on four-class classification of alcohol.

Table 1: Average accuracies from all the training-testing splits for classifying different types of honey. For the non-coded illuminants, many of the classifications gave the same output label regardless of the test input. Thus there are many cases with 40% accuracy because the test sets in those cases had 60% of one class versus 40% of the other class.

Classification/Illuminant	5000K LED	Fluorescent Light	Incandescent	<b>Coded (Proposed)</b>
Acacia vs. Canadian Clover	50.00%	50.67%	50.67%	<b>91.17%</b>
Acacia vs. Orange	40.00%	40.00%	40.00%	<b>78.26%</b>
Canadian Clover vs. Orange	41.67%	40.00%	45.00%	<b>52.03%</b>

Table 2: Average accuracies from all the training-testing splits for classifying various types of Whisky vs. Brandy. Similar to the tests on honey, the classifiers in the case of non-coded light classified all input data as some type of whisky so there are many cases of 66.67% accuracy.

Classification/Illuminant	5000K LED	Fluorescent Light	Incandescent	<b>Coded (Proposed)</b>
Scotch vs. Cognac	66.67%	66.67%	66.67%	<b>100.00%</b>
Bourbon vs. Brandy	66.67%	66.67%	66.67%	<b>71.44%</b>
Bourbon vs. Cognac	66.67%	66.67%	66.67%	<b>95.97%</b>
Scotch vs. Brandy	66.67%	66.67%	66.67%	<b>99.43%</b>
Whisky vs. Cognac	66.67%	66.67%	66.67%	<b>99.50%</b>
Whisky vs. Brandy	66.67%	66.67%	66.67%	<b>93.00%</b>

Multiple training-testing splits were chosen in each case such that the test sets would have four test cases, one from each class. We can see that the multiclass classification accuracies using our coded illumination approach indicate we can differentiate between four different types of alcohol. For the non-coded illuminants, many of the classifications gave the same output label regardless of the test input. We found that many cases using the non-coded illuminants resulted in 25% accuracy, which is the same as random guessing.

Overall, our proposed coded illumination fluorescence-based classification approach showed significant improvement over using conventional light sources. As expected, when intraclass variation is high for both the classes in question, the accuracies are lower. Likewise, separation of different categories of items with very similar characteristics was, as expected, difficult. However, in all our tests, we only used one sample per class and tested on more instances than training data. Thus the overall good classification performance despite the difficult tests shows the effectiveness of our approach.

The experiments presented so far made use of real images but they were narrowband images that were used to simulate an RGB camera using a given set of spectral response functions. This allowed us to perform a large number of extensive tests. We now demonstrate a single-shot setup using a PointGrey GS3-U3-23S6 RGB camera and Nikon ELS programmable light source to gen-

Table 3: Average accuracies on the training-testing splits for Whisky vs. Whisky classification. In the table, “Whisky” is used to denote generic whisky, which are then classified against specific types of whisky such as bourbon or scotch.

Classification/Illuminant	5000K LED	Fluorescent Light	Incandescent	<b>Coded (Proposed)</b>
Whisky vs. Bourbon	52.58%	50.00%	44.27%	<b>61.10%</b>
Whisky vs. Scotch	50.00%	50.00%	55.00%	<b>77.30%</b>
Scotch vs. Bourbon	50.00%	50.00%	50.00%	<b>58.16%</b>

Table 4: Average accuracies from all the training-testing splits for classifying four-class types of alcohol.

Classification/Illuminant	<b>Coded (Proposed)</b>
Whisky vs. Scotch vs. Bourbon vs. Brandy	<b>55.10%</b>
Whisky vs. Scotch vs. Bourbon vs. Cognac	<b>51.04%</b>
Whisky vs. Scotch vs. Brandy vs. Cognac	<b>62.65%</b>
Whisky vs. Bourbon vs. Brandy vs. Cognac	<b>68.93%</b>
Scotch vs. Bourbon vs. Brandy vs. Cognac	<b>64.84%</b>

erate coded illuminants (Fig. 6). We chose to compare the results from our programmable light source setup to two training-testing splits from our previous tests. The results are presented in Table 5. In the table, we can see the two training-testing splits and the average pixel classification accuracies. The column denoted “ideal” shows results from our tests using the real captured narrowband images that are then used to simulate RGB images using spectral response functions. In this case, the tests show what accuracies ideally generated illuminants could yield. Not surprisingly the coded illuminant generated by the Nikon ELS results in a lower accuracy for the honey classification test. It is interesting that in some cases, such as in the bourbon versus scotch test, the results were similar between the programmable light source and ideal setups. Future work will investigate the differences between the programmable light source setup and ideal light setup.

Table 5: Comparison of ideal setup vs. programmable light source for two specific training-testing splits. “Ideal” denotes the results from our narrowband images that were used to simulate RGB images under coded illuminants.

	Ideal	Programmable Light Source
Acacia vs. Canadian Clover	92.69%	85.96%
Bourbon vs. Scotch	85.06%	85.75%




Fig. 5: Coded Illumination for Alcohol Classification




Fig. 6: Setup with Programmable Light Source

## 7 Conclusion

We have demonstrated the effectiveness of learning coded illumination to leverage the particular excitation-emission characteristics of substances in materials for classification purposes. In addition, our system only requires a single image under one illuminant and thus is applicable for use in such settings as factory quality and safety control. We also demonstrated the use of a programmable light source to show that coded illuminants can be generated in reality. There are some cases where our system could not classify well. These are likely due to a combination of high intraclass variability and low interclass difference (e.g. differentiating different kinds of whisky). In the future, we will investigate ways to capture unique excitation-emission characteristics with more detail. One possible approach is to learn coded camera spectral responses instead of just weighting the RGB channels. Building a larger dataset to obtain more training data may also allow us to build stronger classifiers.

## Acknowledgements

This work was supported in part by JSPS KAKENHI Grant Numbers JP15H05918, JP16H01676, and JP17H01766.

## References

1. Gu, J., Liu, C.: Discriminative illumination: Per-pixel classification of raw materials based on optimal projections of spectral BRDF. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Volume 00. (06 2012) 797–804
2. Liu, C., Gu, J.: Discriminative illumination: Per-pixel classification of raw materials based on optimal projections of spectral BRDF. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) **36**(1) (Jan 2014) 86–98
3. Wang, C., Okabe, T.: Joint optimization of coded illumination and grayscale conversion for one-shot raw material classification. In: British Machine Vision Conference (BMVC). (2017)
4. Blasinski, H., Farrell, J., Wandell, B.: Designing illuminant spectral power distributions for surface classification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (July 2017)
5. Sugiyama, J., Fujita, K., Yoshimura, M., Tsuta, M., Shibata, M., Kokawa, M.: Detection of food hazards using fluorescence fingerprint. IFAC Proceedings Volumes **46**(18) (2013) 70 – 74 4th IFAC Conference on Modelling and Control in Agriculture, Horticulture and Post Harvest Industry.
6. Karoui, R., Dufour, E., Schoonheydt, R., Baerdemaeker, J.D.: Characterisation of soft cheese by front face fluorescence spectroscopy coupled with chemometric tools: Effect of the manufacturing process and sampling zone. Food Chemistry **100**(2) (2007) 632–642
7. Chabreyrie, D., Chauvet, S., Guyon, F., Salagoity, M.H., Antinelli, J.F., Medina, B.: Characterization and quantification of grape variety by means of shikimic acid concentration and protein fingerprint in still white wines. Journal of Agricultural and Food Chemistry **56**(16) (2008) 6785–6790 PMID: 18624410.
8. Dufour, E., Frenca, J.P., Kane, E.: Development of a rapid method based on front-face fluorescence spectroscopy for the monitoring of fish freshness. Food Research International **36**(5) (2003) 415–423
9. Lenhardt, L., Zekovic, I., Dramicanin, T., D. Dramicanin, M., Bro, R.: Determination of the botanical origin of honey by front- face synchronous fluorescence spectroscopy. Applied Spectroscopy **68**(5) (2014) 557–563
10. Lam, A., Sato, I.: Spectral modeling and relighting of reflective-fluorescent scenes. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (June 2013) 1452–1459
11. Fu, Y., Lam, A., Sato, I., Okabe, T., Sato, Y.: Separating reflective and fluorescent components using high frequency illumination in the spectral domain. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) **38**(5) (May 2016) 965–978
12. Fu, Y., Lam, A., Sato, I., Okabe, T., Sato, Y.: Reflectance and fluorescence spectral recovery via actively lit RGB images. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) **38**(7) (2016) 1313–1326
13. Zheng, Y., Fu, Y., Lam, A., Sato, I., Sato, Y.: Separating fluorescent and reflective components by using a single hyperspectral image. In: IEEE International Conference on Computer Vision (ICCV). (Dec 2015) 3523–3531
14. Cortes, C., Vapnik, V.: Support-vector networks. Machine Learning **20**(3) (September 1995) 273–297