

16-385 Computer Vision, Spring 2019

Homework Assignment 4

Physics-based vision

Due Date: Wed March 27, 2019 23:59

Instructions

1. **Integrity and collaboration:** Students are encouraged to work in groups but each student must submit their own work. If you work as a group, include the names of your collaborators in your write up. Code should **NOT** be shared or copied. Please **DO NOT** use external code unless permitted. Plagiarism is strongly prohibited and will lead to failure of this course.
2. **Start early!** Especially if you are not familiar with MATLAB.
3. **Questions:** If you have any question, please look at piazza first. Other students may have encountered the same problem, and is solved already. If not, post your question on the discussion board. TAs will respond as soon as possible.
4. **Write-up:** Items to be included in the writeup are mentioned in each question, and summarized in the Writeup section. Please note that we **DO NOT** accept handwritten scans for your write-up in this assignment. Please type your answers to theory questions and discussions for experiments electronically.
5. **Handout:** The handout zip file contains 3 items. `assgn4.pdf` is the assignment handout. `data` contains normals for the Stanford bunny images, photometric stereo data, an image of a fruitbowl, and the spectral radiance of bananas. `matlab` contains a script that you will make use of in this homework.
6. **Submission:** Your submission for this assignment should be a zip file, `<andrew-id>.zip`, composed of your write-up, your MATLAB implementations (including helper functions), and your implementations and results for extra credit (optional). For each of the questions 1.3-1.5 (and 1.6 if you do the optional credit part), please make sure to include a Matlab script named `part_1.x.m` with your solution for that part. Each script can call other helper functions, which you should make sure to include in your submission.

1 Overview


The purpose of this homework is to go over the topics of radiometry, reflectance, photometric stereo, and color. Given the short nature of this homework, and unlike previous assignments, it consists of a few smaller and independent problems, rather than a sequence of related problems.

1.1 Lambertian albedo (10 points)

For Lambertian surfaces, the BRDF is a constant function of the input and output directions. For such a material, we often describe the reflectance in terms of its *albedo*, which is given the symbol ρ . For a Lambertian surface, the BRDF and albedo are related by $f_r(\hat{\mathbf{v}}_i, \hat{\mathbf{v}}_o) = \rho/\pi$. Using conservation of energy, prove that $0 \leq \rho \leq 1$.

1.2 Foreshortening (10 points)

A small Lambertian source dA is centered at P and emits radiance L . The orientation of this patch is the same as that of a plane containing two points, X_1 and X_2 . The point X_1 is the point on this plane that is closest to P , and the distance from P to X_1 is D as shown.



1. Calculate the solid angle subtended by dA at points X_1 and X_2 .
2. Calculate the irradiance E incident on the plane at points X_1 and X_2 , and calculate the ratio $E(X_1)/E(X_2)$.

1.3 Simple rendering (10 points)

Obtain the file `bunny.mat` from the data folder of this assignment and load it into Matlab. There is a single variable in this file; the variable `N` is an $h \times w \times 3$ array of surface normals. `N(i,j,1)`, `N(i,j,2)`, and `N(i,j,3)` are the x , y , and z components of the surface normal at the ij^{th} surface point, as observed by an orthographic camera with view direction $(0, 0, 1)$.

- Use `quiver` to display the x - y components of the normal field and print the result.

- Compute and display the radiance emitted from each point assuming Lambertian reflectance and constant albedo (equal to one), with a distant point light source in direction $\hat{s} = (0, 0, 1)$, by typing `imshow(N(:, :, 3))`. Explain why this works.
- Compute, display and print the emitted radiance for three different light source directions which are rotated i) 45° up, ii) 45° right, and iii) 75° right from the frontal direction $\hat{s} = (0, 0, 1)$. Can you spot errors in the field of surface normals? What are the illumination effects being ignored in this calculation of scene radiance?




Figure 1: A few bunny renderings.

In your write-up: Please include figures for the normal visualizations and the three renderings of the bunny, as well as answers to the associated questions.

1.4 Photometric stereo (20 points)

The data folder of this assignment contains a set of seven photometric stereo images along with a MAT file containing the light source vectors. You can load the images into Matlab and convert them to double, grayscale format.

1. Estimate the surface normal and grayscale albedo for each pixel. Submit your Matlab code as well as the results (using `imshow` for the albedo values and `quiver` for the surface normals.)
2. Note the poor estimates of the albedo (and surface normal) in the area surrounding the nostrils. What is the source of this error? Describe one method for finding a better estimate of this information from these seven images.
3. Use the recovered surface information to predict what the person would look like (in grayscale) if illuminated from direction $\hat{s} = (0.58, -0.58, -0.58)$ and from direction $\hat{s} = (-0.58, -0.58, -0.58)$. Submit your results and your code.

4. The function `integrate_frankot.m` in the `matlab` folder can be used to recover a surface $z(x, y)$ from your surface normals $\hat{\mathbf{n}}(x, y)$, and the surface can be displayed in Matlab. Display and submit two views of the recovered shape.



Figure 2: Two views of the recovered face shape.

In your write-up: Include figures for the recovered normals and albedo, the re-rendered face images, and the final reconstructed surface, as well as answers to all associated questions.

1.5 Dichromatic reflectance

(30 points)

A reasonable reflectance model for dielectric (non-conducting) surfaces is the so-called dichromatic model, according to which the spectral BRDF is written as a linear combination of a Lambertian diffuse component f_d and a wavelength-independent specular component f_s :

$$f(\lambda, \hat{\omega}_i, \hat{\omega}_o) = f_d(\lambda) + f_s(\hat{\omega}_i, \hat{\omega}_o).$$

When we image a scene consisting of such surfaces under a directional light source with spectral power distribution $I(\lambda)$ and direction $\hat{l} = (l_x, l_y, l_z)$, the RGB values $\vec{C}(\vec{u}) = (C_R(\vec{u}), C_G(\vec{u}), C_B(\vec{u}))$ recorded at pixel $\vec{u} = (u, v)$ can be expressed as

$$\vec{C}(\vec{u}) = \langle \hat{n}(\vec{u}), \hat{l} \rangle \vec{d}(\vec{u}) + g_s(\vec{u}) \vec{s},$$

where $\langle \cdot, \cdot \rangle$ is the inner product operation, $\hat{n}(\vec{u})$ is the surface normal at the scene point imaged by pixel \vec{u} , and $g_s(\vec{u})$ is a function that depends non-linearly on $\hat{n}(\vec{u})$ (as well as view and lighting directions) through the specular component of the BRDF.

1. Assuming that the spectral sensitivities of a camera's three filters are $(c_R(\lambda), c_G(\lambda), c_B(\lambda))$ and that the BRDF at the surface point imaged at pixel \vec{u} is $f(\lambda, \hat{\omega}_i(\vec{u}), \hat{\omega}_o(\vec{u})) = f_d(\lambda, \vec{u}) + f_s(\hat{\omega}_i(\vec{u}), \hat{\omega}_o(\vec{u}))$, write expressions for the elements of the *diffuse color vectors* $\vec{d}(\vec{u})$ and the *source color vector* \vec{s} .
2. Suppose you are given two unit-length three-vectors \hat{r}_1 and \hat{r}_2 that are orthogonal to \vec{s} . Show that the two-channel image¹ given by the per-pixel inner products $\vec{J}(\vec{u}) = (\langle \hat{r}_1, \vec{C}(\vec{u}) \rangle, \langle \hat{r}_2, \vec{C}(\vec{u}) \rangle)$:
 - (a) does not depend on the specular components of the BRDFs, $f_s(\hat{\omega}_i(\vec{u}), \hat{\omega}_o(\vec{u}))$.
 - (b) depends linearly on the surface normals, $\hat{n}(\vec{u})$.
3. Show that the two properties from part (b) are also satisfied by the single-channel (grayscale) image $J(\vec{u}) = \|\vec{J}(\vec{u})\|$.
4. Write a function `imout=makeLambertian(im,s)` that takes an RGB image `im` and a source color vector `s` and computes the grayscale image $J(\vec{u})$ from part (c). Run this function on the image `fruitbowl.png` from the data folder using `s=[0.6257 0.5678 0.5349]` and submit the results. Explain in three sentences or less why `imout` might be more useful than `im` to a computer vision system. Note that a simple way to numerically obtain a set of $N - 1$ unit-length vectors that are orthogonal to N -vector is to use the `null` command in Matlab.



Figure 3: Left: Original color image. Right: Image converted to grayscale with specularities removed.

¹This image can have pixels with negative values, and these negative values are just as useful as the positive ones.

1.6 Color metamers (Extra Credit)

(20 points)

Two distinct spectral distributions are *metameric* if they induce the same responses in the cones of a human retina. Said another way, *metamers* are distinct spectral distributions that map to the same tristimulus vector (in CIE XYZ or any other linear color space). If you own a grocery store that sells bananas, it is in your interest to choose your light source so that overripe bananas look that same as those in their prime. That is, you want the light reflected from ripe and overripe bananas to be metameric.

Let \vec{f} and \vec{g} be $N \times 1$ vectors that come from discretizing the spectral reflectance functions $f(\lambda)$ and $g(\lambda)$ of ripe and overripe bananas, respectively. (For example, if we sampled from 400nm to 700nm in increments of 10nm, we'd have $N = 31$.) Let R be the $3 \times N$ matrix whose rows are obtained by discretizing the CIE XYZ color matching functions.

1. Show that the tristimulus vector $\vec{C}_f = (X_f, Y_f, Z_f)$ for a ripe banana $f(\lambda)$ under illuminant spectrum $l(\lambda)$ can be written as $\vec{C}_f = L_f \vec{l}$ where L_f is a $3 \times N$ matrix and \vec{l} is the discretization of $l(\lambda)$.
2. Given \vec{f} and \vec{g} , a good way to choose \vec{l} is to minimize the distance, in a least-squares sense, of the two resulting tristimulus vectors. Write an expression for the Euclidean distance between the two tristimulus vectors in terms of L_f , its counterpart L_g , and \vec{l} .
3. Write a Matlab function `T=metamericLight(f,g)` that finds the temperature T in the closed interval $T \in [2500\text{K}, 10,000\text{K}]$, with precision $\pm 50\text{K}$, of the blackbody radiator that minimizes the distance derived in part (b). The spectral power distribution for a blackbody radiator (see FP Sect. 6.1.2) with temperature T is

$$L_T(\lambda) \propto \frac{1}{\lambda^5} \frac{1}{e^{\frac{a}{\lambda T}} - 1}$$

where $a = 1.4388 \times 10^{-2}(\text{m} \cdot \text{K})$.

4. The file `bananas.mat` from the data folder contains two spectral reflectance functions, `ripe` and `overripe`, which are sampled in 10nm increments from 400nm to 700nm and correspond roughly to the data from Day 5 and Day 7 of Fig. 4. Use your function to find the optimal metamer-inducing blackbody temperature for these materials, and plot the distance as a function of temperature, and the spectral power distribution of the best illuminant. You will need to use the file `ciexyz64_1.csv` from the data folder, which contains tabulated data for the CIE XYZ color matching functions. In Matlab, you can read data from a CSV file using the `csvread` command. Remember to normalize the spectral power distributions of blackbody radiators in a way that is meaningful for the minimization criterion of part (b).




Figure 4: Relative spectral reflectance curves for different colors of bananas during ripening. Day five is the first day of shelf life. (Taken without permission from Morita et al., “Evaluation of Change in Quality of Ripening Bananas Using Light Reflectance Technique”, in *Memoirs of the Faculty of Agriculture, Kagoshima University*, 1992.)